A Hybrid Representation of Knowledge and Belief

Ricardo Wehbe

Institut für Informatik und angewandte Mathematik, Universität Bern,
Neubrückstrasse 10, CH-3012 Bern, Switzerland. E-mail: wehbe@iam.unibe.ch.

Abstract. Agents working in an environment with incomplete information may need not only knowledge, but also beliefs to supplement their information. Default Logics have been frequently used to represent beliefs. Since inconsistent beliefs give rise to different extensions (scenarios), priorities are introduced to establish a preference among scenarios. We present a hybrid framework in which beliefs are represented by both monotonic and non-monotonic sets of clauses, avoiding thus the need for introducing priorities. Finally we give a glimpse of further work in process and comment briefly on the possibility of updating the database with new information.

1 Introduction

Agents working in an environment with incomplete information may need not only knowledge, but also beliefs to supplement their information. In this paper we present a representation of the knowledge and beliefs of one agent. Knowledge is monotonically represented as a set of clauses (similar to Logic Programming but without NAF and with classical negation) and beliefs as a combination of stablish a preference among scenarios. We present a hybrid framework in which beliefs are represented by both monotonic and non-monotonic sets of clauses, avoiding thus the need for introducing priorities.

The paper is organised as follows. In sections 2 and 3 we define respectively the syntax and semantics of sets of clauses, which represent the knowledge of an agent, and of sets of defaults, which represent its beliefs. In both sections we emphasize a constructive view of the semantics, providing ways to construct the knowledge and belief sets. In section 4 we combine the monotonic approach (i.e., adding clauses to the doxastic part of the database) with the non-monotonic approach. This way, the monotonic part provides a natural way of establishing preferences between extensions without having to use priorities. Section 5 concludes with a comment on further developments in progress.

2 Syntax and Semantics of Sets of Clauses

We assume a set of propositional symbols Π. An *atom* will be either a propositional symbol of Π (a *positive* atom) or the negation of a propositional symbol of Π (a *negative* atom.) A *sequent* will be a (possibly empty) sequence of atoms, which will be denoted by Greek capital letters. Given a set S, the notation $\Gamma \in S$ will be a shorthand to denote that for all atoms p in the sequence Γ, $p \in S$. We will use the notation p, $\overline{p}$ to refer to two complementary atoms when we do not specify which one is positive and which one is negative. Otherwise, we will write simply p, $\neg p$. The set of all atoms that can be formed with the propositional symbols in Π will be denoted by $\mathsf{AT}(\Pi)$. Given a set of propositional symbols Π, a *normal clause* (or a *clause* for short) on Π has the form $\Gamma \Rightarrow p$, where Γ is a sequent and $\Gamma, p \in \mathsf{AT}(\Pi)$.

A clause with an empty sequent will be called a *fact*. We will omit the symbol $\Rightarrow$ when dealing with facts whenever no ambiguity arises.

If φ is a set of clauses, we will denote by Π_φ the set of propositional symbols occurring in φ. The *Herbrand base* of φ, denoted H_φ, is $\mathsf{AT}(\Pi_\varphi)$.

The intended meaning of a clause $p_1, \ldots, p_n \Rightarrow q$ is that if the agent knows $p_1 \wedge \ldots \wedge p_n$, then the agent knows q. Facts are thus trivially true.

Definition 1. *Let Π be a set of propositional symbols and let $S \subseteq \mathsf{AT}(\Pi)$. We say that S is* consistent *iff it does not contain complementary atoms. It is* maximally consistent *with respect to Π iff for each propositional symbol $p \in \Pi$, either $p \in S$ or $\neg p \in S$ but not both.*

It is easy to see that if S is a maximally consistent set with respect to Π, we cannot add an atom $q \in \mathsf{AT}(\Pi)$ without making it inconsistent, unless it is already contained in it. Maximally consistent sets will also be called *worlds*.

Definition 2. *Let φ be a set of clauses and let Π_φ be the set of propositional symbols that occur in φ. A* model *of φ is a set $\mathcal{M} \subset AT(\Pi_\varphi)$ that is maximally consistent with respect to Π_φ such that for any normal clause $\Gamma \Rightarrow q \in \varphi$, if $\Gamma \in \mathcal{M}$, then $q \in \mathcal{M}$.*

A set of clauses does not necessarily have a model. Take, for instance, the set consisting of two complementary facts.

Definition 3. *Let φ be a set of clauses. We say that φ is* consistent *iff it has a model.*

Now we turn to a more "operational" notion of semantics. The following definition introduces the important concept of invariant of a set of clauses.

Definition 4. *Let φ be a consistent set of clauses. The* invariant *of φ, denoted $\mathcal{J}(\varphi)$, is the intersection of all its models.*

The invariant has some important properties. Since it contains the atoms that are true in all models, it may be used to represent the knowledge of an agent. We are interested in finding a way to construct this set.

Definition 5. *Let φ be a set of clauses. An* answer set *of φ is a set $\mathcal{A} \subseteq H_\varphi$ such that for any normal clause $\Gamma \Rightarrow p \in \varphi$, if $\Gamma \in \mathcal{A}$, then $p \in \mathcal{A}$. An answer set is* minimal *if there is no answer set that is a proper subset of it.*

Since sets of clauses may be considered as propositional logic programs with classical negation, it is easy to see that the concept of minimal answer sets corresponds to the minimal model in those programs. Hence, every set of clauses has a minimal answer set (see for instance [5].) The difference is that minimal answer sets may be inconsistent. We will denote the minimal answer set of a set φ of clauses by $\mathcal{A}_\varphi$.

We will be interested only in consistent sets of clauses. We will now relate the notions of answer sets, invariants, and models.

Definition 6. *Let φ be a set of clauses. The operator $T_\varphi : 2^{H_\varphi} \mapsto 2^{H_\varphi}$ is defined as follows: $q \in T_\varphi(S)$ if there is a clause $\Gamma \Rightarrow p \in \varphi$ and $\Gamma \in S$.*

A well-known result of logical program theory is that the T_φ operator has a minimal fixpoint and that this fixpoint is $\mathcal{A}_\varphi$. The proof will be omited. Besides, if we define:

- $T_\varphi^0 = \emptyset$
- $T_\varphi^{k+1} = T_\varphi(T_\varphi^k)$

Then, $\mathcal{A}_\varphi = \bigcup_i T_\varphi^i$, and the fixpoint is reached after a finite number of iterations. For a proof, see [5].

Summing up, every set of clauses has a minimal answer set and this set can be computed within a finite number of steps. Of course, the existence of such minimal answer set does not guarantee the existence of a model. We will establish the relation between the minimal answer set, the invariant and the models.

Lemma 1. *Let φ be a consistent set of clauses. Then for any model $\mathcal{M}$ of φ, $\mathcal{A}_\varphi \subseteq \mathcal{M}$.*

Proof. *Induction on the construction of $\mathcal{A}_\varphi = \mathsf{lfp}(T_\varphi)$.*

Corollary 1. *Let φ be a consistent set of clauses. Then $\mathcal{A}_\varphi \subseteq \mathcal{J}(\varphi)$.*

Proof. *By definition 4, $\mathcal{J}(\varphi)$ is the intersection of all models $\mathcal{M}$ of φ. The result follows from lemma 1.*

The reason of the gap between the minimal answer set and the invariant is that there may be atoms that belong to any model, not because they are the consequent of a clause whose antecedent is true, but because this is the only possibility to avoid inconsistency, as shown next.

Example 1. In the following cases, we have atoms that must be included in all models of φ, although they may not be in $\mathsf{lfp}(T_\varphi)$.

1. If $\overline{p} \Rightarrow p \in \varphi$, then $p \in \mathcal{J}(\varphi)$.
2. If $\overline{p} \Rightarrow a \in \varphi$ and $\overline{p} \Rightarrow \overline{a} \in \varphi$, then $p \in \mathcal{J}(\varphi)$.
3. If $\overline{q} \Rightarrow p \in \varphi$ and $q \Rightarrow p \in \varphi$, then $p \in \mathcal{J}(\varphi)$.

These cases illustrate the three basic cases in which atoms in the invariant may be not in $\mathsf{lfp}(T_\varphi)$. Adding *ad-hoc* rules to T_φ would be of no use, since these cases may appear "hidden" (for instance, the clauses $p \Rightarrow a$ and $\neg a \Rightarrow p$ are an instance of the first case.) In all three cases there occur complementary atoms. We will show that if no complementary atoms occur in φ, then $\mathcal{J}(\varphi) = \mathcal{A}_\varphi$.

Definition 7. *Let φ be a consistent set of clauses. An atom $p \in H_\varphi$ is* bound *in φ if $\varphi \cup \{p\}$ is consistent and $\varphi \cup \{\overline{p}\}$ is inconsistent. A proposition $p \in \Pi_\varphi$ is* free *in φ if both $\varphi \cup \{\neg p\}$ and $\varphi \cup \{p\}$ are consistent.*

Lemma 2. *Let φ be a consistent set of clauses, let $\mathcal{M} \subseteq H_\varphi$. Then $\mathcal{M}$ is a model of φ iff for all $p \in \mathcal{M}$, it is a model of $\varphi \cup \{p\}$.*

Proof. *Omitted.*

Lemma 3. *Let φ be a consistent set of clauses. Then $\mathcal{J}(\varphi)$ contains exactly the bound atoms of φ.*

Proof. *Let us assume that p is bound in φ. Thus $\varphi \cup \{\overline{p}\}$ is inconsistent. Then, by lemma 2 no model may include $\overline{p}$. Thus, $p \in \mathcal{J}_\varphi$.*

Assume now that $p \in \mathcal{J}(\varphi)$. Then, p belongs to all models of φ. Hence, $\varphi \cup \{\overline{p}\}$ is inconsistent, and thus p is bound in φ.

Lemma 4. *Let φ be a set of clauses such that the set of consequents of φ contains no complementary atoms. Then φ is consistent.*

Proof. *Let Π_φ the set of propositional symbols occurring in φ. By hypothesis, the set of consequents of φ is consistent. Any extension of this set to a maximally consistent set with respect to Π_φ is a model of φ.*

Lemma 5. *Let φ be a consistent set of clauses such that no complementary atoms occur in φ. Then, any proposition p such that $p, \neg p \notin \mathcal{A}_\varphi$ is free in φ.*

Proof. *We must show that both $\varphi \cup \{p\}$ and $\varphi \cup \{\neg p\}$ are consistent. Assume $\varphi \cup \{p\}$ is inconsistent. Since by lemma 4 φ is consistent, and $\neg p \notin \mathcal{A}_\varphi$, there hould be a clause whose consequent is $\neg p$ and whose antecedent turns true by the addition of the fact p. But since φ does not contain complementary atoms, it cannot contain such a clause. The same reasoning applies to $\varphi \cup \{\neg p\}$.*

Proposition 1. *Let φ be a set of clauses with no occurrence of complementary atoms. Then $\mathcal{A}_\varphi = \mathcal{J}(\varphi)$.*

Proof. *By lemma 5, if φ does not contain complementary atoms, all bound atoms in φ are in $\mathcal{A}_\varphi$. The result follows from corollary 1.*

The restriction of sets of clauses by not allowing complementary clauses is too strong. We will show a constructive way to compute the invariant set of clauses even when it contains complementary atoms. Some new definitions will be needed first.

Definition 8. *Let φ be a consistent set of clauses and let $S \subseteq H_\varphi$. The* reduction *of φ with respect to S, denoted by $\mathcal{R}_S(\varphi)$, is the set of clauses constructed from φ as follows:*

1. *Eliminate all clauses whose consequents are in S.*
2. *Eliminate all clauses containing in their antecedents atoms whose complements are in S.*
3. *Eliminate all other occurrences of atoms in S.*
4. *Rewrite all clauses $\Gamma, \overline{q} \Rightarrow \overline{a}$ where $a \in \mathcal{A}_\varphi$ as $\Gamma \Rightarrow q$.*

Intuitively, if S contains atoms which are taken to be true in a set of clauses, the reduction is the set of the clauses whose consequents are still unsolved.

Lemma 6. *Let φ be a consistent set of clauses and let $\mathcal{M}$ be a model of φ. Then $\mathcal{M}$ is a model of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$.*

Proof. *A clause of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ is either:*

1. *A clause $\Gamma \Rightarrow p$ of φ. Then, if $\Gamma \in \mathcal{M}$, then $p \in \mathcal{M}$.*
2. *A clause $\Gamma' \Rightarrow p$ which is obtained from a clause $\Gamma \Rightarrow p$ of φ by eliminating from Γ all atoms that belong to $\mathcal{A}_\varphi$. Then, since $\mathcal{A}_\varphi \subseteq \mathcal{M}$, we have eliminated atoms that are true in Γ. Thus, if $\mathcal{M}$ is a model of the original clause, it is a model of the modified one.*
3. *A clause $\Gamma \Rightarrow p$ which corresponds to a clause $\Gamma, \overline{p} \Rightarrow \overline{a}$ of φ. Since $a \in \mathcal{A}_\varphi$, either $\Gamma \notin \mathcal{M}$ or $p \in \mathcal{M}$. In both cases, $\mathcal{M}$ is a model of the modified normal clause.*
4. *A normal clause $\Gamma, \Sigma \Rightarrow p$ which corresponds to a clause $\Gamma, \overline{p}, \Sigma \Rightarrow p$ in φ. It is immediate that if $\mathcal{M}$ is a model of the original normal clause, it is a model of the modified one.*

Lemma 7. *Let φ be a consistent set of clauses. Then, any model of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ which includes $\mathcal{A}_\varphi$ is a model of φ.*

Proof. *Let $\mathcal{M}$ be a model of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ such that $\mathcal{A}_\varphi \subseteq \mathcal{M}$. Let us consider the following cases:*

1. *Let $\Gamma \Rightarrow p$ be a clause of φ that was eliminated using rule 1. Then, since $p \in \mathcal{A}_\varphi$, $\mathcal{M}$ is a model of the clause.*
2. *Let $\Gamma \Rightarrow p$ be a clause of φ that was eliminated using rule 2. Then, since there is at least some $\overline{p}$ in Γ such that $p \in \mathcal{A}_\varphi$, $\mathcal{M}$ is a model of the clause.*
3. *Let $\Gamma \Rightarrow p$ be a clause of φ that was rewritten using rule 3. Then, since the rewritten clause is stronger than the original one, if $\mathcal{M}$ is a model of the former it is a model of the latter.*
4. *Let $\Gamma, \overline{p} \Rightarrow \overline{a}$ be a clause of φ that was rewritten using rule 4 yielding $\Gamma \Rightarrow p$. Then, if $\Gamma \in \mathcal{M}$ then $p \in \mathcal{M}$. Since $a \in \mathcal{A}_\varphi$, it follows that $\mathcal{M}$ is a model of the original clause.*

Corollary 2. *Let φ be a consistent set of clauses. If p is bound in φ, then either $p \in \mathcal{A}_\varphi$ or p is bound in $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$.*

Proof. *From lemmas 6 and 7, it is immediate that a set $\mathcal{M}$ is a model of φ iff it is a model of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ that includes $\mathcal{A}_\varphi$. Thus, if p is bound in φ there are no models of φ which include $\overline{p}$. Thus there are no models of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ which include $\overline{p}$.*

The preceding result is important because it allows the construction of the invariant of a set of clauses.

Definition 9. *Let φ be a set of clauses. We define the operator $\mathcal{T}_\varphi = \bigcup_k \mathcal{T}_\varphi^k$, where*

$$\mathcal{T}_\varphi^0 = \emptyset \qquad \varphi_0 = \varphi$$

$$\mathcal{T}_\varphi^{k+1} = \mathcal{A}_{\varphi_k} \; \varphi_{k+1} = \mathcal{R}_{\mathcal{T}_\varphi^{k+1}}(\varphi_k)$$

The operator $\mathcal{T}_\varphi$ is the union of least fixpoints, since each term $\mathcal{T}_\varphi^k$ is the least fixpoint of T_{φ_k}. The following proposition shows that the union has only finitely many elements.

Proposition 2. *Let φ be a set of clauses. Then there is a finite natural number m such that $m > n$ implies $\mathcal{T}_\varphi^{m+1} = \mathcal{T}_\varphi^m$.*

Proof. *The proof (here omitted) is based on the fact that the reduction of a set of clauses has either less normal clauses or less atoms than the original set. Since we have always less atoms, the process cannot go on forever.*

The problem is still what to do with the complementary atoms. It is possible that the $\mathcal{T}$ operator does not find all bound atoms, as we saw in example 1. The strategy will be based on successive "splittings" of the set of clauses. For each pair of complementary atoms $p, \overline{p}$ appearing in a set of clauses φ, two separate sets will be considered, $\varphi \cup \{\Rightarrow p\}$ and $\varphi \cup \{\Rightarrow \overline{p}\}$. Some previous results will be needed. The following lemma is a straightforward extension of previous results.

Lemma 8. *Let φ be a consistent set of clauses. Then:*

1. *If p is bound in φ, either $p \in \mathcal{T}_\varphi$ or p is bound in $\mathcal{R}_{\mathcal{T}_\varphi}$.*
2. *Any model of $\mathcal{R}_{\mathcal{T}_\varphi}$ that includes $\mathcal{T}_\varphi$ is a model of φ.*

Proof. *(Part 1) Induction on the construction of $\mathcal{T}_\varphi$.*
Base case: immediate, since $\mathcal{T}_\varphi^0 = \emptyset$ and $\varphi_0 = \varphi$.
Induction step: $\mathcal{T}_\varphi^{k+1} = \mathcal{A}_{\varphi_k}$, and $\varphi_{k+1} = \mathcal{R}_{\mathcal{A}_{\varphi_k}}$. By corollary 2, we have that if p is bound in φ_k, then either $p \in \mathcal{T}_\varphi^{k+1}$ or p is bound in φ_{k+1}.
(Part 2). Induction on the construction of $\mathcal{T}_\varphi$.
Base case: immediate, since $\mathcal{T}_\varphi^0 = \emptyset$ and $\varphi_0 = \varphi$.
Induction step: $\mathcal{T}_\varphi^{k+1} = \mathcal{A}_{\varphi_k}$, and $\varphi_{k+1} = \mathcal{R}_{\mathcal{A}_{\varphi_k}}$. By lemma 7, any model of φ_{k+1} that includes $\mathcal{T}_\varphi^{k+1}$ is a model of φ_k.

Proposition 3. *Let φ be a set of clauses. Then, if $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ contains no complementary atoms, either $\mathcal{J}(\varphi) = \mathcal{T}_\varphi$ or $\mathcal{T}_\varphi$ is inconsistent.*

Proof. *Assume first that φ is consistent. Observe that $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ contains no facts, because $\mathcal{T}_\varphi$ has reached a fixpoint. If it does not contain complementary atoms, we may construct two models of $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$, one setting all atoms to true and the other one setting all atoms to false. Thus, for any propositional symbol p occurring in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$, both $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{p\}$ and $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{\neg p\}$ are consistent. By lemma 8, any atom that is bound in φ is either in $\mathcal{T}_\varphi$ or is bound in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ If all atoms in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ are free, then $p \in \mathcal{T}_\varphi$. Thus, $\mathcal{J}(\varphi) = \mathcal{T}_\varphi$.*

Assume now that φ is inconsistent. Since $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ contains no complementary atoms, it is consistent. By lemma 8, any model of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ that includes $\mathcal{T}_\varphi$ is a model of φ. Thus, $\mathcal{T}_\varphi$ must be inconsistent.

Corollary 3. *Let φ be a consistent set of clauses.*

1. *An atom p is bound in φ iff for any atom q that is free in φ, p is bound in $\varphi \cup \{q\}$ and in $\varphi \cup \{\overline{q}\}$.*
2. *An atom p is bound in φ iff for any atom q that is bound in φ, then p is bound in $\varphi \cup \{q\}$.*

Proof. *(Part 1) Let p be bound in φ. Then all models $\mathcal{M}$ of φ include p. Since q is free, there are models including q and models including $\overline{q}$. By lemma 2, the former are the models of $\varphi \cup \{q\}$ and the latter are the models of $\varphi \cup \{\overline{q}\}$. Since all of them include p, it is bound in both sets of clauses. Now let p be bound in $\varphi \cup \{\Rightarrow q\}$ and in $\varphi \cup \{\Rightarrow \overline{q}\}$. Then all models of φ that include q include p and all models of φ that include $\overline{q}$ include p. Thus, p is bound in φ.*
(Part 2) Immediate from lemma 2.

Corollary 4. *Let φ be an inconsistent set of clauses. Then either $\mathcal{A}_\varphi$ contains complementary atoms or $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ is inconsistent.*

Proof. *First observe that the atoms in $\mathcal{A}_\varphi$ do not occur in $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$. Thus, they should be free therein. If φ is inconsistent, then there are no models of $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ containing $\mathcal{A}_\varphi$. Thus, all the complements of the atoms in $\mathcal{A}_\varphi$ should be bound in $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$, contradicting the fact that they are free.*

Corollary 5. *Let φ be an inconsistent set of clauses such that $\mathcal{R}_{\mathcal{A}_\varphi}(\varphi)$ contains no complementary atoms. Then $\mathcal{A}_\varphi$ contains complementary atoms.*

Proof. *Immediate from corollary 4.*

We are ready to provide a procedure for the construction of the invariant of a set of clauses. In the next definitions we assume some ordering in the Herbrand base of a set of clauses.

Definition 10. *Let φ be a set of clauses. We put $\varphi_0 = \varphi$ and define:*

- *$\mathcal{I}(\varphi) = \mathcal{T}_\varphi$ if $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ has no complementary atoms.*
- *$\mathcal{I}(\varphi) = \mathcal{T}_\varphi \cup (\mathcal{I}(\varphi_1) \cap \mathcal{I}(\varphi_2))$, where $\varphi_1 = \mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{p\}$, $\varphi_2 = \mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{\neg p\}$ and $p, \neg p$ is the first pair of complementary atoms appearing in $\mathcal{R}_{\mathcal{J}_\varphi}(\varphi)$.*

Note that the process is finite, since we eliminate one pair of complementary atoms in each step.

Proposition 4. *Let φ be a set of clauses. Then $\mathcal{J}(\varphi) = \mathcal{I}(\varphi)$.*

Proof. *Induction on the structure of $\mathcal{J}(\varphi)$.*
Base case: if $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ has no complementary atoms, then by lemma 3 $\mathcal{J}(\varphi) = \mathcal{I}(\varphi)$.
Induction step: if $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ has complementary atoms p, $\neg p$, then by lemma 8, $\mathcal{J}(\varphi) = \mathcal{T}_\varphi \cup \mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi))$. Let us suppose that p is bound in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$, and let q, $\neg q$ be two complementary atoms occurring in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$. Then, either one of them is bound in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$ or both are free. Suppose q is bound. Then by corollary 3, part 2, we have that $p \in \mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{q\})$, and that $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{\neg q\}$ is inconsistent. Thus, by induction hypothesis $\mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)) = \mathcal{I}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{q\})$ and $\mathcal{J}(\mathcal{R}_{\mathcal{J}_\varphi}(\varphi) \cup \{\neg q\}) = H_\varphi$. Hence, $\mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)) = \mathcal{I}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{q\}) \cap \mathcal{I}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{\neg q\})$.

Assume now that both q and $\neg q$ are free in $\mathcal{R}_{\mathcal{T}_\varphi}(\varphi)$. By corollary 3, part one, we have that $p \in \mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi))$ implies $p \in \mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{q\})$ and $p \in \mathcal{J}(\mathcal{R}_{\mathcal{T}_\varphi}(\varphi) \cup \{\neg q\})$. The result follows by induction hypothesis.

Example 2. Let φ be

$$
\begin{aligned}
q &\Rightarrow p \\
q &\Rightarrow \neg p \\
\neg q, s, t &\Rightarrow r \\
&\Rightarrow s \\
&\Rightarrow t
\end{aligned}
$$

And assume an ordering $\{p, \neg p, q, \neg q, r, \neg r, s, \neg s, t, \neg t\}$ in the Herbrand base. We have:

$$\mathcal{J}(\varphi) = \{s, t\} \cup (\{p, \neg q, r\} \cap \{\neg p, \neg q, r\}) = \{\neg q, r, s, t\}$$

Finally we mention some properties of invariants that will be useful later.

Proposition 5. *Let φ_1 and φ_2 be two consistent sets of clauses such that $\varphi_1 \subseteq \varphi_2$. Then $\mathcal{J}(\varphi_1) \subseteq \mathcal{J}(\varphi_2)$.*

Proof. *Let $\mathcal{M}$ be a of φ_2. Then $\mathcal{M} \backslash H_{\varphi_1}$ is a model of φ_1. Thus for all p, $p \in \mathcal{M}$ and $p \in H_{\varphi_1}$ implies $p \in \mathcal{J}(\varphi_1)$. Thus, $\mathcal{J}(\varphi_1) \subseteq \mathcal{J}(\varphi_1)$.*

Proposition 6. *Let φ be a set of clauses. Then $\mathcal{J}(\varphi) = \mathcal{J}_\varphi(E \cup \mathcal{J}(\varphi))$.*

Proof. *Let $\mathcal{M}$ be a model of $\varphi \cup \mathcal{J}(\varphi)$. Then, it is a model of φ. Thus, the intersection of all models of $\varphi \cup \mathcal{J}(\varphi))$ is the intersection of all models of φ.*

Corollary 6. *Let φ_1 and φ_2 be two sets of clauses. Then $\mathcal{J}(\varphi_1 \cup \mathcal{J}(\varphi_1 \cup \varphi_2)) = \mathcal{J}(\varphi_1 \cup \varphi_2)$.*

Proof. *Take $\varphi = \varphi_1 \cup \varphi_2$. The result follows from lemmas 6 and 5.*

3 A Non-Monotonic Representation of Belief

We will use clauses to represent knowledge. The invariant will represent the knowledge of an agent, in the sense that it consists of the atoms which are true in all possible worlds. We will distinguish *knowledge* from *belief* in the sense that knowledge is true, whereas belief may be false. This assumption, although usual, is rather strong: we have problems when several agents are concerned, since private communication might lead to inconsistency.

We will represent beliefs by *defaults*, which we define next.

Definition 11. *Given a set of propositional symbols Π, a* normal default *(or a* default *for short) on Π has the form $\Gamma : p \Rightarrow p$ where Γ is a sequent and $\Gamma, p \in \mathsf{AT}(\Pi)$. The sequent Γ is the* prerequisite *of the default and p is the* justification *of the default (left side) and the* consequent *of the default (right side.)*

If D is a set of defaults, then we will denote by $\mathsf{CONS}(D)$ the set of consequents of D.

Strictly speaking, we will use a subset of defaults, namely the so-called *normal defaults* [6], [3]. In normal defaults, the consequent is the justification; general defaults allow arbitrary sequents as justifications. Since we will use only normal defaults, we call them simply "defaults."

The intended meaning of a default $\Gamma : p \Rightarrow p$ is that if Γ is true and p is not inconsistent with the knowledge the agent has, then it will be taken to be true. Of course, it will not be in the same level as the atoms of the invariant; the latter will be *known*; the former will just be *believed*. The use of defaults implies that the representation is no longer monotonic.

Definition 12. *Let Π be a set of propositions. A* knowledge and belief database *(*KB-database *for short) is a pair $\zeta = (E, D)$ where E is a set of clauses and D is a set of defaults on Π. The set of clauses E is the* epistemic part *of the database and the set D is the* doxastic part *of the database. A KB-database where E is consistent will be said to be* epistemically consistent*.*

As before, given a KB-database ζ, we will denote by Π_ζ the set of propositional symbols occurring in the epistemic and the doxastic parts of ζ. The Herbrand base will be defined in the same way as for sets of clauses: the *Herbrand base* of ζ, denoted by H_ζ, is $\mathsf{AT}(\Pi_\zeta)$.

As usual, semantics of sets of defaults will be based on the concept of *extensions*. Informally speaking, an extension is the set of beliefs we may form starting from a KB-Database. We point out that when we apply defaults, we assume some external circumstances, expressed as a set of atoms. This will be called a *context*. Now we define the belief sets we may form in the presence of a context. This will lead to the formal definition of extensions.

Definition 13. *Let $\zeta = (E, D)$ be a KB-database and let $S \subseteq H_\zeta$. Then $\Lambda_\zeta(S)$ is defined as the smallest set such that the following properties are fulfilled:*

1. $\mathcal{J}(E) \subseteq \Lambda_\zeta(S)$
2. *If* $p \in \mathcal{J}(E \cup \Lambda_\zeta(S))$ *then* $p \in \Lambda_\zeta(S)$
3. *If* $\Gamma : p \Rightarrow p \in D$ *and* $\Gamma \in \Lambda_\zeta(S)$ *and* $\overline{p} \notin S$*, then* $p \in \Lambda_\zeta(S)$

Informally, $\Lambda_\zeta(S)$ is the minimal set of beliefs that an agent whose KB-database is ζ may have in view of the context S. The set $\Lambda_\zeta(S)$ may be inconsistent, as the following example shows.

Example 3. Let $\zeta = (\emptyset, \{: p \Rightarrow p, : \neg p \Rightarrow \neg p\})$. Then, $p, \neg p \in \Lambda_\zeta(\emptyset)$.

A natural way to avoid this problem is to take the fixpoint.

Definition 14. *Let* $\zeta = (E, D)$ *be a KB-Database. A set* $S \subseteq H_\zeta$ *is an* extension *for* ζ *iff* S *is a fixpoint of* Λ_ζ*, i.e.,* $S = \Lambda_\zeta(S)$*.*

The concept of extensions may seem elusive. We will give several characterisations of it.

Proposition 7. *Let* $\zeta = (E, D)$ *be a KB-database and let* S *be an extension for* ζ*. Then* S *is inconsistent iff* E *is inconsistent.*

Proof. It is trivial that if E *is inconsistent, then* S *is inconsistent for the first condition of the definition of* $\Lambda_\zeta(S)$*. First we show that if* S *is inconsistent, then* $S \subseteq \mathcal{J}(E)$*. It is clear that the first and the second conditions ot the definition of* $\Lambda_\zeta(S)$ *are fulfilled by* $\mathcal{J}(E)$*, since* $\mathcal{J}(E) = \mathcal{J}(E \cup \mathcal{J}(E))$*. Besides, the third condition is also fulfilled, since all justifications of defaults belong to* S*. Thus,* $\Lambda_\zeta(S) \subseteq \mathcal{J}(E)$ *by the minimality of* $\Lambda_\zeta(S)$*. Since* S *is inconsistent, so must be* $\mathcal{J}(E)$*.*

Corollary 7. *Let* ζ *be a KB-database. If* S *has an inconsistent extension, then it has no other extension.*

Proof. Since all extensions for ζ *include* $\mathcal{J}$ *and this invariant is inconsistent by proposition 7, then all extensions must also be inconsistent.*

Proposition 7 states that the addition of defaults may not turn a KB-database inconsistent, unless it is epistemically inconsistent.

Now we relate the notion of extensions to that of invariants. We need some more notation first.

Definition 15. *Let* $\zeta = (E, D)$ *be a KB-database and let* S *be an extension for* ζ*. Then the set of* generating defaults *for* S *in* ζ*, denoted by* $\mathsf{GD}_\zeta(S)$*, is the set* $\mathsf{GD}_\zeta(S) = \{\Gamma : p \Rightarrow p \in D \mid \Gamma \in S \text{ and } \overline{p} \notin S\}$

Lemma 9. *Let* $\zeta = (E, D)$ *be a KB-database and let* S *be an extension for* ζ*. Then* $S = \mathcal{J}(E \cup \mathsf{CONS}(\mathsf{GD}_\zeta(S)))$*.*

Proof. *Let $p \in \mathsf{CONS}(\mathsf{GD}_\zeta(S))$. Then, there is a default $\mathit{\Gamma} : p \Rightarrow p \in D$ such that $\mathit{\Gamma} \in S$ and $\overline{p} \notin S$. Thus, $p \in S$. Therefore, $\mathsf{CONS}(\mathsf{GD}_\zeta(S)) \subseteq S = \mathit{\Lambda}_\zeta(S)$ and thus $\mathcal{J}(E \cup \mathsf{CONS}(\mathsf{GD}_\zeta(S))) \subseteq S$.*

Now let $\mathit{\Phi}$ be an abbreviation for $\mathsf{CONS}(\mathsf{GD}_\zeta(S))$. We will show that $\mathcal{J}(E \cup \mathit{\Phi})$ includes $\mathit{\Lambda}_\zeta(S)$.

On the one hand, we have by proposition 5 that $\mathcal{J}(E) \subseteq \mathcal{J}(E \cup \mathit{\Phi})$. On the other hand, by proposition 6, $\mathcal{J}(E \cup \mathcal{J}(E \cup \mathit{\Phi})) = \mathcal{J}(E \cup \mathit{\Phi})$.

Assume now that there is a default $\mathit{\Gamma} : p \Rightarrow p \in D$ such that $\mathit{\Gamma} \in \mathcal{J}(E \cup \mathit{\Phi})$ and $\overline{p} \notin S$. Since $\mathcal{J}(E \cup \mathit{\Phi}) \subseteq S$, then $\mathit{\Gamma} \in S$. Thus, $\mathit{\Gamma} : p \Rightarrow p \in \mathsf{GD}_\zeta(S)$. Therefore, $p \in \mathcal{J}(E \cup \mathit{\Phi})$. Thus all three conditions of the definition of $\mathit{\Lambda}_\zeta(S)$ are fulfilled and $\mathit{\Lambda}_\zeta(S) \subseteq \mathcal{J}(E \cup \mathit{\Phi})$. The result follows immediately.

The following lemma states an important property of defaults.

Lemma 10. *Let $\zeta_1 = (E, D_1)$ and $\zeta_2 = (E, D_2)$ be two KB-databases such that $D_1 \subseteq D_2$ and let S_1 be an extension for ζ_1. Then there is an extension S_2 for ζ_2 such that $S_1 \subseteq S_2$.*

Proof. *We construct first the sequence of sets of defaults $\mathit{\Delta}_0, \ldots, \mathit{\Delta}_m$ such that:*

- $\mathit{\Delta}_0 = \mathsf{GD}_{\zeta_1}(S_1)$
- $\mathit{\Delta}_{k+1} = \{\mathit{\Gamma} : p \Rightarrow p \in D_2 \mid \mathit{\Gamma} \in \bigcup_{i=0}^{k}(\mathsf{CONS}(\mathit{\Delta}_i))$ and $\overline{p} \notin \bigcup_{i=0}^{k+1}(\mathsf{CONS}(\mathit{\Delta}_i))\}$

The sequence is finite, since the set of defaults is finite. Now let us define the abbreviation $\mathit{\Sigma} = \mathcal{J}(E \cup (\bigcup_i \mathsf{CONS}(\mathit{\Delta}_i)))$. We show first that $\mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma}) \subseteq \mathit{\Sigma}$.

It is clear that $\mathcal{J}(E) \subseteq \mathit{\Sigma}$. Besides, by proposition 6, $\mathcal{J}(E \cup \mathit{\Sigma}) = \mathcal{J}(\mathit{\Sigma})$.

Consider now a default $\mathit{\Gamma} : p \Rightarrow p \in D_2$, with $\mathit{\Gamma} \in \mathit{\Sigma}$ and $\overline{p} \notin \mathit{\Sigma}$. Then there is some k such that $\mathit{\Gamma} \in \mathcal{J}(E \cup (\bigcup_{i=0}^{k} \mathsf{CONS}(\mathit{\Delta}_i)))$, and $\overline{p} \notin \mathcal{J}(E \cup (\bigcup_{i=0}^{k+1} \mathsf{CONS}(\mathit{\Delta}_i)))$. Hence, $\mathit{\Gamma} : p \Rightarrow p \in \mathit{\Delta}_{k+1}$ and $p \in \mathit{\Sigma}$. Thus, by the minimality of $\mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma})$ we get $\mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma}) \subseteq \mathit{\Sigma}$.

Now assume that $\mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma}) \neq \mathit{\Sigma}$. Then there is some k such that $\mathcal{J}(E \cup (\bigcup_{i=0}^{k} \mathsf{CONS}(\mathit{\Delta}_i))) \subseteq \mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma})$, even though $\mathcal{J}(E \cup (\bigcup_{i=0}^{k+1} \mathsf{CONS}(\mathit{\Delta}_i))) \not\subseteq \mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma})$. Thus there is a default $\mathit{\Gamma} : p \Rightarrow p \in D_2$ such that $p \in \mathit{\Sigma}$ and $p \notin \mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma})$. Hence $\mathit{\Gamma} \in \mathcal{J}(E \cup (\bigcup_{i=0}^{k} \mathsf{CONS}(\mathit{\Delta}_i)))$ and $\overline{p} \notin \mathcal{J}(E \cup (\bigcup_{i=0}^{k+1} \mathsf{CONS}(\mathit{\Delta}_i)))$. But then, $\mathit{\Gamma} \in \mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma})$. Besides, $\overline{p} \notin \mathit{\Sigma}$ and thus p must be in $\mathit{\Lambda}_{\zeta_2}(\mathit{\Sigma})$, which contradicts the hypothesis.

Finally, we have that by construction $\mathcal{J}(E \cup \mathsf{CONS}(\mathsf{GD}_{\zeta_1}(S_1)) \subseteq \mathit{\Sigma}$ and by lemma 9 we get the result.

The last lemma shows a property that is sometimes called "semimonotonicity" [3]. It is also important because it leads to the following result.

Corollary 8. *Let $\zeta = (E, D)$ be a KB-database and let E be a consistent set of clauses. Then ζ has an extension.*

Proof. *Since E is consistent, then clearly the KB-Database $(E, \emptyset)$ has an extension, namely $\mathcal{J}$. Then by lemma 10, it follows that ζ has an extension.*

We have seen that a KB-database that is epistemically consistent has at least one extension and in general a family of extensions. Now we will characterise the families of extensions that a KB-database may have.

Lemma 11. *Let $\zeta = (E, D)$ be a KB-database and let us define:*

- $S^0 = \mathcal{J}(E)$
- $S^{k+1} = \mathcal{J}(E \cup \{\mathsf{CONS}(\Gamma : p \Rightarrow p \in D) \mid \Gamma \in S^k \text{ and } \overline{p} \notin S\}$

Then S is an extension for ζ iff $S = \bigcup_i (S^i)$.

Proof. *First we show that $\Lambda_\zeta(S) \subseteq \bigcup_i (S^i)$. We have that $\mathcal{J}(E) \subseteq \bigcup_i (S^i)$ and by proposition 6 we have that $\mathcal{J}(E \cup \bigcup_i (S^i)) = \mathcal{J}(\bigcup_i (S^i))$. Now let us consider a default $\Gamma : p \Rightarrow p \in D$ such that $\Gamma \in \bigcup_i (S^i)$ and $\overline{p} \notin \bigcup_i (S^i)$. Then $p \in \bigcup_i (S^i)$.*

Now we show that $\bigcup_i (S^i) \subseteq \Lambda_\zeta(S)$. Assume there is some $p \in \bigcup_i (S^i)$ such that $p \notin \Lambda_\zeta(S)$. Then there is some k such that $S^k \subseteq \Lambda_\zeta(S)$ although $S^{k+1} \not\subseteq \Lambda_\zeta(S)$. Thus, there must be a default $\Gamma : p \Rightarrow p \in D$ such that $\Gamma \in S^k$, $\overline{p} \notin S$ and $p \notin \Lambda_\zeta(S)$. Since $S^k \subseteq \Lambda_\zeta(S)$, then p must be in $\Lambda_\zeta(S)$ contradicting thus the hypothesis.

Lemma 12. *Let $\zeta = (E, D)$ be a KB-database and let S_1 and S_2 be two extensions for ζ such that $S_1 \subseteq S_2$. Then $S_1 = S_2$.*

Proof. *By lemma 11 we have: $S_1 = \bigcup_i (S_1^i)$ and $S_2 = \bigcup_i (S_2^i)$. The lemma is proved by induction on k.*
Base case: $S_1^0 = S_2^0 = \mathcal{J}(E)$
Induction step: $S_2^{k+1} = \mathcal{J}(E \cup \{\mathsf{CONS}(\Gamma : p \Rightarrow p \in D) \mid \Gamma \in S_1^k \text{ and } \overline{p} \notin S_2\})$.
By induction hypothesis, $S_1^k = S_2^k$ and, since $S_1 \subseteq S_2$, then $\overline{p} \notin S_2$ implies $\overline{p} \notin S_1$. Hence $S_1^{k+1} = S_2^{k+1}$.

Lemma 13. *Let $\zeta = (E, D)$ be a KB-database and let S_1, S_2 be two distinct extensions for ζ. Then $S_1 \cup S_2$ is inconsistent.*

Proof. *If S_1 and S_2 are extensions, we have by lemma 10 that $S_1 = \bigcup_i S_1^i$ and $S_2 = \bigcup_i S_2^i$. We have also by lemma 9 that $S_1 = \mathcal{J}(E \cup \mathsf{CONS}(\mathsf{GD}_\zeta(S_1)))$ and $S_2 = \mathcal{J}(E \cup \mathsf{CONS}(\mathsf{GD}_\zeta(S_2)))$. f we take $\Delta = \mathsf{GD}_\zeta(S_1) \cap \mathsf{GD}_\zeta(S_2)$, then we have that $\mathsf{GD}_\zeta(S_1) \backslash \Delta \neq \emptyset$ and $\mathsf{GD}_\zeta(S_2) \backslash \Delta \neq \emptyset$, since otherwise we would be in the case of lemma 12. Let us suppose that there is an ordering of the defaults of D and that the first k defaults are those in Δ. Let thus δ_{k+1} be the first default such that $\delta_{k+1} \in \mathsf{GD}_\zeta(S_1) \backslash \Delta$. Then we have that $S_1^k = S_2^k$ and $S_1^{k+1} = \mathcal{J}(E \cup \{\mathsf{CONS}(\Gamma : p \Rightarrow p \in D) \mid \Gamma \in S_1^k \text{ and } \overline{p} \notin S_1\})$.*

Since $\Gamma \in S_1^k$, $\Gamma \in S_2^k$ by the construction of the sequence. Hence, if $p \notin S_2$, the only possibility is that $\overline{p} \in S_2$. Thus $S_1 \cup S_2$ contains complementary atoms.

The process to construct the extensions of a KB-database $\zeta = (E, D)$ will be a "naïve" one, similar to one the proposed in [6]. Starting from the invariant $\mathcal{J}(E)$, we apply one default each time until no further defaults are applicable. This process terminates, since the set of defaults is finite. The question is, whether this process is sound (i.e., if it yields an extension) and whether it is complete (i.e., if all extensions can be obtained this way.)

Definition 16. *Let $\zeta = (E, D)$ be a KB-database. Then we define the operator $\chi(\zeta) = \bigcup_i \chi^i(\zeta)$ as follows:*

– $\chi^0(\zeta) = \mathcal{J}(E)$

– $$\chi^{k+1}(\zeta) = \begin{cases} \mathcal{J}(E \cup \chi^k(\zeta) \cup \{p\}) & \text{if there is a default } \Gamma : p \Rightarrow p \in D \\ & \text{such that } \Gamma \in \mathcal{J}\chi^k(\zeta) \\ & \text{and } p \text{ is free in } E \cup \chi^k(\zeta) \\ \\ \chi^k(\zeta) & \text{otherwise} \end{cases}$$

Note that this definition describes a family of sets rather than one.

Lemma 14. *Let $\zeta = (E, D)$ be a KB-Database. Then if E is consistent, so is $\chi(\zeta)$.*

Proof. *Induction on the structure of $\chi(\zeta)$.*
Base case: immediate, since $\chi^0(\zeta) = \mathcal{J}(E)$ and E is consistent by hypothesis.
Induction step: $\chi^{k+1}(\zeta) = \mathcal{J}(E \cup \chi^k(\zeta) \cup \{p\})$ where there is a default $\Gamma : p \Rightarrow p \in D$ such that $\Gamma \in \chi^k(\zeta)$ and p is free in $E \cup \chi^k(\zeta)$. By induction hypothesis, $\chi^k(\zeta)$ is consistent and by proposition 6, so is $E \cup \chi(\zeta)$. Thus, since p is free in $E \cup \chi^k(\zeta)$, then $E \cup \chi^k(\zeta) \cup \{p\}$ is consistent.

Proposition 8. *Let $\zeta = (E, D)$ be a KB-Database and let E be consistent. Then $\chi(\zeta)$ is an extension for ζ.*

Proof. *We prove first that $\Lambda_\zeta(\chi(\zeta)) \subseteq \chi(\zeta)$. We have that $\mathcal{J}(E) \subseteq \chi(\zeta)$. Besides, by proposition 6 we have that $\mathcal{J}(E \cup \chi(\zeta)) = \chi(\zeta)$. Now consider a default $\Gamma : p \Rightarrow p \in D$ such that $\Gamma \in \chi(\zeta)$ and $\overline{p} \notin \chi(\zeta)$. Then there is some k such that $\Gamma \subseteq \chi^k(\zeta)$ and also $\overline{p} \notin \chi^k(\zeta)$. Thus, $p \in \chi^{k+1}(\zeta) \subseteq \chi(\zeta)$. Hence, $\Lambda_\zeta(\chi(\zeta)) \subseteq \chi(\zeta)$.*

Assume now that $\chi(\zeta) \nsubseteq \Lambda_\zeta(\chi(\zeta))$. Since $\chi^0(\zeta) \subseteq \Lambda_\zeta(\chi(\zeta))$ and for all $j \leq 0$ $\chi^j(\zeta) \subseteq \chi^{j+1}(\zeta)$, we have that there must be some k such that $\chi^k(\zeta) \subseteq \Lambda_\zeta(\chi(\zeta))$ even though $\chi^{k+1}(\zeta) \nsubseteq \Lambda_\zeta(\chi(\zeta))$. Then there is some default $\Gamma : p \Rightarrow p \in D$ such that $\Gamma \in \chi^k(\zeta)$ and p is free in $\chi^k(\zeta)$. But then, since $p \notin \Lambda_\zeta(\chi(\zeta))$, then $\overline{p} \in \Lambda_\zeta(\chi(\zeta))$ and since $\Lambda_\zeta(\chi(\zeta)) \subseteq \chi(\zeta))$, then $\overline{p} \in \chi(\zeta)$. Hence, $\chi(\zeta)$ would be inconsistent, contradicting thus lemma 14.

We have thus that the application of the operator $\chi(\zeta)$ indeed yields an extension for ζ. The question is now whether all extensions may be found with some application of this operator.

Proposition 9. *Let $\zeta = (E, D)$ be a KB-Database. Then for any extension S for ζ, there is a set $\chi(\zeta)$ such that $\chi(\zeta) = S$.*

Proof. *Let S be an extension for ζ. We define $\chi(\zeta)$ such that the choice of the defaults is restricted to defaults whose consequents are in S. We show that $\chi(\zeta) \subseteq S$ by induction on the construction of $\chi(\zeta)$.*
Base case: $\chi^0(\zeta) = \mathcal{J}(E) \subseteq S$.

Induction step: let $q \in \chi^{k+1}(\zeta)$. Then $q \in \mathcal{J}(E \cup \chi^k(\zeta) \cup \{\Rightarrow q\})$, where $\Gamma : q \Rightarrow q \in D$ and $\Gamma \in S^k$ and q is free in $E \cup \chi(\zeta)$. Besides, we have imposed the condition $q \in S$. Thus, using the induction hypothesis, we get that $\chi^{k+1}(\zeta) \subseteq S$.

Now we show that $S \subseteq \chi(\zeta)$. Since S is an extension, $S = \bigcup_i S^i$. We prove the inclusion by induction on i.
Base case: $S^0 = \mathcal{J}(E) \subseteq \chi(\zeta)$.
Induction step: $S^{k+1} = \mathcal{J}(E \cup \{\mathsf{CONS}(\Gamma : p \Rightarrow p \in D) \mid \Gamma \in S^k \text{ and } \overline{p} \notin S\})$

Since $\chi(\zeta) \subseteq S$, $\overline{p} \notin S$ implies $\overline{p} \notin \chi(\zeta)$. Besides, by induction hypothesis we have that $S^k \subseteq \chi(\zeta)$; hence, $\Gamma \in \chi(\zeta)$ and thus $\mathsf{CONS}(\Gamma : p \Rightarrow p) \subseteq \chi(\zeta)$.

Therefore, $\chi(\zeta) = S$.

Summing up: we have the representation of the knowledge and beliefs of an agent by a KB-database $\zeta = (E, D)$, where E is a (monotonic) set of clauses that represents knowledge and D is a set of defaults that represents the beliefs of the agent. We have a constructive way to obtain the invariant $\mathcal{J}(E)$ and the extensions of D by means of the $\chi(\zeta)$ operator. The following example shows that this representation may not correspond to the intuition of what a belief is.

Example 4. Let ζ be

$$\begin{array}{rrl} & & \Rightarrow p \\ & q & \Rightarrow r \\ p: & q & \Rightarrow q \\ : & \neg q & \Rightarrow \neg q \\ r: & s & \Rightarrow s \end{array}$$

We have here two extensions: $\mathit{Ext}_1 = \{p, q, r, s\}$ and $\mathit{Ext}_2 = \{p, \neg q\}$.

In the example we cannot say that the agent "believes" either q or $\neg q$; it rather takes these possible scenarios under consideration. If the agent believes q, then some kind of priority must be given to extension E_1. One possibility is to add some extra-logical features, such as priorities [6]. In the next section we will combine the monotonic and the non-monotonic approaches.

4 A Hybrid Representation of Belief

The representation of belief through defaults does not correspond exactly to the usual notion of beliefs. For instance, lemma 13 allows different and incompatible sets of beliefs (as a matter of fact, if two belief sets are different, they *must* be incompatible.) This approach corresponds to the consideration of possible scenarios rather than a set of possible states that is considered more plausible than other.

A possible approach to establish some hierarchy between possible worlds is to give priorities to the defaults. Another one, which is the one we propose here, is to combine the monotonic and the non-monotonic approach. In this case, we re-define a KB-Database as follows:

Definition 17. *Let Π be a set of propositions. A* hybrid knowledge and belief database *(*HKB-database *for short) is a triple $\zeta = (E, D_{cl}, D_{def})$ where E and D_{cl} are sets of clauses and D_{def} is a set of defaults on Π. The set of clauses E is the* epistemic part *of the database and the sets D_{cl} and D_{def} constitute the* doxastic part *of the database. An HKB-database where E is consistent is* epistemically consistent. *An HKB-database where both E and D_{cl} are consistent is* doxastically consistent.

Note that doxastic consistency implies the consistency of both E and D_{cl} and not only the consistency of D_{cl}. This is because everything that is known is also believed. The knowledge of an agent is defined as $\mathcal{J}(E)$ (as before) and the belief of an agent are defined as the extensions constructed starting from $\chi_\zeta^0 = \mathcal{J}(E \cup D_{cl})$. In this way, only the extensions of ζ that are consistent with $\mathcal{J}(E \cup D_{cl})$ are taken into account.

Example 5. Let $\zeta = (E, D_{cl}, D_{def})$ be

$$\begin{array}{rl} E & \left\{ \begin{array}{r} \Rightarrow p \\ q \Rightarrow r \end{array} \right. \\ D_{cl} & \left\{ \begin{array}{r} \Rightarrow q \end{array} \right. \\ D_{def} & \left\{ \begin{array}{r} p : q \Rightarrow q \\ : \neg q \Rightarrow \neg q \\ r : s \Rightarrow s \end{array} \right. \end{array}$$

We have here that $\mathcal{J}(E) = \{p\}$, $\mathcal{J}(E \cup D_{cl}) = \{p, q, r\}$. The only extension is $Ext = \{p, q, r, s\}$.

What is the rôle of the default $: \neg q \Rightarrow \neg q$? It seems that the clause in D_{cl} has made it redundant and that it could be eliminate. This would be true in a static environment, in which the knowledge and beliefs of the agent are permanent. But in a dynamic environment, where new information may become available and the beliefs must be eventually revised, this default may become relevant if new information gives preference to $\neg q$ over q. Recall that belief, in contrast to knowledge, is not necessarily true.

5 Discussion and Future Work

We have presented a representation of the knowledge and beliefs of an agent by a hybrid representation that includes a monotonic part (a set of clauses) and a non-monotonic part (a set of defaults.) The corresponding knowledge- and belief-sets may be iteratively computed. If we represent beliefs purely by defaults, we have no way of selecting among the different extensions, unless we include explicit priorities. The division of the doxastic part into a monotonic and a non-monotonic one provides a natural way to introduce preferences among the different extensions.

We are currently working on several directions. A natural extension of this representation in the case of multiple agents is the introduction of *introspection*, that is, agents have knowledge about other agents' knowledge. This poses several interesting problems, such as *common knowledge* [4], [2]. Besides, private communication may entail inconsistency, since knowledge of an agent a about the knowledge of agent b may turn false as a result of a private communication from a third party. Thus, the axiom that knowledge is always true should be adapted.

Besides, we are interested in the effect of updating the databases. Since beliefs may be wrong, it may be necessary to change the doxastic part. In the case of a pure non-monotonic representation, this poses no problems, since this update amounts to the creation of a new extension, which is incompatible with the other existing ones. In the case of a hybrid representation, it may be possible that some older beliefs must be given up to preserve consistency. The idea is then to decide whether to accept the new belief and incorporate it with minimal changes to the belief database [1] or to reject it.

Finally, we want to explore the situation in which the state of the world changes. Then, unless agents are instantly informed of any change (not always a realistic assumption), their knowledge database may become inconsistent. Once more, the axiom that knowledge is always true may be too strong for practical purposes.

References

1. Alchourrón, Carlos; Gärdenfors, Peter; Makinson, David: *On the Logic of Theory Change: Partial Meet Contraction and Revision Functions*, The Journal of Symbolic Logic 50 (2), 1985, pp. 510-530.
2. Baltag, Alexandru; Moss, Lawrence; Solecki, Sławomir: *The Logic of Public Announcements, Common Knowledge, and Private Suspicions*,CWI Technical Report SEN-R9922, Amsterdam, 1999.
3. Besnard, Philippe: *An Introduction to Default Logic*, Springer Verlag, Berlin-Heidelberg, 1989.
4. Fagin, Ronald; Halpern, Joseph; Moses, Yoram; Vardi, Moshe: *Reasoning About Knowledge*, The MIT Press, Cambridge, MA, 1996.
5. Lloyd, John Wylie: *Foundations of Logic Programming*, Springer Verlag, Berlin, 1987.
6. Marek, Wiktor; Mirosław Truszczyński: *Nonmonotonic Logic*, Springer-Verlag, Berlin, 1993.