Consistent Belief Reasoning in the Presence of Inconsistency

Jinxin (Jason) Lin
Dept. of Computer Science
University of Toronto
Toronto, M5S 1A4, Canada
Email: lin@db.toronto.edu

Abstract

Classical logics are unable to deal with inconsistency, so they are not useful for modeling the reasoning of an agent who has inconsistent beliefs. In this paper, we differentiate consistent belief from inconsistent belief. We propose two belief operators $\mathbf{B^C}$ and $\mathbf{B}$, standing for *consistent belief* and *belief*, respectively. $\mathbf{B^C}$ has the modus ponens property, by which the agent is able to reasoning with consistent beliefs as normal and make consistent conclusions. $\mathbf{B}$ tolerates inconsistency, and with it the agent can reason about his inconsistent beliefs as well. The concept of consistent belief and our logical formalism for it are entirely new, in that reasoning *consistently* about the information in an inconsistent knowledge base is possible. We also present a complete axiomatization for the logic and discuss the application of $\mathbf{B^C}$ and $\mathbf{B}$ in reasoning about the knowledge implicit in a group of agents and inconsistency elimination in knowledge base systems.

1 Introduction

As pointed out by Belnap [Bel75], information processing is often threatened by inconsistency, since computers obtain their data from a variety of sources, none of which can be assumed to be that "paragon of paragons, a universal truth-teller". Ioannidis and Sellis [IS] state that the rules in deductive database systems of most organizations do contain inconsistencies. Halpern [Hal86] concludes that reasoning in the presence of inconsistency is an issue which will need to be considered in the design of knowledge bases.

As everything is a consequence of an inconsistency in classical logic, a single contradiction may ruin all the information in the knowledge base. For instance, if $\text{KB} = \{p, \neg p \wedge q, q \supset r\}$, $\text{KB} \models \alpha$ for any formula α in classical logic. But intuitively q, r are irrelevant to the inconsistency. It is desirable that q and r be distinguished from p and be treated as valuable information.

Since a *knowledge base* could be inconsistent while knowledge in the real world is never inconsistent, it is more appropriate to view a knowledge base to be the reflection of an agent's *belief* rather than the true state of affairs in the real world. Then, we can differentiate two kinds of beliefs: consistent beliefs and inconsistent beliefs. In the knowledge base $\text{KB} = \{p, \neg p \wedge q, q \supset r\}$, we say q and $q \supset r$ are *consistently believed* (thus they are *consistent beliefs* of the agent), since intuitively they are not involved in the inconsistency. On the contrary, we say p and $\neg p$ are *inconsistently believed* (thus they are *inconsistent beliefs* of the agent). Inconsistent beliefs have no meaning and they represent the beliefs that are damaged by the inconsistency, while consistent beliefs are still valuable despite the existence of the inconsistency.

The motivation of our logic is to provide some ways for reasoning about consistent beliefs. As mentioned, in real life a knowledge base may be inconsistent. However, checking and removing inconsistencies in a large knowledge base is difficult and expensive. This is because as we know inconsistency may not lie on the surface and in most cases there is no single solution for eliminating it. Furthermore, the knowledge base may be in use for quite some time before the inconsistency is ever detected [BKMS92]. This can occur in a medical expert system that obtains its knowledge from multiple physician experts, in which it is common that different experts hold conflicting views on their domain of expertise. Then the question is: given the conflict (inconsistency) in the expert system, can *correct* diagnoses be made upon the patients' diseases? Suppose the knowledge base is KB $= \{p, \neg p, q, q \supset r\}$ and p means "the patient has hepatitis". Then it is clear that the diagnosis that the patient has hepatitis is a bad one since another piece of knowledge in the knowledge base ($\neg p$) has already denied it. So is the diagnosis that the patient has not hepatitis. Both could have a detrimental effect on the patient, i.e., the former results that the patient goes through unnecessary treatment and the later results the treatment of the patient being delayed. This shows that a decision should not be made on the basis of inconsistent beliefs. However, if it is not p but r that means "the patient has hepatitis", and $q \supset r$ is a rule of the meaning "if a patient has jaundice then he/she has hepatitis". Then we would not have problem to assert that "the patient has hepatitis" according to the fact that he/she has jaundice. Hence it is essential to be able to reason about consistent beliefs, and a decision or assertion of some fact can be made if it can be " consistently believed" by the knowledge base. This is the main motivation of our logic.

Many many inconsistency-tolerant logic systems have been proposed, e.g., [Bel75, Roz89, dC74, Pri79, Gin88, KL89, Lin87]. In these systems, a proposition generally could have more than two values so that there are interpretations for inconsistent formulas and it is no longer the case that everything is a consequence of an inconsistency. Three or four values were involved in the logic systems of [Bel75, Roz89, dC74, Pri79, Lin87] to avoid triviality in the presence of inconsistency, while more than four values may be involved in the systems of [Gin88, KL89] for modeling non-monotonic reasoning, truth maintenance systems, etc. However, these systems cannot be used to reason about consistent beliefs, since they cannot reason *consistently* about the information in an inconsistent knowledge base. They may, on one hand, conclude a statement and, on the other hand, conclude the negation of it. For example, given the aforesaid KBs, both p and $\neg p$ will be deduced by these systems. While in our logic, as we shall see, the consistent belief operator $\mathbf{B^c}$ is always consistent.

Also, there are modal logics of belief, e.g., [Lev, Lak90, FH88, Var86, FHV90] where inconsistency is not as damaging as in classical logic. Among them, [Lev] is most related to our work. Levesque [Lev] proposed a logic of explicit belief that allows the agent to have inconsistent explicit beliefs without explicitly believing every sentence. Fagin *et al.* [FH88] presented a logic of local reasoning where an agent was viewed as a society of minds, each with its own set of beliefs, which may contradict each other. However, since in these logics the agents' beliefs are generally not closed under implication, they cannot reason about consistent beliefs. For example, given the aforementioned KBs it is not possible to derive r from q and $q \supset r$ in these logics, while in our logic r can be obtained by the modus ponens property of the consistent belief operator.

The work on combining knowledge bases [BKMS92] is also related to the topic of this paper. In [BKMS92], a cautious semantics is defined for inconsistent theories. The cautious semantics considers all maximally consistent subsets of the inconsistent theory, and a sentence is considered true if it is true in all maximally consistent subsets of the theory. Despite the simplicity of this approach, it is dependent on the syntax of the knowledge base. For example, consider a syntactic variant of the two KBs above, KB$' = \{p \wedge \neg p \wedge q \wedge (q \supset r)\}$, then nothing can be derived from KB$'$ since the maximal consistent subset of KB$'$ is empty. This means that if a knowledge base is expressed as a conjunction of sentences then the entire KB is of no use even though the inconsistency may locate only in a single proposition. While in our approach in this paper, KB$'$ is treated exactly the same as the two KBs above and the problem of losing all the information in KB$'$ is avoided.

Also related to our topic are belief revision systems ([G88̈, FKUV86, KM91]), and truth-maintenance systems (e.g., [Doy79]). The approaches of belief revision systems are different from ours. In theirs some knowledge in the knowledge base is given up to resolve inconsistency, while in ours nothing is given up and reasoning is performed in the presence of inconsistency. And as mentioned before, in many cases there is no single solution for how to revise the knowledge base. Moreover, the knowledge base may be in use for quite some time before the inconsistency is ever detected and eliminated. Truth-maintenance systems need to keep track of justifications which induces substantial maintenance cost. And they assume that the initial knowledge base is consistent and that the inconsistency is due to some *assumptions* being introduced during the evolution of the knowledge base. In this paper, we do not have this assumption.

2 The Basic Part of the Logic

The formulas in the language we consider are formed from a set of atomic sentences $\mathcal{P}$ using the three standard connectives $\neg$, $\wedge$, and $\vee$, as usual, and two modal operators (belief operators) $\mathbf{B}$ and $\mathbf{B^c}$. The modal operator $\mathbf{B^c}$ is to capture consistent beliefs and $\mathbf{B}$ to capture beliefs that include not only consistent beliefs but also possibly inconsistent ones. A formula such as $\mathbf{B}\alpha$ is read "the agent believes α", $\mathbf{B^c}\alpha$ is read "the agent consistently believes α". We only consider a single agent and the formulas that do not have nesting belief operators. We choose modal operators to represent beliefs so that it is possible to express both the agent's beliefs and the facts in the real world, e.g., a description $\{\mathbf{B}p, q\}$ means that the agent believes p, and q is true in the real world. By this way we are able to reason not only about an agent's beliefs but also about the true state of affairs in the real world. Note reasoning about the real world is different from reasoning about the beliefs of an agent, because what is in an agent's head may be different from what is happening in the real world.

Other connectives like $\supset$ and $\equiv$ are defined in terms of $\neg$, $\wedge$, and $\vee$. The sentence $\alpha \supset \beta$ is taken as the abbreviation of $\neg\alpha \vee \beta$, and $\alpha \equiv \beta$ as the abbreviation of $(\alpha \supset \beta) \wedge (\beta \supset \alpha)$. Following [Lev90], we give the term *objective* sentences to those without any belief operator, *subjective* sentences to those where all atomic sentences occur within the scope of a belief operator. Since we only consider the formulas that do not have nesting belief operators, it should be understood that formulas appearing within the scope of a belief operator are objective, throughout this paper. And by KB we usually mean a finite set of objective sentences. We usually use p, q, r, a, b, c to denote atomic sentences (atoms for short), l to denote literal (a *literal* is an atom or the negation of an atom). For convenience, we use the notation $-l$ for $\neg p$ if $l = p$, and p if $l = \neg p$, $p \in \mathcal{P}$.

Since no *possible world* is an interpretation of an inconsistent KB, we should generalize the notion of possible worlds, to *situations*, in order to interpret the KB. A *situation* is a partial possible world that supports the truth or falsity or both for every atomic sentence in the underlying language [1]. Incoherent situations are allowed. An *incoherent* situation is one that supports both p and $\neg p$ for some atomic sentence p. Intuitively, a possible world is a situation that is coherent, i.e., it supports either p or $\neg p$ but not both for every atomic sentence p.

Formally, a situation is a pair $(\mathcal{P}, v)$ where v is a mapping: $\mathcal{P} \rightarrow \{\{t\}, \{f\}, \{t, f\}\}$. We use v_s to identify the v associated with situation s, and we call $v_s(p)$ the *set-up* for p in situation s. For a situation s, we use $\Theta(s)$ to denote the set of atoms whose set-ups in s are $\{t, f\}$, i.e., $\Theta(s) = \{p \mid v_s(p) = \{t, f\}, p \in \mathcal{P}\}$. Then intuitively, a possible world is a situation s such that $\Theta(s)$ is empty. We denote the set of all situations by $\mathcal{S}$, and the set of all possible worlds by $\mathcal{W}$. According to what is believed by the agent, we can identify $M \subseteq \mathcal{S}$ to be the set of situations that could be the actual ones, which we call the *belief model* of the agent.

Given a belief model M and a situation $s = (\mathcal{P}, v)$, we now define the support relations $\models_T$ and $\models_F$ between (M, s) and formulas in the language. The symbol $\models_T$ has the meaning of "support the truth of" and $\models_F$ has the meaning of "support the falsity of". The following part of this section is essentially the one used by Levesque [Lev] for defining the basic part of logic of implicit and explicit belief:

[1] The *situation* here is different from the one in [Lev], where Levesque also consider a situation that may support neither p nor $\neg p$ for some atomic sentence p.

1. $M, s \models_T p$ iff $t \in v_s(p)$, where p is an atomic sentence.
   $M, s \models_F p$ iff $f \in v_s(p)$, where p is an atomic sentence.

2. $M, s \models_T (\alpha \wedge \beta)$ iff $M, s \models_T \alpha$ and $M, s \models_T \beta$
   $M, s \models_F (\alpha \wedge \beta)$ iff $M, s \models_F \alpha$ or $M, s \models_F \beta$

3. $M, s \models_T (\alpha \vee \beta)$ iff $M, s \models_T \alpha$ or $M, s \models_T \beta$
   $M, s \models_F (\alpha \vee \beta)$ iff $M, s \models_F \alpha$ and $M, s \models_F \beta$

4. $M, s \models_T \neg\alpha$ iff $M, s \models_F \alpha$
   $M, s \models_F \neg\alpha$ iff $M, s \models_T \alpha$

And (M, s) supports the truth of a set of formulas iff it supports the truth of every formula in the set, and supports the falsity of the set iff it at least supports the falsity of a formula in the set.

In the following, we use the notation $[\![\alpha]\!]$ to denote the set of situations in $\mathcal{S}$ that support the truth of α. And when there is no risk of confusion, we simply use 'support α' for 'support the truth of α'.

3 The Semantics of B^c and B

In order to define the support relations between situations and a formula involving $\mathbf{B^c}$ or $\mathbf{B}$, we need the notion of *most-complete situations.*

3.1 Most-Complete Situations

Return to the example KB $= \{p, \neg p, q, q \supset r\}$, we would like to obtain $\mathbf{B^c}q$ and $\mathbf{B^c}(q \supset r)$ since intuitively q and r has nothing to do with the inconsistency in KB. This means that the effect of the inconsistency should be minimized (in this case, minimized on atom p) so that as much information as possible is obtained as consistent belief. Furthermore, we would like to conclude $\mathbf{B^c}r$ from $\mathbf{B^c}q$ and $\mathbf{B^c}(q \supset r)$ since inconsistency should not disrupt the normal reasoning process of consistent beliefs. And because the agent's beliefs include consistent beliefs, a consistent belief is also a belief, and we want to conclude $\mathbf{B}r$ as well.

However, consider the situations that support KB $= \{p, \neg p, q, q \supset r\}$ as follows:

v	p	q	r
s_1	$\{t, f\}$	$\{t\}$	$\{t\}$
s_2	$\{t, f\}$	$\{t\}$	$\{t, f\}$
s_3	$\{t, f\}$	$\{t, f\}$	$\{t\}$
s_4	$\{t, f\}$	$\{t, f\}$	$\{f\}$
s_5	$\{t, f\}$	$\{t, f\}$	$\{t, f\}$

There is a situation s_4 that supports KB but does not support r. On closer examination of s_4, we observe that in s_4 not only p receives the inconsistent set-up $\{t, f\}$ (which is justified since both p and $\neg p$ are in KB), but also q receives $\{t, f\}$. While in s_1 (which supports r), only p receives $\{t, f\}$. In fact in s_1 fewer atoms receive $\{t, f\}$ than in any other situation, and therefore s_1 is the situation among the five that is "closest" to a possible world (in which no atom receives $\{t, f\}$). Since the inconsistency about p should not have any influence in q and r, we prefer s_1 to the other situations, which means we should minimize the inconsistent set-ups in the set of situations. This suggests that we restrict our attention to the set of situations that are "closest" to possible worlds, which we refer to as *most-complete situations.*

We define a situation s to be more *complete* than s' (denoted by $s <_c s'$) iff $\Theta(s) \subset \Theta(s')$. Define the set of most-complete situations in S to be $\mathcal{MC}(S) = \{s \mid s \in S, \nexists s' \in S$ such that

$s' <_c s\}$. This order of comparing two situations about their closeness to possible worlds is essentially the one used by Priest [Pri91] to define his minimally inconsistent logic, and the one used by Kifer and Lozinskii [KL89] to define their notions of *epistemic entailment*.

And we have the following property of most-complete situation:

Proposition 3.1 *If* KB *is consistent then* $\mathcal{MC}(\llbracket \text{KB} \rrbracket)$ *is the set of possible worlds that satisfy* KB.

Proof: Follows from the fact that for all $w \in \mathcal{W}$ and $s \in \mathcal{S} - \mathcal{W}$, $w <_c s$, and there is no $w' \in \mathcal{W}$ such that $w' <_c w$. ■

3.2 The Definitions of $\mathbf{B^c}$ and $\mathbf{B}$

In our logic, an agent's belief model is a set of situations. As mentioned, we like to restrict our attention to the *most-complete* situations in the set so that the inconsistency is localized and the inference can function as normally as possible. Then it is easy to define that the agent *believes* α iff every most-complete situation in his belief model supports α. For $\mathbf{B^c}$, we require that each of these most-complete situations *coherently* supports α (i.e., supports the truth but not the falsity of α) in order for the agent to *consistently* believes α. We have more to say about the motivation of such definition of $\mathbf{B^c}$ shortly. The semantics for $\mathbf{B}$ and $\mathbf{B^c}$ are:

6. $M, s \models_T \mathbf{B}\alpha$ iff for every $s' \in \mathcal{MC}(M)$, $M, s' \models_T \alpha$. [2]
   $M, s \models_F \mathbf{B}\alpha$ iff $M, s \not\models_T \mathbf{B}\alpha$.

7. $M, s \models_T \mathbf{B^c}\alpha$ iff for every $s' \in \mathcal{MC}(M)$, $M, s' \models_T \alpha$ and $M, s' \not\models_F \alpha$.
   $M, s \models_F \mathbf{B^c}\alpha$ iff $M, s \not\models_T \mathbf{B^c}\alpha$.

We say α is *valid* (write $\models \alpha$) iff $M, w \models_T \alpha$ for all $w \in \mathcal{W}$ and for all $M \subseteq \mathcal{S}$. For an objective α, we may abbreviate $M, s \models_T \alpha$ by $s \models_T \alpha$ since M is not relevant. For a subjective σ, we may write $M \models_T \sigma$ for $M, s \models_T \sigma$ since s is not in question.

The semantics of $\mathbf{B^c}$ requires that *every* situation in $\mathcal{MC}(M)$ supports the truth but not the falsity of α in order for the agent to *consistently* believe α. This is because even a situation in $\mathcal{MC}(M)$ supporting both α and $\neg\alpha$ means α is also involved in an inconsistency. For example, suppose KB $= \{p, \neg p \vee \neg q, q, r\}$, then both p and q *equally* contribute to the inconsistency of KB. And to be conservative, in such a case neither p nor q should be consistently believed. It is easy to verify that $\mathcal{MC}(\llbracket \text{KB} \rrbracket)$ consists of two situations: one assigning $\{t\}$ to p and $\{t, f\}$ to q, and another assigning $\{t, f\}$ to p and $\{t\}$ to q. There is one situation in $\mathcal{MC}(\llbracket \text{KB} \rrbracket)$ that supports both the truth and the falsity of p (or q) and one situation that supports the truth but not the falsity of p (or q, respectively). To rule out the possibility that the proposition which participates (equally) in the cause of inconsistency being consistently believed, we require that *every* situation in $\mathcal{MC}(M)$ must support the truth but not the falsity of α for $\mathbf{B^c}\alpha$ to be true. We notice that in contrast to p and q, both situations support the truth but not the falsity of r, and r is consistently believed, conforming with our intuition.

This actually explains (in part) of the choice of our definition (for $\mathbf{B^c}$) over another possibility which was brought up by one of the reviewers. That is to defined $\mathbf{B^c}$ in terms of $\mathbf{B}$, i.e., $M, s \models_T \mathbf{B^c}\alpha$ iff $M, s \models_T \mathbf{B}\alpha$ and $M, s \not\models_T \mathbf{B}\neg\alpha$. Although this seems to be a plausible alternative, some of its result may not be intuitive. To see this, suppose the knowledge base of the agent is (one that is similar to the above) KB $= \{p, \neg p \vee \neg q, q\}$. And so $M = \llbracket \text{KB} \rrbracket$. Then by this definition, it is easy to verify that $M \models_T \mathbf{B^c}p$ and $M \models_T \mathbf{B^c}q$ (since every situation in $\mathcal{MC}(M)$ supports p and

[2] We may think equivalently in terms of accessibility relation, that for $\mathbf{B}$ and $\mathbf{B^c}$ the set of situations $\mathcal{MC}(M)$ is accessible from any situation, since we restrict our attention to the set of the most complete situations of M. It has been shown that for the modal logic weak S_5 the same fixed set of worlds can be thought of as being accessible from any world.

q and there exist in $\mathcal{MC}(M)$ a situation that does not support $\neg p$ and a situation that does not support $\neg q$). Hence p and q is consistently believed, although each of them contribute in part to the inconsistency of the knowledge base. Then by our intuition $p \wedge q$ should also be consistently believed (since a good property we like to have for consistent beliefs is that consistently believing a conjunction is equivalent to consistently believing each conjunct). However, the agent believes $\neg p \vee \neg q$ which is the negation of $p \wedge q$! An agent certainly cannot consistently believes a sentence if he also believes the negation of the sentence.

4 Properties of $\mathbf{B}^c$ and $\mathbf{B}$

First of all, it is easy to see that an objective sentence is valid iff it is a standard tautology. So the propositional subset of the language is handled correctly, which means that reasoning about the real world is dealt with correctly.

For consistent belief, it is intuitive that it should satisfy the following conditions: (1) consistent belief is consistent, and the set of consistent beliefs is consistent. (2) Consistent belief is closed under implication. (3) The damages caused by the inconsistency is minimal, so that as much information as possible is obtained as consistent beliefs. Condition (1) is rather intuitive and needs not to be explained, while the Condition (2) and (3) have been argued before.

Indeed, for $\mathbf{B^c}$ we have the axioms:

$$\neg \mathbf{B^c}(\neg\alpha), \text{ where } \alpha \text{ is a tautology. }^3 \tag{A1}$$

$$\mathbf{B^c}\alpha \wedge \mathbf{B^c}\beta \equiv \mathbf{B^c}(\alpha \wedge \beta). \tag{A2}$$

$$\mathbf{B^c}\alpha \wedge \mathbf{B^c}(\alpha \supset \beta) \supset \mathbf{B^c}\beta. \tag{A3}$$

The first axiom demonstrates that the agent cannot consistently believe any inconsistent formula (or falsehood). It follows from this axiom that consistent belief is always consistent. The second axiom shows that consistently believing a conjunction is equivalent to consistently believing each conjunct. From these two axioms we know that the set of consistent beliefs is *consistent*, satisfying Condition (1) above. The third axiom is the version of modus ponens for $\mathbf{B^c}$, which shows consistent belief is closed under implication, satisfying Condition (2). This axiom makes the derivations such as inferring $\mathbf{B^c}r$ from $\mathbf{B^c}q$ and $\mathbf{B^c}(q \supset r)$ possible, so the agent is able to draw conclusions from consistent beliefs. Condition (3) is satisfied from the semantics of $\mathbf{B^c}$ which uses the set of most-complete situations where the inconsistency in the knowledge base is localized (or minimized), and the propositions irrelevant to the inconsistency are obtained as consistent beliefs.

For the relationships between belief and consistent belief, we have:

$$\mathbf{B^c}\alpha \supset \mathbf{B}\alpha. \tag{A4}$$

$$\mathbf{B^c}\alpha \supset \neg\mathbf{B}\neg\alpha. \tag{A5}$$

The axiom A4 shows the intuition that beliefs include consistent beliefs. The axiom A5 reflects the following fact. If the agent consistently believes a formula, then he must not believe the negation of the formula for otherwise he believes both the formula and its negation, contradicting the fact that he consistently believes the formula. So we can be assured that consistent belief is never be the cause of inconsistency. For example suppose the knowledge base is $\{p, \neg p \wedge q, q \supset r\}$, then the inconsistency is located at p and the agent does not consistently believe p (or $\neg p$) since he believes $\neg p$ (or p, respectively).

For belief, we have:

$$\mathbf{B}\alpha, \text{ where } \alpha \text{ is a tautology.} \tag{A6}$$

$$\mathbf{B}\alpha \wedge \mathbf{B}\beta \equiv \mathbf{B}(\alpha \wedge \beta). \tag{A7}$$

[3] By tautology we mean standard propositional tautology. Hence here $\neg\alpha$ means a formula that is inconsistent in the sense of classical logic.

The axiom A6 implies that our agent is "perfect" in believing tautologies, as in classical logic. However, $\mathbf{B}$ is different from classical logic because the agent can have inconsistent beliefs without believing every sentence. For example $\{\mathbf{B}p, \mathbf{B}\neg p, \neg\mathbf{B}q\}$ is satisfiable. Axiom A7 is the property of conjunctive belief.

We have limited versions of modus ponens between $\mathbf{B^c}$ and $\mathbf{B}$:

$$\mathbf{B^c}\alpha \wedge \mathbf{B}(\alpha \supset \beta) \supset \mathbf{B}\beta. \tag{A8}$$

$$\mathbf{B}\alpha \wedge \mathbf{B^c}(\alpha \supset \beta) \supset \mathbf{B^c}\beta. \tag{A9}$$

Both axioms are rather intuitive. They say that whenever the antecedent of the implication or the implication itself is *consistently* believed, a conclusion can be drawn through the implication. Indeed, in our daily life if there is no controversy about a fact or a rule, we can usually make conclusions from them.

5 A Complete Axiomatization

In this section we concentrate on obtaining a complete axiomatization for our logic. We have presented some axioms for $\mathbf{B^c}$ and $\mathbf{B}$ in the last section[4]. We shall introduced additional axioms to complete the axiom system.

5.1 The O Operator

To obtain a complete axiomatization for $\mathbf{B^c}$, we need a new belief operator $\mathbf{O}$ to specify the totality of what the agent believes. Intuitively, $\mathbf{O}\alpha$ is read as "α is all that is believed". [5] As we have known, $\mathbf{B}\alpha$ means that the agent believes α, but tends to believe more. We need $\mathbf{O}$ to fix the upper bound of what is believed by the agent so that the set of consistent beliefs can be determined. The semantics of $\mathbf{O}$ is similar to the one in [Lak90, Lev90].

8. $M, s \models_T \mathbf{O}\alpha$ iff for every $s' \in M \Longleftrightarrow M, s' \models_T \alpha$.
   $M, s \models_F \mathbf{O}\alpha$ iff $M, s \not\models_T \mathbf{O}\alpha$.

Since the agent believes α, it is obvious that every situation in M must support α. For the other direction, the idea is that the larger the belief model M is, the less knowledge the agent has. Thus every situation that supports α must be in M in order for the agent to *only believe* α.

With the $\mathbf{O}$ operator, we can further investigate the relationship of $\mathbf{B}$ and $\mathbf{B^c}$ when KB is consistent. If KB is consistent then for any objective α we have:

$$\models \mathbf{O}\textsc{kb} \supset \mathbf{B}\alpha \text{ iff } \models \mathbf{O}\textsc{kb} \supset \mathbf{B^c}\alpha \text{ iff } \textsc{kb} \models \alpha.$$

This is because for a consistent KB, $\mathcal{MC}([\![\textsc{kb}]\!])$ is the set of possible worlds that satisfy KB. Hence if there is no inconsistency in what is believed, belief is identical to consistent belief and is closed under (classical) logical consequence, showing that both $\mathbf{B^c}$ and $\mathbf{B}$ are faithful extensions of classical logic.

5.2 The Complete Axiomatization

The axiom system should include the axioms and rules of inference of the standard propositional logic, as follows:

All instances of tautologies. (A10)

From α and $\alpha \supset \beta$, infer β (modus ponens). (A11)

[4] Note that some axioms are redundant, e.g., A3 can be obtained from A9 and A4, and A1 can be obtained from A6 and A5.

[5] The notion of 'only believing' is first introduced in [Lev90].

In addition, we need axioms stating that $\mathbf{B^c}$, $\mathbf{B}$ and $\mathbf{O}$ respect the standard properties, such as commutativity, associativity, distributivity, De Morgan's laws and double negation. For example, if $\alpha \vee \beta$ is consistently believed then so is $\beta \vee \alpha$, and if α is consistently believed then so is $\neg\neg\alpha$, etc. These properties distinguishes our logic from syntactic ones which suffer from the shortcomings mentioned in [Lev]. We summarize these properties in the following axiom (CNF), where α_{CNF} and α_{DNF} are the conjunctive normal form and disjunctive normal form of α, respectively, and $\mathbf{L}$ stands for either one of $\mathbf{B^c}$, $\mathbf{B}$ or $\mathbf{O}$:

$$\mathbf{L}\alpha \equiv \mathbf{L}\alpha_{\mathrm{CNF}} \equiv \mathbf{L}\alpha_{\mathrm{DNF}}. \tag{A12}$$

$$\mathbf{L}(\alpha \wedge (\alpha \vee \beta)) \equiv \mathbf{L}\alpha. \quad \text{(Axiom of subsuming)} \tag{A13}$$

Also for $\mathbf{O}$, we have the following axiom of "absorption of tautology":

$$\mathbf{O}\alpha \equiv \mathbf{O}(\alpha \wedge \beta), \text{ where } \beta \text{ is a tautology.} \tag{A14}$$

Using these axioms for $\mathbf{O}$, we can transform every sentence within an $\mathbf{O}$ operator to its *standard form*, while preserving validity. We define the *standard form* of an objective sentence α (denoted by α_{STD}) to be the set of clauses [6] in α_{CNF} such that no clause is a standard tautology or subsumed by the other. That is $\alpha_{\mathrm{STD}} = \{C \mid C \in \alpha_{\mathrm{CNF}}, \nexists p$ such that $p, \neg p \in C$, and $\nexists C' \in \alpha_{\mathrm{CNF}}$ such that $C' \subset C\}$. The significance of the *standard form* is that it completely determines whether two objective sentences have the same set of supporting situations. More specifically, we have $\alpha_{\mathrm{STD}} = \beta_{\mathrm{STD}}$ iff $[\![\alpha]\!] = [\![\beta]\!]$, see Appendix Lemma A.1. Then we have the following axiom to generate negative 'only-belief' ($\neg\mathbf{O}$) from 'only-belief':

$$\mathbf{O}\alpha \supset \neg\mathbf{O}\beta, \text{ where } \beta \text{ is any objective formula such that } \beta_{\mathrm{STD}} \neq \alpha_{\mathrm{STD}}. \tag{A15}$$

The axiom that describes the relationship of $\mathbf{O}$ and $\mathbf{B}$ is easy:

$$\mathbf{O}\psi \supset \mathbf{B}\psi. \tag{A16}$$

Finally, some axioms for generating consistent belief and negative consistent belief ($\neg\mathbf{B^c}$) from $\mathbf{O}$, are needed. First, we give some definitions. Let the disjunctive normal form of ψ be ψ_{DNF} which is a set of dual clauses. For a dual clause D, we define D^u to be the set of literals l in D such that $-l$ is also in D, and $D^c = D - D^u$. Then define $\Omega(\psi) = \{D \mid D \in \psi_{\mathrm{DNF}}, \nexists Q \in \psi_{\mathrm{DNF}}\ s.t.\ Q^u \subset D^u\}$ (in the appendix, we call the dual clauses in $\Omega(\psi)$ the *dual clauses of least inconsistency* in ψ_{DNF}). Define $\psi^* = \bigvee_{D \in \Omega(\psi)} D^c$. [7]

Then we have:

$$\mathbf{O}\psi \supset \mathbf{B^c}\psi^* \tag{A17}$$

$$\mathbf{O}\psi \supset \neg\mathbf{B^c}\neg D, \text{ for all } D \in \Omega(\psi). \tag{A18}$$

We have the following soundness and completeness result:

Theorem 5.1 *The axioms A1- A18 give a sound and complete axiomatization for the above logic.*

Proof: See Appendix. ■

6 Applications

In addition to reasoning in the presence of inconsistency, the logic is useful in reasoning about implicit knowledge in a group of agents, and in inconsistency elimination from a knowledge base.

[6] We will omit the primitive connectives $\vee$ and $\wedge$ in α_{CNF} and α_{DNF}, and therefore α_{CNF} is a set of *clauses* (a *clauses* is a set of literals where $\vee$ is implied), and α_{DNF} is a set of *dual clauses* (a *dual clauses* is a set of literals where $\wedge$ is implied).

[7] Here D^c should be understood as $\bigwedge D^c$, for $\wedge$ is implied among the elements of a dual clause. We will continue to omit $\wedge$ as if the meaning of a dual clause can be understood.

6.1 Reasoning about Implicit Knowledge in a Group of Agents

In many cases an expert system needs to encode the knowledge of multiple experts, so that the resulting expert system (or knowledge base) will be able to derive the facts that neither of the experts individually knows [BKMS92]. In other words, it is desirable to be able to reason about the knowledge that is *implicit* in a group of agents (called *implicit knowledge* in [HM85]). For example, if an agent knows a and another agent knows $a \rightarrow b$ then together they have implicit knowledge of b, even though it might be the case that neither of them individually knows b. If there is no contradiction among the agents' knowledge, implicit knowledge can be easily obtained by computing the logical consequences of the union of all agents' knowledge. However, different agents (or experts) can, and often do, hold conflicting views on their domain of expertise. Then the union of all agents' knowledge is not meaningful in classical logic. We need some kind of mechanism to resolve (or eliminate) the conflicts and still be able to infer the knowledge that is implicit in the group of agents.

With the above logic, we can give an interpretation to the union of all agent's knowledge. Let $\mathrm{KB}_1, ..., \mathrm{KB}_n$ be the knowledge bases of n agents, we then can define that α is implicit knowledge if and only if $\models \mathbf{O}(\mathrm{KB}_1 \cup ... \cup \mathrm{KB}_n) \supset \mathbf{B^c}\alpha$. Immediately we have the following useful properties of this definition: (1) the set of implicit knowledge is consistent; (2) if $\mathrm{KB}_1 \cup ... \cup \mathrm{KB}_n$ is consistent, then α is implicit knowledge iff $\mathrm{KB}_1 \cup ... \cup \mathrm{KB}_n \models \alpha$.

Example 6.1 Suppose $\mathrm{KB}_1 = \{a, b\}$, $\mathrm{KB}_2 = \{\neg a, b \rightarrow c\}$, and $\mathrm{KB}_3 = \{c \rightarrow d\}$, then b, c, d are all implicit knowledge of the three agents. We can see that the conflicts among the agents are minimized and the knowledge *implied* in the group of agents is inferred.

If we do not require that the set of implicit knowledge is consistent, we can use the semantics: α is implicit knowledge iff $\models \mathbf{O}(\mathrm{KB}_1 \cup ... \cup \mathrm{KB}_n) \supset \mathbf{B}\alpha$. As in the last definition, in the case that $\mathrm{KB}_1 \cup ... \cup \mathrm{KB}_n$ is consistent, α is implicit knowledge iff $\mathrm{KB}_1 \cup ... \cup \mathrm{KB}_n \models \alpha$.

Example 6.2 Suppose three knowledge bases are as in Example 6.1, then $a, \neg a, b, c, d$ are all implicit knowledge of the three agents. We can see the set of implicit knowledge concluded may be inconsistent, but the knowledge *implied* in the group of agents is inferred.

It is of interest to notice that such definitions of implicit knowledge are independent of the exact *syntax form* of each knowledge base [8]. For example, the three knowledge bases $\mathrm{KB}'_1 = \{a \wedge b\}$, $\mathrm{KB}'_2 = \{\neg a \wedge (b \rightarrow c)\}$, and $\mathrm{KB}'_3 = \{c \rightarrow d\}$ have the same set of implicit knowledge as in Example 6.1. This would not be the case if the syntactical approaches in [BKMS92] are used. By those approaches the implicit knowledge c and d cannot be derived from the result of merging $\mathrm{KB}'_1, \mathrm{KB}'_2$, and KB'_3, since c and d are not the logical consequences of all maximal consistent subsets of $\mathrm{KB}'_1 \cup \mathrm{KB}'_2 \cup \mathrm{KB}'_3$.

Other methods of merging the knowledge of multiple agents are discussed in [LMa].

6.2 Inconsistency Elimination

As most existing inference systems are based on classical logic, it is essential to eliminate any contradiction (inconsistency) once it arises so that the inference can function properly. The question is what information in the knowledge base should be eliminated in order to restore consistency and what information should be preserved. Since $\mathbf{B^c}$ represents those that are unrelated to the cause of the inconsistency, we may want consistent beliefs to be preserved, i.e., we want the result of elimination to imply the set of consistent beliefs. For example, suppose $\mathrm{KB} = \{p, \neg p, q, q \rightarrow r\}$, then since q and $q \rightarrow r$ are consistent beliefs, the intuitive result is $\{p, q, q \rightarrow r\}$ or $\{\neg p, q, q \rightarrow r\}$. Depending on the preference about inconsistent beliefs (p and $\neg p$), we can choose one of them.

[8] Two knowledge bases are treated equivalently as long as they have the same set of supporting situations.

Suppose only consistent beliefs are preserved (and the inconsistent beliefs are eliminated) then the result is $\{q, q \rightarrow r\}$. Again, we notice that this procedure is independent of the exact syntax form of the knowledge base, as in the last subsection. For example, the knowledge base $\mathrm{KB}' = \{p \wedge \neg p \wedge q \wedge (q \rightarrow r)\}$ will have the same result of elimination as the above KB. Details will be explored in [LMb].

7 Conclusion

The above logic provides a formal semantic for reasoning under inconsistency in terms of two belief operators $\mathbf{B^c}$ and $\mathbf{B}$. The phenomenon of inconsistent beliefs seems to occur even in science. As cited in [FH88], the physicist Eugene Wigner noted that the two great theories physicists reason with are the theory of quantum phenomena and the theory of relativity. However, Wigner thought that the two theories might well be incompatible! The incompatibility should not prevent us from reasoning about other facts irrelevant to the theory of quantum phenomena and the theory of relativity, e.g., inferring that Sam is not immortal from the fact that he is a man. The logic represented by the $\mathbf{B^c}$ operator provides us such a way of reasoning *consistently* about the information in a knowledge base although the knowledge base may be inconsistent.

The logic represented by the $\mathbf{B}$ operator allows the agent to hold inconsistent beliefs. However the effect of inconsistency is localized (or minimized) in the sense that conclusions can still be drawn through implication from consistent beliefs. Hence $\mathbf{B}$ can be used to model the reasoning of the agent who may have beliefs that contradict each other.

Future research directions include generalizing the logic to the first-order case and/or multiple agent case, considering those formulas that have nested belief operators (care should be taken when we consider nested beliefs and positive or negative introspection [GMR92]), and investigating the usefulness of the logic in the area of counterfactuals.

Acknowledgments

I am indebted to Alberto Mendelzon for his detailed and enlightening comments on the drafts of this paper. I also thank Anthony Bonner, Hector Levesque, Fangzhen Lin, Ray Reiter, and Danny Zilio for their helpful comments. This work was supported by the Institute for Robotics and Intelligent Systems, Canada.

A The Proof of Soundness

Theorem A.1 *The axioms A1- A18 are sound.*

Proof: **A1:** Follows from axiom A6 which will be proved below and axiom A5 whose soundness is easy to see.

A3: Let M be a set of situations such that $M \models_T \mathbf{B^c}\alpha \wedge \mathbf{B^c}(\neg\alpha \vee \beta)$. Consider $s \in \mathcal{MC}(M)$. From $M \models_T \mathbf{B^c}\alpha$ we have $s \models_T \alpha$, and $s \not\models_F \alpha$. Hence $s \models_F \neg\alpha$ and $s \not\models_T \neg\alpha$. By $M \models_T \mathbf{B^c}(\neg\alpha \vee \beta)$ we have $s \models_T \neg\alpha \vee \beta$ and $s \not\models_F \neg\alpha \vee \beta$. From $s \models_T \neg\alpha \vee \beta$ and $s \not\models_T \neg\alpha$ it follows that $s \models_T \beta$. And from $s \not\models_F \neg\alpha \vee \beta$ and $s \models_F \neg\alpha$ it follows that $s \not\models_F \beta$. Hence $M \models_T \mathbf{B^c}\beta$.

A6: Let M be any set of situations and α be a tautology. Consider $s \in \mathcal{MC}(M)$. Let w be a possible world such that for all $p \in \mathcal{P}$, $v_w(p) = \{t\}$ if $v_s(p) = \{t, f\}$ and $v_w(p) = v_s(p)$ otherwise. Then we have $v_w(p) \subseteq v_s(p)$ for all atom $p \in \mathcal{P}$. Since α is a tautology which is supported by every possible world, we have $w \models_T \alpha$. It follows that $s \models_T \alpha$, and since s is any situation in $\mathcal{MC}(M)$, we have $M \models_T \mathbf{B}\alpha$. Since M is any set of situations we have $\models \mathbf{B}\alpha$.

A8,A9: similar to the proof of A3.

A15: Let M be any set of situations. For the case that $M \not\models_T \mathbf{O}\alpha$, by the semantics we know $M \models_T \neg\mathbf{O}\alpha$ and the soundness of the axiom follows trivially. So we consider the case that $M \models_T \mathbf{O}\alpha$. By the semantics we know $M = [\![\alpha]\!]$. Let β be an objective formula such that $\beta_{\text{STD}} \neq \alpha_{\text{STD}}$. Then from Lemma A.1 we have $[\![\alpha]\!] \neq [\![\beta]\!]$. Hence $M \neq [\![\beta]\!]$. By the semantics of $\mathbf{O}$ we have $M \not\models_T \mathbf{O}\beta$, i.e., $M \models_T \neg\mathbf{O}\beta$.

A17: Let M be a set of situations. As above, we only need to consider the case that $M \models_T \mathbf{O}\psi$. We have $M \models_T \mathbf{O}\psi_{\text{DNF}}$. By the semantics, $M = [\![\psi_{\text{DNF}}]\!]$. Let $s \in \mathcal{MC}(M)$ then according to Lemma A.2, there is $D \in \Omega(\psi)$ such that $s \models_T D$. Hence $s \models_T D^c$, and consequently $s \models_T \psi^*$. Now we want to prove $s \not\models_T \neg D^c$. Assume to the contrary that $s \models_T \neg D^c$, then $s \models_T -l$ for some $l \in D^c$. Meanwhile $s \models_T l$ since $s \models_T D^c$. Therefore $v_s(p) = \{t, f\}$, where p is the atom appearing in l. Let s' be the situation whose set-ups are identical everywhere to that of s except p in which case $v_{s'}(p)$ is $\{t\}$ if $l = p$, is $\{f\}$ if $l = \neg p$. Then $s' \models_T D$, and it follows that $s' \models_T \psi_{\text{DNF}}$, i.e., $s' \in M$. But it is easy to see $\Theta(s') \subset \Theta(s)$, contradicting the fact that $s \in \mathcal{MC}(M)$. Hence $s \not\models_T \neg D^c$, and it follows that $s \not\models_T \neg\psi^*$. Since s is an arbitrary one in $\mathcal{MC}(M)$, we have $M \models_T \mathbf{B^C}\psi^*$. Thus $M \models_T \mathbf{O}\psi \supset \mathbf{B^C}\psi^*$.

A18: Let M be a set of situations. We only need to consider the case that $M \models_T \mathbf{O}\psi$. We have $M \models_T \mathbf{O}\psi_{\text{DNF}}$. By the semantics, $M = [\![\psi_{\text{DNF}}]\!]$. Let $\alpha \in \Omega(\psi)$, then according to Lemma A.2, there is some $s \in \mathcal{MC}(M)$ such that $s \models_T \alpha$. It follows that $M \not\models_T \mathbf{B^C}\neg\alpha$. By the semantics, we have $M \models_T \neg\mathbf{B^C}\neg\alpha$. Hence $\models \mathbf{O}\psi \supset \neg\mathbf{B^C}\neg\alpha$, for all $\alpha \in \Omega(\psi)$.

It is easy to see the other axioms are sound. ■

Lemma A.1 *For any two objective formulas α and β, $[\![\alpha]\!] = [\![\beta]\!]$ iff $\alpha_{\text{STD}} = \beta_{\text{STD}}$.*

Proof: "$\Leftarrow$": $\alpha_{\text{STD}} = \beta_{\text{STD}}$, want to prove $[\![\alpha]\!] = [\![\beta]\!]$.

Let $s \in [\![\alpha]\!]$, then $s \models_T \alpha_{\text{CNF}}$. We have $s \models_T \alpha_{\text{STD}}$, so $s \models_T \beta_{\text{STD}}$ from the fact that $\alpha_{\text{STD}} = \beta_{\text{STD}}$. It follows that $s \models_T C$, for any $C \in \beta_{\text{STD}}$. For the clause $C \notin \beta_{\text{STD}}$ but $C \in \beta_{\text{CNF}}$, we know either C is a tautology or there is $C' \in \beta_{\text{STD}}$ such that $C' \subset C$. If C is a tautology then it is trivial that $s \models_T C$. If there is $C' \in \beta_{\text{STD}}$ such that $C' \subseteq C$, then since $s \models_T C'$ we also have $s \models_T C$. So $s \models_T C$, for all $C \in \beta_{\text{CNF}}$. Hence $s \models_T \beta_{\text{CNF}}$. It follows that $s \models_T \beta$, i.e., $s \in [\![\beta]\!]$. Similarly we can prove for all $s \in [\![\beta]\!]$, $s \in [\![\alpha]\!]$. Thus $[\![\alpha]\!] = [\![\beta]\!]$.

"$\Rightarrow$": Let α and β be two objective formulas such that $[\![\alpha]\!] = [\![\beta]\!]$.

Let $\alpha_{\text{CNF}} = \{C_1, \ldots, C_n\}$ be the conjunctive normal form of α. Then $\alpha_{\text{STD}} = \{C \mid C \in \alpha_{\text{CNF}}, \nexists p$ such that $p, \neg p \in C$, and $\nexists C' \in \alpha_{\text{CNF}}$ such that $C' \subset C\}$. From the proof of "$\Leftarrow$" we have $[\![\alpha]\!] = [\![\alpha_{\text{STD}}]\!]$. Similarly we have $[\![\beta]\!] = [\![\beta_{\text{STD}}]\!]$. Hence $[\![\alpha_{\text{STD}}]\!] = [\![\beta_{\text{STD}}]\!]$

We now prove that for any $C \in \alpha_{\text{STD}}$, there is a $Q \in \beta_{\text{STD}}$ such that $C = Q$, and vice versa. We first prove for any $C \in \alpha_{\text{STD}}$, there is a $Q \in \beta_{\text{STD}}$ such that $Q \subseteq C$. Suppose not, then for all $Q \in \beta_{\text{STD}}$ there is a literal in Q but not in C. Let s be the situation such that for any $p \in \mathcal{P}$,
$v_s(p) = \{f\}$ if $p \in C$,
$v_s(p) = \{t\}$ if $\neg p \in C$, and
$v_s(p) = \{t, f\}$ otherwise.
Notice we can assign such set-ups for s, since there is no $C \in \alpha_{\text{STD}}$ that contains both p and $\neg p$ for any $p \in \mathcal{P}$ otherwise C is a tautology which α_{STD} does not contain. Then $s \not\models_T C$, consequently $s \not\models_T \alpha_{\text{STD}}$. But since for any $Q \in \beta_{\text{STD}}$ there is a literal in Q but not in C, we have $s \models_T Q$,for any $Q \in \beta_{\text{STD}}$. It follows that $s \models_T \beta_{\text{STD}}$. Hence $s \in [\![\beta_{\text{STD}}]\!]$, but since $s \notin [\![\alpha_{\text{STD}}]\!]$ it contradicts the fact that $[\![\alpha_{\text{STD}}]\!] = [\![\beta_{\text{STD}}]\!]$.

Similarly we can prove for $Q \in \beta_{\text{STD}}$ there is a $C' \in \alpha_{\text{STD}}$ such that $C' \subseteq Q$. Hence $C' \subseteq Q \subseteq C$. Since there is no element in α_{STD} that is subsumed by some other elements in it, we have $C = C' = Q$. The proof for the other direction is similar; thus we have for all $C \in \alpha_{\text{STD}}$, there is a $Q \in \beta_{\text{STD}}$ such that $C = Q$, and vice versa. This is equivalent to say $\alpha_{\text{STD}} = \beta_{\text{STD}}$ since we regard α_{STD} and β_{STD} to be sets where the order of elements is irrelevant.
■

Lemma A.2 *For a set of literals D, let $\Phi(D) = \{p \mid p \in \mathcal{P}$ such that $p, \neg p \in D\}$. Let ψ be in disjunctive normal form (DNF). Then*
(1) for each $s \in \mathcal{MC}(\llbracket\psi\rrbracket)$ there exists a dual clause of least inconsistency $D \in \psi$ such that $s \models_T D$ and $\Theta(s) = \Phi(D)$;
(2) for each dual clause of least inconsistency $D \in \psi$ there exists $s \in \mathcal{MC}(\llbracket\psi\rrbracket)$ such that $s \models_T D$ and $\Theta(s) = \Phi(D)$.

Proof: (1). Let $s \in \mathcal{MC}(\llbracket\psi\rrbracket)$. Since $s \models_T \psi$, we have $s \models_T D$, for some $D \in \psi$. Suppose D is not of least inconsistency, then $\exists D' \in \psi$ such that $\Phi(D') \subset \Phi(D)$. Consider the situation s^* such that for $p \in \mathcal{P}$, $v_{s^*}(p) = \{t, f\}$ if $p, \neg p \in D'$, $v_{s^*}(p) = \{t\}$ if $p \in D', \neg p \notin D'$, $v_{s^*}(p) = \{f\}$ if $\neg p \in D', p \notin D'$, and $v_{s^*}(p) = \{t\}$ otherwise. Then it is easy to see $s^* \models_T D'$, thus $s^* \models_T \psi$. And $\Theta(s^*) = \Phi(D') \subset \Phi(D) \subseteq \Theta(s)$. Therefore $s^* <_c s$. This contradicts the assumption that $s \in \mathcal{MC}(\llbracket\psi\rrbracket)$.

(2). Suppose $D \in \psi$ is of least inconsistency, we show that there exists $s \in \mathcal{MC}(\llbracket\psi\rrbracket)$ such that $s \models_T D$ and $\Theta(s) = \Phi(D)$. Let s be the situation such that for $p \in \mathcal{P}$, $v_s(p) = \{t, f\}$ if $p, \neg p \in D$, $v_s(p) = \{t\}$ if $p \in D, \neg p \notin D$, $v_s(p) = \{f\}$ if $\neg p \in D, p \notin D$, and $v_s(p) = \{t\}$ or $\{f\}$ otherwise. Then it is easy to see $s \models_T D$, and $\Theta(s) = \Phi(D)$. Assume $s \notin \mathcal{MC}(\llbracket\psi\rrbracket)$, then $\exists s' \in \mathcal{MC}(\llbracket\psi\rrbracket)$ such that $\Theta(s') \subset \Theta(s)$. From (1) we know $\exists D' \in \psi$ of least inconsistency such that $s' \models_T D'$ and $\Phi(D') = \Theta(s')$. Since $\Theta(s') \subset \Theta(s) = \Phi(D)$ we have $\Phi(D') \subset \Phi(D)$, contradicting the fact that D is of least inconsistency in ψ. ■

B The Proof of Completeness

First, we give some definitions and notation. For a set of sentences, we use $\bigwedge$ (or $\bigvee$) before the set to denote the conjunction (or disjunction, respectively) over the sentences in the set. The notion of *theorem* is defined as usual, that is, α is a *theorem* iff it is provable from an empty set (denoted by $\vdash \alpha$). We write $\Sigma \vdash \alpha$ if there is a finite subset Σ' of Σ such that $\vdash \bigwedge \Sigma' \supset \alpha$. For two sentences α and β, we also use $\alpha \vdash \beta$ as a convenient abbreviation for $\vdash \alpha \supset \beta$. A formula α is *inconsistent* if $\neg\alpha$ is a theorem, and *consistent* otherwise. A set of sentences Σ is *inconsistent* if it contains a finite subset Σ' such that $\bigwedge \Sigma'$ is inconsistent. And a set of sentences is a *maximal consistent set* if it is consistent and any proper superset is inconsistent. We will use the following usual properties of maximal consistent set (see, for example, [?]):

- Any consistent set can be extended to a maximal consistent set.
- If Σ is a maximal consistent set, then for any formulas ψ and ϕ:
  - $\psi \wedge \phi \in \Sigma$ iff $\psi \in \Sigma$ and $\phi \in \Sigma$, and either $\psi \in \Sigma$ or $\neg\psi \in \Sigma$.
  - If $\vdash \psi$ then $\psi \in \Sigma$.
  - If $\Sigma \vdash \psi$ then $\psi \in \Sigma$.

We shall prove in Part 1 some preparatory lemmas that will be used in the proof of the main theorem in Part 2.

B.1 Part 1

Lemma B.1 *Suppose $M = \llbracket\alpha\rrbracket$. Then: (1). If $M \models_T \mathbf{O}\beta$ then $\mathbf{O}\alpha \vdash \mathbf{O}\beta$.*
(2). If $M \models_T \neg\mathbf{O}\beta$ then $\mathbf{O}\alpha \vdash \neg\mathbf{O}\beta$.

Proof: (1). Let $M = \llbracket\alpha\rrbracket$ and $M \models_T \mathbf{O}\beta$. Then by the semantics of $\mathbf{O}$ we have $M = \llbracket\alpha\rrbracket = \llbracket\beta\rrbracket$. Hence $\alpha_{\mathrm{STD}} = \beta_{\mathrm{STD}}$ by Lemma A.1. Since transforming every sentence within an $\mathbf{O}$ operator to its standard form preserves validity, we have $\vdash \mathbf{O}\alpha \equiv \mathbf{O}\alpha_{\mathrm{STD}}$ and $\vdash \mathbf{O}\beta \equiv \mathbf{O}\beta_{\mathrm{STD}}$, it follows that $\mathbf{O}\alpha \vdash \mathbf{O}\beta$.

(2). Let $M = [\![\alpha]\!]$ and $M \models_T \neg\mathbf{O}\beta$, then by the semantics of $\mathbf{O}$ we have $[\![\alpha]\!] \neq [\![\beta]\!]$. Hence $\alpha_{\mathrm{STD}} \neq \beta_{\mathrm{STD}}$ by Lemma A.1. By axiom A15, we have $\mathbf{O}\alpha \vdash \neg\mathbf{O}\beta$. ■

Lemma B.2 *Suppose $M = [\![\psi]\!]$. Then: (1). If $M \models_T \mathbf{B^c}\alpha$ then $\mathbf{O}\psi \vdash \mathbf{B^c}\alpha$.*
(2). If $M \models_T \neg\mathbf{B^c}\alpha$ then $\mathbf{O}\psi \vdash \neg\mathbf{B^c}\alpha$.

Proof: Let $\psi_{\mathrm{DNF}} = \{D_1, \ldots, D_n\}$ be the disjunctive normal form of ψ, in particular $D_1 \ldots D_k, k \leq n$ be the dual clauses of least inconsistency in ψ_{DNF}.

(1). Suppose $M \models_T \mathbf{B^c}\alpha$. Let α_{CNF} be the conjunctive normal form of α. Then $M \models_T \mathbf{B^c}\alpha_{\mathrm{CNF}}$ and hence $M \models_T \mathbf{B^c}C, \forall C \in \alpha_{\mathrm{CNF}}$. We now prove for all C in α_{CNF}, $\mathbf{O}\psi \vdash \mathbf{B^c}C$. There are two cases:

- There is no $p \in \mathcal{P}$ such that both p and $\neg p$ are in C. We now show for all $D_i, i \in [1,k]$, D_i^c contains at least a literal of C. Suppose not for some D_i, then from Lemma A.2 we know $\exists s \in \mathcal{MC}(M), s \models_T D_i$. Let s' be the situation such that for all $p \in \mathcal{P}$,
  $v_{s'}(p) = \{f\}$ if $p \in (C - D_i^u)$,
  $v_{s'}(p) = \{t\}$ if $\neg p \in (C - D_i^u)$, and
  $v_{s'}(p) = v_s(p)$ otherwise.

  Notice we can assign such set-ups for s', since there is no $p \in \mathcal{P}$ such that both p and $\neg p$ are in C. Then $s' \models_T D_i^c$ since $s \models_T D_i^c$ and D_i^c contains no literal of C. Also $s' \models_T D_i^u$ since $s \models_T D_i^u$. Hence $s' \models_T D_i$, and it follows that $s' \models_T \psi_{\mathrm{DNF}}$. Thus $s' \in M$. Since $s \in \mathcal{MC}(M)$ and $\Theta(s') \subseteq \Theta(s)$, we have $s' \in \mathcal{MC}(M)$. By the construction of s', we have for all $l \in (C - D_i^u)$, $s' \models_T -l$. For all literal $l \in C \cup D_i^u$, $l \in D_i^u$ and hence $-l \in D_i^u$. From $s' \models_T D_i^u$ it follows that $s' \models_T -l$ also. Thus $s' \models_T \neg C$, which contradicts the fact that $M \models_T \mathbf{B^c}C$.

  Hence each $D_i^c, i \in [1,k]$ contains at least a literal of C, which we denote as l_i. Then we have $\vdash \mathbf{B^c}\psi^* \supset \mathbf{B^c}(l_1 \vee \ldots \vee l_k)$. By the axiom of "subsuming", we have $\vdash \mathbf{B^c}\psi^* \supset \mathbf{B^c}C$. Since $\mathbf{O}\psi \vdash \mathbf{B^c}\psi^*$, it follows that $\mathbf{O}\psi \vdash \mathbf{B^c}C$.

- There is some $p \in \mathcal{P}$ such that both p and $\neg p$ are in C. Let $C^u = \{l \mid l \in C, -l \in C\}$. Then C^u is not empty. We now prove for each $D_i, i \in [1,k]$, if D_i^c does not contain any literal of C then there exists $q \in C^u$ such that $q \notin D_i^u$. Assume to the contrary that there is $D_i^c, i \in [1,k]$ which does not contain any literal of C and $C^u \subseteq D_i^u$. Then from Lemma A.2 we know $\exists s \in \mathcal{MC}(M), s \models_T D_i$. Let s' be the situation such that for all $p \in \mathcal{P}$,
  $v_{s'}(p) = \{f\}$ if $p \in (C - D_i^u)$,
  $v_{s'}(p) = \{t\}$ if $\neg p \in (C - D_i^u)$, and
  $v_{s'}(p) = v_s(p)$ otherwise.

  Since there is no $p \in \mathcal{P}$ such that both p and $\neg p$ are in $C - C^u$ and $C^u \subseteq D_i^u$, there is no $p \in \mathcal{P}$ such that both p and $\neg p$ are in $C - D_i^u$. Hence we can assign such set-ups for s'.

  Then $s' \models_T D_i^c$ since $s \models_T D_i^c$ and D_i^c contains no literals of C. Also $s' \models_T D_i^u$ since $s \models_T D_i^u$. Hence $s' \models_T D_i$, and it follows that $s' \models_T \psi_{\mathrm{DNF}}$. Thus $s' \in M$. Since $s \in \mathcal{MC}(M)$ and $\Theta(s') \subseteq \Theta(s)$, we have $s' \in \mathcal{MC}(M)$. By the construction of s', we have for all $l \in (C - D_i^u)$, $s' \models_T -l$. For all literal $l \in C \cup D_i^u$, $l \in D_i^u$. Therefore $-l \in D_i^u$. From $s' \models_T D_i^u$ it follows that $s' \models_T -l$ also. Thus $s' \models_T \neg C$, which contradicts the fact that $M \models_T \mathbf{B^c}C$.

  Therefore for each $D_i, i \in [1,k]$, if D_i^c does not contain any literal of C then there exists $q \in C^u$ such that $q \notin D_i^u$. Since $q \vee \neg q \vee D_1 \vee \ldots \vee D_{i-1} \vee D_{i+1} \vee \ldots \vee D_n$ is a tautology,

$$\vdash \mathbf{O}\psi_{\mathrm{DNF}} \equiv \mathbf{O}((D_1 \vee \ldots \vee D_n) \wedge (q \vee \neg q \vee D_1 \vee \ldots \vee D_{i-1} \vee D_{i+1} \vee \ldots \vee D_n))$$

Let $X_i = D_i \cup \{q\}$ and $Y_i = D_i \cup \{\neg q\}$. Let $\psi' = D_1 \vee \ldots \vee D_{i-1} \vee X_i \vee Y_i \vee D_{i+1} \vee \ldots \vee D_n$. Then by (CNF) transformation,

$$\vdash \mathbf{O}\psi_{\mathrm{DNF}} \equiv \mathbf{O}\psi'.$$

Since $q \notin D_i^u$ and D_i^c does not contain any literal of C, we have $q \notin D_i$. It follows that $X_i^u = Y_i^u = D_i^u$, and hence X_i^u and Y_i^u are the dual clauses of least inconsistency in ψ'. And each of X_i^u and Y_i^u contains a literal of C, since $q, \neg q \in C$.

For a set of literals D, let us call D^c temporarily its *consistent part*. Hence now D_i whose consistent part does not contain any literal of C has been substituted by two dual clauses whose consistent parts contain at least one literal of C. By this way we can find a formula ϕ such that $\vdash \mathbf{O}\psi_{\mathrm{DNF}} \equiv \mathbf{O}\phi$ and every dual clause of least inconsistency in ϕ whose consistent part contains at least one literal of C. Then similar to the last case, we can prove $\mathbf{B^c}\phi^* \vdash \mathbf{B^c}C$. Since $\mathbf{O}\psi \vdash \mathbf{O}\phi$ and $\mathbf{O}\phi \vdash \mathbf{B^c}\phi^*$, it follows that $\mathbf{O}\psi \vdash \mathbf{B^c}C$.

So in both cases we have $\mathbf{O}\psi \vdash \mathbf{B^c}C$. Similarly, we have $\mathbf{O}\psi \vdash \mathbf{B^c}C'$ for any other element C' in α_{CNF}. By the axiom of conjunctive beliefs we have $\mathbf{O}\psi \vdash \mathbf{B^c}\alpha_{\mathrm{CNF}}$. Thus $\mathbf{O}\psi \vdash \mathbf{B^c}\alpha$.

(2). Suppose $M \models_T \neg\mathbf{B^c}\alpha$. Then there is $s \in \mathcal{MC}(M)$ such that $s \models_T \neg\alpha$. Let $\beta = \neg\alpha$ and $\beta_{\mathrm{DNF}} = \{Q_1, \ldots, Q_m\}$ be the disjunctive normal form of β. Then $s \models_T Q_i$, for some $i \in [1, m]$. From Lemma A.2 we know $s \models_T D_j$, for some $j \in [1, k]$, and $\Theta(s) = \Phi(D_j)$. Now there are two cases:

- $Q_i \subseteq D_j$. Since D_j is a dual clause of least inconsistency in ψ_{DNF}, by axiom A18 we have $\vdash \mathbf{O}\psi \supset \neg\mathbf{B^c}(\neg D_j)$. Since $Q_i \subseteq D_j$ we have $\vdash \mathbf{B^c}(\neg Q_i) \supset \mathbf{B^c}(\neg D_j)$. It follows that $\vdash \mathbf{O}\psi \supset \neg\mathbf{B^c}(\neg Q_i)$. We also have $\vdash \mathbf{B^c}(\neg\beta_{\mathrm{DNF}}) \supset \mathbf{B^c}(\neg Q_i)$, and $\vdash \mathbf{B^c}\alpha \equiv \mathbf{B^c}(\neg\beta_{\mathrm{DNF}})$. Hence $\mathbf{O}\psi \vdash \neg\mathbf{B^c}\alpha$.

- $Q_i \not\subseteq D_j$. Then for all $l \in Q_i - D_j$, $-l \notin D_j$. Assume to the contrary that there is $l \in Q_i - D_j$, $-l \in D_j$. Then $s \models_T -l$ since $s \models_T D_j$, and $s \models_T l$ since $s \models_T Q_i$. It follows from $\Theta(s) = \Phi(D_j)$ that $l \in D_j$, since only the atoms appearing in D_j^u receive set-up $\{t, f\}$ in s. This contradicts the fact that $l \in Q_i - D_j$. Also there is no l in $Q_i - D_j$ such that $-l$ is also in $Q_i - D_j$, for otherwise we also have $s \models_T l$ and $s \models_T -l$. This contradicts the fact that $l \in Q_i - D_j$.

  Let $Q_i - D_j = \{l_1, \ldots, l_h\}$. Let $\delta = ((l_1 \vee -l_1) \wedge (l_2 \vee -l_2) \wedge \ldots \wedge (l_h \vee -l_h)) \vee D_1 \vee \ldots \vee D_{j-1} \vee D_{j+1} \vee \ldots \vee D_n$. Since δ is a tautology,

  $$\vdash \mathbf{O}\psi_{\mathrm{DNF}} \equiv \mathbf{O}((D_1 \vee \ldots \vee D_n) \wedge \delta).$$

  Let ϕ be the disjunctive normal form of $(D_1 \vee \ldots \vee D_n) \wedge \delta$. Then by (CNF) transformation, $\vdash \mathbf{O}\psi_{\mathrm{DNF}} \equiv \mathbf{O}\phi$. Notice now ϕ contains a dual clause $X = D_j \cup \{l_1, \ldots, l_h\}$. And $X^u = D_j^u$, since neither of $\{-l_1, \ldots, -l_h\}$ are in D_j or in $\{l_1, \ldots, l_h\}$. Hence X is a dual clause of least inconsistency in ϕ. Notice $Q_i \subseteq X$. Then similar to the last case, we can prove $\mathbf{O}\phi \vdash \neg\mathbf{B^c}\alpha$. Since $\vdash \mathbf{O}\psi \equiv \mathbf{O}\phi$, we also have $\mathbf{O}\psi \vdash \neg\mathbf{B^c}\alpha$.

■

Lemma B.3 *Suppose $M = [\![\psi]\!]$. Then: (1). If $M \models_T \mathbf{B}\alpha$ then $\mathbf{O}\psi \vdash \mathbf{B}\alpha$. (2). If $M \models_T \neg\mathbf{B}\alpha$ then $\mathbf{O}\psi \vdash \neg\mathbf{B}\alpha$.*

Proof: Let $\psi_{\mathrm{DNF}} = \{D_1, \ldots, D_n\}$ be the disjunctive normal form of ψ, in particular $D_1 \ldots D_k$, $k \leq n$ be the dual clauses of least inconsistency in ψ_{DNF}.

(1). Suppose $M \models_T \mathbf{B}\alpha$.

By $\mathbf{O}\psi \supset \mathbf{B}\psi$ and Axiom (CNF) we have $\mathbf{O}\psi \vdash \mathbf{B}\psi_{\mathrm{DNF}}$.

Let α_{CNF} be the conjunctive normal form of α. Then $M \models_T \mathbf{B}\alpha_{\mathrm{CNF}}$. Consequently $M \models_T \mathbf{B}C, \forall C \in \alpha_{\mathrm{CNF}}$. We now prove for all C in α_{CNF}, $\mathbf{O}\psi \vdash \mathbf{B}C$. There are two cases:

- There is some $p \in \mathcal{P}$ such that both p and $\neg p$ in C. Then C is a tautology. By A6 we know $\mathbf{O}\psi \vdash \mathbf{B}C$.

- There is no $p \in \mathcal{P}$ such that both p and $\neg p$ in C. We now show each $D_i, i \in [1, k]$ contains at least a literal of C. Suppose not for some D_i, then from Lemma A.2 we know $\exists s \in \mathcal{MC}(M), s \models_T D_i$. Let s' be the situation such that,
  $v_{s'}(p) = \{f\}$ if $p \in C$,
  $v_{s'}(p) = \{t\}$ if $\neg p \in C$, and
  $v_{s'}(p) = v_s(p)$ otherwise.

  Then $s' \models_T D_i$ since all literals in C are not in D_i. Hence $s' \models_T \psi_{\mathrm{DNF}}$. It follows that $s' \in M$. And by the construction of s', we have $\Theta(s') \subseteq \Theta(s)$. Since $s \in \mathcal{MC}(M)$ it follows that $s' \in \mathcal{MC}(M)$. But $s' \not\models_T C$, contradicting $M \models_T \mathbf{B}C$.

  Hence for all $D_i, i \in [1, k]$, D_i contains at least a literal of C, which we denote as l_i. Then we have $\vdash \mathbf{B}\psi_{\mathrm{DNF}} \supset \mathbf{B}(l_1 \vee \ldots \vee l_k \vee D_{k+1} \vee \ldots \vee D_n)$. But $l_1, \ldots, l_k \in C$. By the axiom of "subsuming" for $\mathbf{B}$, we have $\vdash \mathbf{B}\psi_{\mathrm{DNF}} \supset \mathbf{B}(C \vee D_{k+1} \vee \ldots \vee D_n)$.

  Notice $D^u_{k+1}, \ldots, D^u_n$ are not empty since $D_{k+1}, \ldots, D_n$ are not of least inconsistency. By (CNF) transformation and axiom of conjunctive belief, we have $\vdash \mathbf{B}\psi_{\mathrm{DNF}} \supset \mathbf{B}(C \vee D^u_{k+1} \vee \ldots \vee D^u_n)$.

  We now prove for each $D^u_j, j \in [k+1, \ldots, n]$, $M \models_T \mathbf{B^c}(\neg D^u_j)$. Assume to the contrary that there is $j \in [k+1, \ldots, n]$ such that $M \not\models_T \mathbf{B^c}(\neg D^u_j)$. Then $M \models_T \neg\mathbf{B^c}(\neg D^u_j)$. It means there is $s \in \mathcal{MC}(M)$ such that $s \models_T D^u_j$. Hence $s \models_T l$ for all $l \in D^u_j$. It follows that $\Phi(D_j) \subseteq \Theta(s)$. From Lemma A.2 we know there is $D_h, h \in [1, k]$ such that $\Phi(D_h) = \Theta(s)$. Hence $\Phi(D_j) \subseteq \Phi(D_h)$, as a result $D^u_j \subseteq D^u_h$, which contradicts the fact that D_j is not of least inconsistency in ψ_{DNF}.

  Therefore $M \models_T \mathbf{B^c}(\neg D^u_j)$ for all $j \in [k+1, \ldots, n]$. By Lemma B.2 we have $\mathbf{O}\psi \vdash \mathbf{B^c}(\neg D^u_j)$ for all $j \in [k+1, \ldots, n]$.

  Then by repeatedly applying A8 we arrive $\mathbf{O}\psi \vdash \mathbf{B}C$.

Similarly, we have $\mathbf{O}\psi \vdash \mathbf{B}C'$ for any other elements C' in α_{CNF}. By the axiom of conjunctive beliefs we have $\mathbf{O}\psi \vdash \mathbf{B}\alpha_{\mathrm{CNF}}$, and consequently $\mathbf{O}\psi \vdash \mathbf{B}\alpha$.

(2) Let $M \models_T \neg\mathbf{B}\alpha$. By the semantics of $\mathbf{B}$, there exists $s^* \in \mathcal{MC}(M), s^* \not\models_T \alpha$. We now prove $\mathbf{O}\psi \vdash \neg\mathbf{B}\alpha$.

For every situation s, let

$$\sigma^s = \bigvee_{p \in (\mathcal{P} - \Theta(s))} \begin{cases} \neg p & \text{if } v_s(p) = \{t\} \\ p & \text{if } v_s(p) = \{f\}. \end{cases}$$

Let $\beta = \bigwedge_{s \in \mathcal{MC}(M), s \neq s^*} \sigma^s$, then it is easy to see for all $s \in \mathcal{MC}(M)$ and $s \neq s^*, s \not\models_T \beta$. Since $s^* \not\models_T \alpha$, we have for all $s \in \mathcal{MC}(M), s \not\models_T \alpha \wedge \beta$. It follows that $M \models_T \mathbf{B^c}(\neg\alpha \vee \neg\beta)$. We now prove $M \models_T \neg\mathbf{B^c}\neg\beta$. It is sufficient to prove $s^* \models_T \beta$. Suppose $s^* \not\models_T \beta$, then there is $s \in \mathcal{MC}(M), s \neq s^*$ such that $s^* \not\models_T \sigma^s$. This means $v_{s^*}(p) = v_s(p), \forall p \in \mathcal{P} - \Theta(s)$. But since $s \neq s^*$, there is at least one $q \in \mathcal{P}$ such that $v_s(q) \neq v_{s^*}(q)$. It follows that $q \in \Theta(s)$. Hence we have $v_{s^*}(q) = \{t\}$ or $\{f\}$ but $v_s(q) = \{t, f\}$. Consequently $\Theta(s^*) \subset \Theta(s)$. This contradicts the fact that s is a most-complete situation in M. Hence $s^* \models_T \beta$, and it follows that $M \models_T \neg\mathbf{B^c}(\neg\beta)$. From Lemma B.2 we know $\mathbf{O}\psi \vdash \mathbf{B^c}(\neg\alpha \vee \neg\beta)$ and $\mathbf{O}\psi \vdash \neg\mathbf{B^c}(\neg\beta)$. Since $\mathbf{B}\alpha \wedge \mathbf{B^c}(\alpha \supset \neg\beta) \supset \mathbf{B^c}(\neg\beta)$ is a theorem, it follows that $\mathbf{O}\psi \vdash \neg\mathbf{B}\alpha$. ■

B.2 Part 2

Theorem B.1 *The axioms A1-A18 are complete with respect to the above logic.*

Proof: To prove the completeness, we should prove that every valid formula is provable. Equivalently, we prove every consistent formula is satisfiable. The proof is much like the one in [Lev]. Given a consistent formula σ, we can extend $\{\sigma\}$ into a maximal consistent set Σ. Then Σ contains a set of literals, conjunctions, disjunctions and the like, together with a set Γ like:

$$\{\mathbf{O}\alpha_i\} \cup \{\neg\mathbf{O}\beta_i\} \cup \{\mathbf{B}\gamma_i\} \cup \{\neg\mathbf{B}\delta_i\} \cup \{\mathbf{B^c}\rho_i\} \cup \{\neg\mathbf{B^c}\phi_i\}.$$

By the propositional completeness we can find a possible world s satisfies the set of literals in Σ. Suppose we can prove Γ is satisfiable, i.e., we can find a set of situations M supports Γ, then $M, s \models_T \Sigma$ by the inductive definition of support. And since $\sigma \in \Sigma$ it follows that $M, s \models_T \sigma$. This means σ is satisfiable. So every consistent formula is satisfiable. In that case we can conclude that our axiom system is complete.

Now we show Γ is satisfiable. There are two cases:

(1) There is at least an $\mathbf{O}\alpha_i \in \Gamma$. Let $M = [\![\alpha_i]\!]$. Then by the semantics of $\mathbf{O}$, $M \models_T \mathbf{O}\alpha_i$. For the other $\mathbf{O}\alpha_j \in \Gamma, j \neq i$, we have $M \models_T \mathbf{O}\alpha_j$ otherwise $M \models_T \neg\mathbf{O}\alpha_j$ and by Lemma B.1 we have $\mathbf{O}\alpha_i \vdash \neg\mathbf{O}\alpha_j$. This means Σ is inconsistent since it contains both $\mathbf{O}\alpha_j$ and $\neg\mathbf{O}\alpha_j$, a contradiction. For each β_i, $M \models_T \neg\mathbf{O}\beta_i$ otherwise $M \models_T \mathbf{O}\beta_i$ and by Lemma B.1, we have $\mathbf{O}\alpha_i \vdash \mathbf{O}\beta_i$; once again we have Σ is inconsistent.

The proof of $M \models_T \{\mathbf{B}\gamma_i\} \cup \{\neg\mathbf{B}\delta_i\} \cup \{\mathbf{B^c}\rho_i\} \cup \{\neg\mathbf{B^c}\phi_i\}$ is similar. Therefore $M \models_T \Gamma$, which means Γ is satisfiable.

(2) There is no $\mathbf{O}\alpha_i \in \Gamma$, i.e.,

$$\Gamma = \{\neg\mathbf{O}\beta_i\} \cup \{\mathbf{B}\gamma_i\} \cup \{\neg\mathbf{B}\delta_i\} \cup \{\mathbf{B^c}\rho_i\} \cup \{\neg\mathbf{B^c}\phi_i\}.$$

We want to prove there exists a set of situations M that supports Γ. First we need the following lemma:

Lemma B.4 *(a).For each δ in the above consistent set Γ, there exists a situation s such that $s \models_T \bigwedge\gamma_i, s \not\models_T \neg(\bigwedge\rho_i)$ and $s \not\models_T \delta$.*
(b).For each ϕ in the above consistent set Γ, there exists a situation [9] s such that $s \models_T \bigwedge\gamma_i, s \not\models_T \neg(\bigwedge\rho_i)$ and $s \models_T \neg\phi$.

Proof: Let $\mu = \bigwedge\gamma_i, \nu = \neg(\bigwedge\rho_i)$.

First we prove there is a situation s such that $s \models_T \mu, s \not\models_T \nu$. Suppose it is not the case, then for all $s \models_T \mu, s \models_T \nu$. This means $[\![\mu]\!] \subseteq [\![\nu]\!]$. By Lemma B.5, $\mathbf{B}\mu \supset \mathbf{B}\nu$ is a theorem. Since $\mathbf{B}\nu \supset \neg\mathbf{B^c}\neg\nu$ is a theorem, $\mathbf{B}\mu \supset \neg\mathbf{B^c}\neg\nu$ is also a theorem. However by the axiom of conjunctive beliefs, we have $\Gamma \vdash \mathbf{B}\mu \wedge \mathbf{B^c}\neg\nu$. Hence Σ is inconsistent, a contradiction.

Hence there is at least one situation s such that $s \models_T \mu, s \not\models_T \nu$.

Let $\mu_{\text{DNF}} = D_1 \vee \ldots \vee D_n$,and $(\neg\nu)_{\text{DNF}} = Q_1 \vee \ldots \vee Q_m$. In particular, let $Q_1, \ldots, Q_k, k \leq m$ be the dual clauses that do not contain both p and $\neg p$ for any p. Notice there is at least one Q_i which does not contain both p and $\neg p$ for any $p \in \mathcal{P}$, for otherwise ν is a propositional tautology. It follows that $\neg\mathbf{B^c}\neg\nu$ is a theorem. However by the axiom of conjunctive beliefs, we have $\Gamma \vdash \mathbf{B^c}\neg\nu$. This means Σ is inconsistent, a contradiction.

First we prove $\Gamma \vdash \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k)$. Clearly $\vdash \mathbf{B^c}\bigwedge\rho_i \supset \mathbf{B^c}(Q_1 \vee \ldots \vee Q_m)$. Notice each $Q_j, j \in [k+1, m]$ contains p and $\neg p$, for some atom p. From Lemma B.7, we know $\mathbf{B^c}((\alpha\wedge\neg\alpha\wedge\beta)\vee\gamma) \equiv \mathbf{B^c}\gamma$ is a theorem. Hence we can remove the dual clauses $Q_j, j \in [k+1, m]$ from $Q_1 \vee \ldots \vee Q_m$, while preserving validity of consistent belief. That is we have $\vdash \mathbf{B^c}(\bigwedge\rho_i) \supset \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k)$. By the axioms of conjunctive beliefs, we have $\Gamma \vdash \mathbf{B^c}(\bigwedge\rho_i)$ and hence $\Gamma \vdash \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k)$.

[9] The situation here can be different from the one in (a).

(a). Assume to the contrary, there is δ in Γ such that all situations supporting μ but not ν support δ.

Consider $C \in \delta_{\mathrm{CNF}}$ that is not a tautology. Without loss of generality, suppose $D_1, \ldots, D_g$, $Q_1, \ldots, Q_h$ ($g \le n, h \le k$) contain no literal of C (then $D_{g+1}, \ldots, D_n$, $Q_{h+1}, \ldots, Q_k$ contain at least one literal of C). We claim that for $D_i, Q_j, i \in [1,g], j \in [1,h]$, there is $l \in D_i$ such that $-l \in Q_j$. Assume for $D_i, Q_j, i \in [1,g], j \in [1,h]$, there is not $l \in D_i$ such that $-l \in Q_j$. Consider the set of literals $\{-l \mid l \in C\} \cup Q_j$. This set does not contain both p and $\neg p$ for any p, since C is not tautology (hence it does not contain both $p, \neg p$ for any p), Q_j does not contain both $p, \neg p$ for any p, and Q_j contains no literal of C. Hence there is a situation s^* such that for all $p \in \mathcal{P}$,
$v_{s^*}(p) = \{t\}$ if $p \in \{-l \mid l \in C\} \cup Q_j$, and
$v_{s^*}(p) = \{t, f\}$ otherwise.

Then $s^* \not\models_T l, \forall l \in C$ and $s^* \not\models_T -l, \forall l \in Q_j$. Then since D_i contains no literal of C and there is not $l \in D_i$ such that $-l \in Q_j$, $s^* \models_T D_i$. And $s^* \not\models_T \neg Q_j$ but $s^* \not\models_T C$. It follows that $s^* \models_T \mu$, $s^* \not\models_T \nu$ but $s^* \not\models_T \delta$, contradicting our assumption that all situations supporting μ but not ν support δ.

So for any $D_i, Q_j, i \in [1,g], j \in [1,h]$, there is $l \in D_i, -l \in Q_j$. Consider D_1. We know for each $Q_j, j \in [1,h]$, there is $l_j \in D_1, -l_j \in Q_j$. Then by (CNF) transformation and axiom of conjunctive belief,

$$\vdash \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B^c}(-l_1 \vee \ldots \vee -l_h \vee Q_{h+1} \vee \ldots \vee Q_k).$$

And

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \supset \mathbf{B}((l_1 \wedge \ldots \wedge l_h) \vee D_2 \vee \ldots \vee D_n).$$

Notice $\vdash l_1 \wedge \ldots \wedge l_h \equiv \neg(-l_1 \vee \ldots \vee -l_h)$. By Lemma B.6, we have

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B}(D_2 \vee \ldots \vee D_n \vee Q_{h+1} \vee \ldots \vee Q_k).$$

Let η represent the formula $D_{g+1} \vee \ldots \vee D_n \vee Q_{h+1} \vee \ldots \vee Q_k$. Then similarly we can prove

$$\vdash \mathbf{B}(D_2 \vee \ldots \vee D_g \vee \eta) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B}(D_3 \vee \ldots \vee D_g \vee \eta).$$

$$\ldots,$$

$$\vdash \mathbf{B}(D_g \vee \eta) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B}\eta.$$

Hence we have

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B}(D_{g+1} \vee \ldots \vee D_n \vee Q_{h+1} \vee \ldots \vee Q_k).$$

Since $\Gamma \vdash \mathbf{B}(\bigwedge \gamma_i) \wedge \mathbf{B^c}(\bigwedge \rho_i)$, we have $\Gamma \vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k)$. Then $\Gamma \vdash \mathbf{B}(D_{g+1} \vee \ldots \vee D_n \vee Q_{h+1} \vee \ldots \vee Q_k)$. Notice if $g = n$ and at the same time $h = k$ then at the last step of the above procedure, we have

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B}(l_1 \wedge \ldots \wedge l_k) \wedge \mathbf{B^c}(-l_1 \vee \ldots \vee -l_k).$$

Then $\Gamma \vdash \mathbf{B}(l_1 \wedge \ldots \wedge l_k) \wedge \mathbf{B^c}(-l_1 \vee \ldots \vee -l_k)$, but since we have the axiom $\mathbf{B}\alpha \supset \neg\mathbf{B^c}\neg\alpha$ it renders that Σ is inconsistent, a contradiction. So there is at least one disjunct in $D_{g+1} \vee \ldots \vee D_n \vee Q_{h+1} \vee \ldots \vee Q_k$. Notice all elements of $\{D_{g+1}, \ldots, D_n, Q_{h+1}, \ldots, Q_k\}$ contain at least one literal of C. By the axiom (CNF) and subsuming we have $\vdash \mathbf{B}(D_{g+1} \vee \ldots \vee D_n \vee Q_{h+1} \vee \ldots \vee Q_k) \supset \mathbf{B}C$. Hence $\Gamma \vdash \mathbf{B}C$. Similarly we can prove $\Gamma \vdash \mathbf{B}C'$ for the other $C' \in \delta_{\mathrm{CNF}}$ which is not a tautology. For the $C' \in \delta_{\mathrm{CNF}}$ which is a tautology, it is trivial that $\mathbf{B}C'$ is a theorem. By the axiom of conjunctive beliefs and (CNF) we have $\Gamma \vdash \mathbf{B}\delta$, which means Σ is inconsistent, a contradiction.

Therefore, for each δ in the consistent set Γ, there exists a situation $s \in \mathcal{S}$ such that $s \models_T \bigwedge \gamma_i, s \not\models_T \neg(\bigwedge \rho_i)$ and $s \not\models_T \delta$.

(b). Assume to the contrary, all situations which support μ but not ν do not support $\neg\phi$.

Without loss of generality, this time assume $Q_1, \ldots, Q_h, h \leq k$ are the Q in $(\neg\nu)_{\mathrm{DNF}}$ such that for each $D_i, i \in [1,n]$ there exists some literal $l \in D_i, -l \in Q$. Consider Q_1. For each $D_i, i \in [1,n]$, there is $l_i \in D_i, -l_i \in Q_1$. Then by (CNF) transformation and axiom of conjunctive belief, we have

$$\vdash \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B^c}((-l_1 \wedge \ldots \wedge -l_n) \vee Q_2 \vee \ldots \vee Q_h).$$

And $\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \supset \mathbf{B}(l_1 \vee \ldots \vee l_n)$. Notice $\vdash l_1 \vee \ldots \vee l_n \equiv \neg(-l_1 \wedge \ldots \wedge -l_n)$. By axiom A9 we have

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B^c}(Q_2 \vee \ldots \vee Q_k).$$

By repeating the same reasoning, we have

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B^c}(Q_{h+1} \vee \ldots \vee Q_k).$$

Notice $h \neq k$ otherwise at the last step of the above procedure, we have

$$\vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k) \supset \mathbf{B}(l_1 \vee \ldots \vee l_n) \wedge \mathbf{B^c}(-l_1 \wedge \ldots \wedge -l_n).$$

Since $\Gamma \vdash \mathbf{B}(D_1 \vee \ldots \vee D_n) \wedge \mathbf{B^c}(Q_1 \vee \ldots \vee Q_k)$ we have $\Gamma \vdash \mathbf{B}(l_1 \vee \ldots \vee l_n) \wedge \mathbf{B^c}(-l_1 \wedge \ldots \wedge -l_n)$. But since we have the axiom $\mathbf{B}\alpha \supset \neg\mathbf{B^c}\neg\alpha$ it means that Σ is inconsistent, a contradiction. So there is at least one disjunct in $Q_{h+1} \vee \ldots \vee Q_k$. And we have $\Gamma \vdash \mathbf{B^c}(Q_{h+1} \vee \ldots \vee Q_k)$.

Consider $C \in \phi_{\mathrm{CNF}}$, we now prove each $Q_j, j \in [h+1,k]$ contains at least one literal of C. Suppose not for some $Q_j, j \in [h+1,k]$, then there is a situation s^* such that for all $p \in \mathcal{P}$,
$v_{s^*}(p) = \{t\}$ if $p \in Q_j$,
$v_{s^*}(p) = \{f\}$ if $\neg p \in Q_j$,
$v_{s^*}(p) = \{t,f\}$ otherwise.
Such s^* exists because Q_j does not contain both $p, \neg p$ for any p. By the construction of s^*, we have $s^* \not\models_T -l$ iff $l \in Q_j$. So it is easy to see $s^* \not\models_T \neg Q_j$, and hence $s^* \not\models_T \nu$. Notice for the Q_j there is a $D_i, i \in [1,n]$ such that there is no literal $l \in D_i, -l \in Q_j$, and hence $s^* \models_T D_i$. It follows that $s^* \models_T \mu$.

But $s^* \models_T -l, \forall l \in C$ since no literal of C is contained in Q_j. Hence $s^* \models_T \neg C$, from which it follows that $s^* \models_T \neg\phi$, contradicting our assumption that all situations which support μ but not ν do not support $\neg\phi$.

Hence each $Q_j, j \in [h+1,k]$ contains at least one literal of C. Then by the axiom (CNF) and subsuming we have $\vdash \mathbf{B^c}(Q_{h+1} \vee \ldots \vee Q_k) \supset \mathbf{B^c}C$. Similarly we have $\vdash \mathbf{B^c}(Q_{h+1} \vee \ldots \vee Q_k) \supset \mathbf{B^c}C'$ for the other $C' \in \phi_{\mathrm{CNF}}$. By the axiom of conjunctive beliefs and (CNF) we have $\vdash \mathbf{B^c}(Q_{h+1} \vee \ldots \vee Q_k) \supset \mathbf{B^c}\phi$. Since $\Gamma \vdash \mathbf{B^c}(Q_{h+1} \vee \ldots \vee Q_k)$, it follows that $\Gamma \vdash \mathbf{B^c}\phi$, which means Σ is inconsistent, a contradiction.

This concludes the proof that for each ϕ in Γ, there exists a situation s such that $s \models_T \bigwedge\gamma_i, s \not\models_T \neg(\bigwedge\rho_i)$ and $s \models_T \neg\phi$. ■

We now proceed to prove the main theorem. We construct a set of situations S such that for each δ_i in Γ, S contains one situation promised by the above lemma, i.e., a situation s such that $s \models_T \bigwedge\gamma_i, s \not\models_T \neg(\bigwedge\rho_i)$ and $s \not\models_T \delta_i$, and for each ϕ_i in Γ, S contains one situation also promised by the above lemma, i.e., a situation s such that $s \models_T \bigwedge\gamma_i, s \not\models_T \neg(\bigwedge\rho_i)$ and $s \models_T \neg\phi_i$. Suppose $S = \{s_1, \ldots, s_n\}$, and $\{p_1, \ldots, p_n\}$ be a set of atoms that are not mentioned in Γ. For $i \in [1,n]$, let s_i' be the situation whose set-ups are identical to that of s_i, except for those atoms in $\{p_1, \ldots, p_n\}$ in which case, let $v_{s_i'}(p_i) = \{t,f\}$ and $v_{s_i'}(p_j) = \{t\}, \forall j \neq i$ and $j \in [1,n]$.

Let $M = \{s_1', \ldots, s_n'\}$. Then from the set-ups for the atoms in $\{p_1, \ldots, p_n\}$, it is easy to see there is no situation in M that is more complete than any other situation in the set. Hence $\mathcal{MC}(M) = M$. We know every situation in M supports $\bigwedge\gamma_i$, as every situation in S does so.

Hence $M \models_T \mathbf{B}\gamma_i$ for every γ_i. Every situation in M does not support $\neg(\bigwedge \rho_i)$ (then it supports $\bigwedge \rho_i$), as every situation in S does not either. Hence $M \models_T \mathbf{B^c}\rho_i$ for every ρ_i. For each δ_i, there is a situation in M that does not support it, as there is a situation in S that does not either. Hence $M \models_T \neg\mathbf{B}\delta_i$. For each ϕ_i, there is a situation in M that supports $\neg\phi_i$, as there is a situation in S that does so. Hence $M \models_T \neg\mathbf{B^c}\phi_i$. Let s^* be the situation that assigns the set-up $\{t, f\}$ to every atom including those in $\{p_1, \ldots, p_n\}$. Then $s^* \models_T \beta_i$ for each β_i. But $s^* \notin M$, therefore for each β_i, $M \models_T \neg\mathbf{O}\beta_i$. In a word, $M \models_T \Gamma$, so Γ is satisfiable. This concludes the proof of completeness. ■

Lemma B.5 *If* $[\![\mu]\!] \subseteq [\![\nu]\!]$, *then* $\vdash \mathbf{B}\mu \supset \mathbf{B}\nu$.

Proof: Let $\mu_{\text{DNF}} = \{D_1, \ldots, D_n\}$. Then $[\![\mu_{\text{DNF}}]\!] \subseteq [\![\nu_{\text{CNF}}]\!]$.

For each $C \in \nu_{\text{CNF}}$ that is not a tautology, we prove for all $D \in \mu_{\text{DNF}}$, D contains at least one literal of C. Suppose there is $D \in \mu_{\text{DNF}}$ contains no literal of C. Notice there are not $p, \neg p \in C$ for any p otherwise C is a tautology. Let s^* be the situation such that for $p \in \mathcal{P}$,
$v_{s^*}(p) = \{t\}$ if $\neg p \in C$,
$v_{s^*}(p) = \{f\}$ if $p \in C$,
$v_{s^*}(p) = \{t, f\}$ otherwise.
Then $s^* \not\models_T C$. It follows that $s^* \not\models_T \nu_{\text{CNF}}$. But $s^* \models_T D$ since D contains no literal of C. Therefore $s^* \models_T \mu_{\text{DNF}}$, contradicting $[\![\mu_{\text{DNF}}]\!] \subseteq [\![\nu_{\text{CNF}}]\!]$.

Hence for all $D_i, i = 1, \ldots, n$, D_i contains at least a literal of C, which we assume to be l_i. Then by CNF transformation and axiom of conjunctive belief, we have $\vdash \mathbf{B}\mu_{\text{DNF}} \supset \mathbf{B}(l_1 \vee \ldots \vee l_n)$. And by axiom of "subsuming" for $\mathbf{B}$, we have $\vdash \mathbf{B}\mu_{\text{DNF}} \supset \mathbf{B}C$. Similarly we can prove $\vdash \mathbf{B}\mu_{\text{DNF}} \supset \mathbf{B}C'$, for the other $C' \in \nu_{\text{CNF}}$ that is not tautology. Since if C' is a tautology, $\vdash \mathbf{B}C'$. By the axiom of conjunctive beliefs, $\vdash \mathbf{B}\mu_{\text{DNF}} \supset \mathbf{B}\nu_{\text{CNF}}$. By axiom (CNF), we have $\vdash \mathbf{B}\mu \supset \mathbf{B}\nu$. ■

Lemma B.6 $\vdash \mathbf{B}(\alpha \vee \beta) \wedge \mathbf{B^c}(\neg\alpha \vee \gamma) \supset \mathbf{B}(\beta \vee \gamma)$.

Proof: Since $\gamma \vee \neg\gamma \vee \beta$ is a tautology, it follows that $\vdash \mathbf{B}(\gamma \vee \neg\gamma \vee \beta)$. Hence $\vdash \mathbf{B}(\alpha \vee \beta) \equiv (\mathbf{B}(\alpha \vee \beta) \wedge \mathbf{B}(\gamma \vee \neg\gamma \vee \beta))$. By (CNF) transformation, $\vdash \mathbf{B}(\alpha \vee \beta) \equiv \mathbf{B}((\alpha \wedge \neg\gamma) \vee (\alpha \wedge \gamma) \vee \beta)$. Notice $\vdash \alpha \wedge \neg\gamma \equiv \neg(\neg\alpha \vee \gamma)$. By the the modus ponens for $\mathbf{B}$ and $\mathbf{B^c}$, we have $\vdash \mathbf{B}((\alpha \wedge \neg\gamma) \vee (\alpha \wedge \gamma) \vee \beta) \wedge \mathbf{B^c}(\neg\alpha \vee \gamma) \supset \mathbf{B}((\alpha \wedge \gamma) \vee \beta)$.

Then $\vdash \mathbf{B}(\alpha \vee \beta) \wedge \mathbf{B^c}(\neg\alpha \vee \gamma) \supset \mathbf{B}((\alpha \wedge \gamma) \vee \beta)$. Since $\vdash \mathbf{B}((\alpha \wedge \gamma) \vee \beta) \supset \mathbf{B}(\gamma \vee \beta)$, we have $\vdash \mathbf{B}(\alpha \vee \beta) \wedge \mathbf{B^c}(\neg\alpha \vee \gamma) \supset \mathbf{B}(\beta \vee \gamma)$. ■

Lemma B.7 $\vdash \mathbf{B^c}(\alpha \wedge \neg\alpha \wedge \beta \vee \gamma) \equiv \mathbf{B^c}\gamma$.

Proof: It is easy to see $\vdash \mathbf{B^c}\gamma \supset \mathbf{B^c}((\alpha \wedge \neg\alpha \wedge \beta) \vee \gamma)$.

We now prove $\vdash \mathbf{B^c}((\alpha \wedge \neg\alpha \wedge \beta) \vee \gamma) \supset \mathbf{B^c}\gamma$. By axiom (CNF) we have $\vdash \mathbf{B^c}((\alpha \wedge \neg\alpha \wedge \beta) \vee \gamma) \supset \mathbf{B^c}((\alpha \wedge \neg\alpha) \vee \gamma)$. Since $\alpha \vee \neg\alpha$ is a tautology, $\vdash \mathbf{B}(\alpha \vee \neg\alpha)$. Notice $\vdash \alpha \vee \neg\alpha \equiv \neg(\alpha \wedge \neg\alpha)$. By the modus ponens between $\mathbf{B}$ and $\mathbf{B^c}$, we have $\vdash \mathbf{B}(\alpha \vee \neg\alpha) \wedge \mathbf{B^c}((\alpha \wedge \neg\alpha) \vee \gamma) \supset \mathbf{B^c}\gamma$. Hence $\vdash \mathbf{B^c}((\alpha \wedge \neg\alpha \wedge \beta) \vee \gamma) \supset \mathbf{B^c}\gamma$. ■

References

[Bel75] N. D. Belnap. A useful four-valued logic. In J. M. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*, 1975.

[BKMS92] C. Baral, S. Kraus, J. Minker, and V. S. Subrahmanian. Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, 8:45–71, 1992.

[dC74] N. C. A. da Costa. On the theory of inconsistent formal systems. *Notre Dam J. of Formal Logic*, 15:497–510, 1974.

[Doy79] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12:231–272, 1979.

[FH88] R. Fagin and J. Halpern. Belief, awareness, and limited reasoning. *Artificial Intelligence*, 34:39–76, 1988.

[FHV90] R. Fagin, J. Halpern, and M. Y. Vardi. A nonstandard approach to the logical omniscience problem. In *Proceedings of 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, 1990.

[FKUV86] R. Fagin, G. M. Kuper, J. D. Ullman, and M. Y. Vardi. Updating logical databases. *Advances in Computing Research*, 3:1–18, 1986.

[G8̈8] Peter Gärdenfors. *Knowledge in Flux - Modeling the dynamics of epistemic states.* MIT press, 1988.

[Gin88] M. L. Ginsberg. Multivalued logics: A uniform approach to reasoning in artificial intelligence. *Computational Intelligence*, 4:265–316, 1988.

[GM88] Peter Gärdenfors and D. Makinson. Revision of knowledge systems using epistemic entrenchment. In M. Vardi, editor, *Proceedings of 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*, 1988.

[GMR92] G. Grahne, A. O. Mendelzon, and Ray Reiter. On the semantics of belief revision systems. In *Proceedings of 4th Conference on Theoretical Aspects of Reasoning about Knowledge*, 1992.

[Hal86] J. Halpern. Reasoning about knowledge: an overview. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge*, 1986.

[HM85] J. Halpern and Y. O. Moses. A guide to the modal logics of knowledge and belief. In *Proceedings IJCAI-85*, 1985.

[IS] Y. E. Ioannidis and T. K. Sellis. Supporting inconsistent rules in database systems. To appear in the Int. Journal of Intelligent Information Systems. An earlier version appeared under the title "Conflict Resolution of Rules Assigning Values to Virtual Attributes" in Proceedings of the 1989 ACM-Sigmod Conference, pages 205-214, 1989.

[KL89] Michael Kifer and E. L. Lozinskii. Ri: A logic for reasoning with inconsistency. In *Proc. 4th Symposium on Logic in Computer Science*, 1989.

[KM91] H. Katsuno and A. O. Mendelzon. Propositional knowledgebase revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[Lak90] Gerhard Lakemeyer. Models of belief for decidable reasoning in incomplete knowledge base. Ph.D thesis, Univ. of Toronto, Dept. of Computer Science, 1990.

[Lev] Hector J. Levesque. A logic of implicit and explicit belief. FLAIR Texh. Rept. 32, Fairchld Lab. for AI Research, Palp Alto. A preliminary version appears in Proceedings of the 4th National conference of the American Association for Artificial Intelligence, pages 198-202,1984.

[Lev90] Hector J. Levesque. All I know: a study in autoepistemic logic. *Artificial Intelligence*, 42:263–309, 1990.

[Lin] J. Lin. Consistent belief reasoning in the presence of inconsistency. KRR-TR-93-1, Dept. of Computer Science, Univ. of Toronto, Toronto M5S 1A4.

[Lin87] Fangzhen Lin. Reasoning in the presence of inconsistency. In *Proceedings of the 6th National conference of the American Association for Artificial Intelligence*, 1987.

[LMa] J. Lin and A. O. Mendelzon. Merging the knowledge of multiple agents. Forthcoming.

[LMb] J. Lin and A. O. Mendelzon. On inconsistency elimination in knowledge bases. Forthcoming.

[Pri79] G. Priest. Logic of paradox. *Journal of Philosophical Logic*, 8:219–241, 1979.

[Pri91] Graham Priest. Minimally inconsistent LP. *Studia Logica*, L2:321–331, 1991.

[RB80] N. Rescher and R. Brandom. *The logic of inconsistency.* Billing & Sons, Ltd.,Oxford, UK, 1980.

[Roz89] L. I. Rozonoer. On interpretation of inconsistent theories. *Information Sciences*, 47:243–266, 1989.

[Var86] M. Vardi. On epistemic logic and logical omniscience. In *Proc. of the Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 293–305, 1986.