

ITERATED BELIEF CHANGE

ADITYA K. GHOSE¹ AND PABLO O. HADJINIAN

Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, T6G 2H1

ABDUL SATTAR

School of Computer and Information Technology, Griffith University, Nathan, Brisbane, QLD. 41111, Australia

JIA-H. YOU AND RANDY GOEBEL

Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada, T6G 2H1

Most existing formalizations treat belief change as a single-step process, and ignore several problems that become important when a theory, or belief state, is revised over several steps. This paper identifies these problems, and argues for the need to retain all of the multiple possible outcomes of a belief change step, and for a framework in which the effects of a belief change step persist as long as is consistently possible. To demonstrate that such a formalization is indeed possible, we develop a framework, which uses the language of PJ-default logic (Delgrande and Jackson 1991) to represent a belief state, and which enables the effects of a belief change step to persist by propagating *belief constraints*. Belief change in this framework maps one belief state to another, where each belief state is a collection of theories given by the set of extensions of the PJ-default theory representing that belief state. Belief constraints do not need to be separately recorded; they are encoded as clearly identifiable components of a PJ-default theory. The framework meets the requirements for iterated belief change that we identify and satisfies most of the AGM postulates (Alchourrón, Gärdenfors, and Makinson 1985) as well.

Key words: knowledge representation, belief revision, nonmonotonic reasoning.

1. INTRODUCTION

Belief change, the process by which a rational agent acquires new beliefs or retracts previously held ones, is a crucial component of intelligent behavior. Most existing formalizations view it as a single-step process, mapping from one set of beliefs to another (Alchourrón et al. 1985; Gärdenfors 1988; Nebel 1989). Several new challenges emerge, however, when one considers belief change over several steps. The process of belief change can, in general, have many possible outcomes. With a commitment to producing one unique outcome, most existing formalizations use some form of ordering on the beliefs to select some subset of the possible candidate theories, as in (Gärdenfors and Makinson 1988; Gärdenfors 1988), or combine all of the candidate theories to obtain the unique outcome, as in (Nebel 1989). Both approaches have undesirable consequences. In the first case, existing theories do not describe how the ordering relation on beliefs changes as a result of a belief change step. The notion of *iterated belief change* is thus not supported. Also, by requiring that all beliefs are ranked relative to each other (as in Gärdenfors and Makinson 1988), these approaches make unduly strong demands on the amount of information that must be available. In the second case, too much information is lost as a result of combining the mutually incompatible outcomes into a single theory. A related issue is the question of belief persistence. Common sense dictates that the effects of belief change persist as long as there is no evidence to the contrary. However, as a consequence of not recording changes in which beliefs are retracted, the effects of belief change do not persist over iterated steps in most existing formalizations; we call this the *problem of nonpersistent beliefs*.

Motivated by these pitfalls in the state of the art, we develop an alternative framework that obviates the need for selecting among the multiple possible outcomes of a belief change step, does not discard potentially useful information, and formalizes our intuition of belief persistence as well. First, we commit to representing a belief state as a collection of theories,

¹Current affiliation: School of IT and Computer Science, University of Wollongong, NSW 2522 Australia.

so that belief change becomes a mapping from one collection of theories to another. Given that each belief state corresponds to a possibly inconsistent and incomplete picture of the world, nonmonotonic formalisms are obvious candidates for representing such a state of affairs. In this paper, we use the language of PJ-default logic (Delgrande and Jackson 1991) to represent each belief state, such that the collection of theories at that belief state is the set of extensions of the corresponding PJ-default theory. However, any suitable nonmonotonic formalism could be used as the belief representation language using the general principles we describe here. We achieve persistent belief change by recording and maintaining belief constraints. These are of two types: constraints that enforce that some beliefs be necessarily held, and others that enforce that some beliefs may necessarily not be held. The intuition behind our revision operator is that if the new information logically contradicts the original set of beliefs, then its effects are two-fold (Sattar 1991). Some of the original beliefs are not affected at all; these remain as part of the set of facts of the corresponding PJ-default theory. Other beliefs may be contradicted, or brought into question; these beliefs are demoted to the status of tentative beliefs. This is achieved by removing these beliefs from the set of facts and incorporating them as default rules. A three-fold characterization of beliefs is thus introduced. For a PJ-default theory (W, D) representing a belief state, the set of facts W corresponds to the *necessary beliefs*, the set of consequents of the default rules in D corresponds to the set of *tentative beliefs*, while the negations of the conjuncts in the “semi-normal” part of the justification of each PJ-default rule correspond to the set of *necessary disbeliefs*. By addressing pragmatic concerns, such as belief persistence over iterated belief change steps, our framework represents a first step toward bridging the gap between theory and practice in this area. By factoring out the process of choosing amongst the multiple possible outcomes of a belief change step (typically based on orderings on beliefs Gärdenfors and Makinson 1988, Gärdenfors 1988; Nebel 1989, 1991), our approach permits the dynamic prioritization of beliefs. In other words, different orderings could be used at different times to choose among the multiple possible outcomes without causing inconsistencies or unreasonable results. Our approach can also be viewed as an account of how a nonmonotonic theory evolves into a more accurate representation of the world through the process of belief change.

This paper is organized as follows. Section 2 describes two existing formalizations of belief change. In Section 3, we identify some problems with existing frameworks; specifically the problem of nonpersistent beliefs and the absence of a description of how an ordering relation on beliefs should change as a result of a belief revision. Section 4 provides the details of our new framework. Section 5 describes the formal properties of our framework, specially with respect to some well-established rationality criteria. Section 5 discusses the relation of our work to other accounts of belief change.

Preliminary versions of the material in Sections 3 and 4 have appeared in Ghose et al. (1993).

2. BELIEF CHANGE: EXISTING FRAMEWORKS

Alchourrón et al. (1985) have undertaken a systematic study of the dynamics of belief change, resulting in what is currently popularly known as the AGM framework for belief change (Alchourrón et al. 1985; Gärdenfors 1988; Gärdenfors and Makinson 1988). In the AGM framework, the belief state of an agent is represented by a deductively closed, logically consistent set of propositional sentences called a *belief set*. They define three kinds of belief change operations: *expansion*, in which the new belief being added is guaranteed to be consistent with the existing body of beliefs; *contraction*, in which an existing belief is retracted; and *revision*, in which a new belief, which may possibly be inconsistent with existing beliefs,

is added. The operations of contraction and revision can be defined in terms of each other, as shown by the *Levi identity* below (here, K_A^* , K_A^- , and K_A^+ denote, respectively, the revision, contraction, and expansion of K with A):

$$K_A^* = (K_{\neg A}^-)^+.$$

The *Harper identity* (1975) ($K_A^- = K_{\neg A}^* \cap K$) similarly defines contraction in terms of revision. We describe only contraction operators in this section, because the corresponding revision operators follow via the Levi identity. Alchourrón et al. (1985) define a set of rationality postulates for each of the operations of expansion, revision, and contraction. The postulates for contraction are listed below.

1. For any sentence A and any belief set K , K_A^- is a belief set.
2. $K_A^- \subseteq K$.
3. If $A \notin K$, then $K_A^- = K$.
4. If $\not\models A$, then $A \notin K_A^-$.
5. If $A \in K$, then $K \subseteq (K_A^-)_A^+$.
6. If $\models A \leftrightarrow B$, then $K_A^- = K_B^-$.
7. $K_A^- \cap K_B^- \subseteq K_{A \wedge B}^-$.
8. If $A \notin K_{A \wedge B}^-$, then $K_{A \wedge B}^- \subseteq K_A^-$.

Postulate 1 requires that beliefs be represented in the same form before and after a belief change step. Postulate 2 requires that no new beliefs be held as a result of a contraction. Postulate 3 requires that if the belief to be contracted is not held, then no change should be made. Postulate 4 requires that every contraction operation succeed, unless the belief being contracted is a logical truth. Postulate 5 is the principle of recovery, which requires that if a belief held in a given belief state is retracted and then added back to the belief state, the outcome contains the initial belief state, i.e., the initial belief state is recovered. Postulate 6 is the principle of irrelevance of syntax, which requires that the outcome of a contraction operation be independent of the syntactic form of the beliefs being contracted. Postulate 7 requires that the retraction of a conjunction of beliefs should not retire any beliefs that are common to the retraction of the same belief set with each individual conjunct. Postulate 8 requires that, when retracting the conjunct of two beliefs A and B forces us to give up A , then in retracting A , we do not give up any more than in retracting the conjunction of A and B .

Partial meet contraction (Alchourrón et al. 1985) is a representative AGM operator, which satisfies all of the rationality postulates, and which appears to be more reasonable, compared to the other AGM operators by virtue of not discarding too many beliefs and also by not forcing commitment to a truth value for every propositional letter. Let the *removal* of x from A , denoted by $A \downarrow x$, be defined as follows:

$$A \downarrow x = \{B \subseteq A \mid B \not\models x, \forall C : B \subset C \subseteq A \Rightarrow C \models x\}.$$

Let S be a *selection function* that selects a nonempty subset of $K \downarrow x$ (provided $K \downarrow x$ is nonempty, \emptyset otherwise). The partial meet contraction operator \neg_p is defined as follows:

$$K_x^{-p} = \bigcap S(K \downarrow \neg x).$$

Let $M(K)$ stand for the family of all the sets $K \downarrow x$, where x is any proposition in K that is not logically valid. Let \leq be a relation defined on $M(K)$. Let

$$S(K \downarrow x) = \{K' \in K \downarrow x \mid K'' \leq K' \text{ for all } K'' \in K \downarrow x\}.$$

Any partial meet contraction operator for which the selection function S is defined in this manner, and for which the relation \leq is transitive, satisfies all the AGM postulates for contraction (Gärdenfors 1998). A contraction operator, which uses a special class of total orderings (called *epistemic entrenchment*), defined on the entire language, to decide which beliefs to retain is shown to satisfy all eight of the AGM postulates for contraction, and to be equivalent to partial meet contraction. Let $x \leq_K y$, where \leq_K is the epistemic entrenchment relation associated with the belief set K , and denotes that x is at most as entrenched as y . The relation \leq_K must satisfy the following conditions:

1. If $x \leq_K y$ and $y \leq_K z$, then $x \leq_K z$.
2. If $x \models y$ then $x \leq_K y$.
3. For any x and y , $x \leq_K x \wedge y$ or $y \leq_K x \wedge y$.
4. When $K \neq K_\perp$, $x \notin K$ iff $x \leq_K y$ for all y .
5. If $y \leq_K x$ for all y , then $x \models$.

The epistemic entrenchment relation uniquely determines a contraction operation via the following definition:

$$y \in K_x^- \text{ iff } y \in K \text{ and either } x \leq_K x \vee y \text{ or } \models x.$$

Nebel (1991) discusses contraction operators on *beliefbases*, which are finite sets of sentences instead of infinite deductively closed belief sets. The motivations for defining belief change on belief bases are twofold. First, operators defined on the bases of belief are computationally viable (they do not have to operate on infinite sets). Second, belief change operations on belief bases permit reason maintenance, while those on belief sets do not. The base contraction operator \simeq defined as

$$B \simeq x = \begin{cases} (\bigvee_{C \in (B \downarrow x)} C) \wedge (B \vee \neg x) & \text{if } \not\models x \\ B & \text{otherwise} \end{cases}$$

satisfies most, but not all, of the AGM postulates. Note that the term $(B \vee \neg x)$ ensures that the original belief base reappears whenever x becomes true. The corresponding revision operator can be defined, as before, via the Levi identity. In a later paper, Nebel (1991) generalizes this operator to prioritized belief bases. We will not describe this operator in detail here, because the problems we will point out apply in either case.

3. ITERATED BELIEF CHANGE

This work has been motivated by the need to address the inadequacy of most existing belief revision systems in handling the dynamics of iterated belief change, and by the need to develop more expressive systems, which have significant computational advantages as well. The inadequacy of existing systems with regard to iterated belief change stems from two sources. First, contraction operations do not persist in these systems. Second, these systems provide no account of how selection functions (or equivalently, ordering relations on beliefs) evolve over iterated belief change steps.

We motivate our discussion of the dynamics of belief change by arguing for the need for the effects of a belief change step to *persist*. Belief change operators should satisfy a *persistence postulate* requiring that the effects of every belief change persist until there is new evidence to indicate otherwise. A contracted belief should not reappear in a belief set unless there is new evidence accumulated since the contraction that requires this. Similarly, a newly added belief should remain in the belief set until new evidence obtained since the

addition warrants its removal. We use the notion of *evidence* here to denote belief inputs that an agent obtains from the world, together with their prioritization relative to other beliefs. Thus, a newly added belief will be treated as sufficient evidence to contradict a previously added belief only if it brings the previously held belief into question, and has a higher priority than the previously held belief as well.

Example 1. Consider the AGM framework. Let $\{b, f\}$ be the alphabet of our language. Let the initial belief state be $K0 = Cn\{b \rightarrow f\}$. After contracting f , let the outcome be $K1 = Cn\{b \rightarrow f\}$ (since f is not a consequence of the beliefs in $K0$, no change is made to $K0$). Revising $K1$ with b results in the belief state $Cn\{b, b \rightarrow f\}$. Thus, the agent starts believing f again, although the only new information (the belief b) obtained since being told to retract the belief f does not in itself require that f be believed again. A more detailed analysis reveals that when $K1$ is revised with b , three different entities need to be considered:

- A: $b \rightarrow f$ and its consequences are believed.
- B: f is retracted.
- C: b is believed.

Prioritizing these entities informally using a relation $>$, where $x > y$ denotes that x has higher priority over y , a variety of outcomes is possible.

- If $C > A > B$ then $K1_b^* = Cn\{b, b \rightarrow f\}$.
- If $C > B > A$ then $K1_b^* = Cn\{b\}$.
- If $A > B > C$ then $K1_b^* = Cn\{b \rightarrow f\}$.
- If $A > C > B$ then $K1_b^* = Cn\{b, b \rightarrow f\}$.

We do not list all the possibilities here, but clearly the three distinct entities and their relative prioritization need to be considered in generating an outcome. For instance, let b denote that Tweety is a bird and f denote that Tweety flies. Then, starting with a belief in an instance of the rule “birds fly,” after retracting the belief that Tweety flies and then being told that Tweety is a bird, it seems reasonable to remove the “birds fly” rule from the status of a first-class belief, given new information about Tweety’s flying ability and the fact that Tweety is a bird, giving a final belief state $Cn(b)$. This corresponds to the case where $C > B > A$.

It may be argued that $Cn\{b, b \rightarrow f\}$ is a reasonable resultant belief state, because the agent has been told to believe b in a belief state that contains only the consequences of the belief $b \rightarrow f$. However, $\{Cnb \rightarrow f\}$ is an inadequate representation of the belief state that results from contracting f from the initial belief state, because the retraction of f has not been recorded. Note that it is possible that the belief f could be held in the final belief state even if the contraction of f is somehow recorded in the belief state (this would happen if the agent assigned a higher priority to the belief in $b \rightarrow f$ and b over the retraction of f). What is crucial, however, is that the retraction of f should be considered in deciding the subsequent belief states instead of it being forgotten from the next step onward. In general, existing frameworks always record every revision step, but never record a contraction step.

It may be argued that contraction is never an independent belief change operation and is only useful as part of a revision operation (as in the Levi identity). We disagree, because it is important to admit a belief change operation that causes a belief to become undefined, or unknown, without requiring that the negation of the belief be held. Contraction is redundant only in the case of *complete theories* (i.e., theories that must commit to p or $\neg p$ for any p). An independent contraction operation can be crucial in a variety of situations, such as those in which negative examples are accumulated in an incremental inductive logic programming setting. We therefore believe that contraction should be accorded the status of an independent

belief change operation, at par with revision, and that contractions should be recorded in the representation of a belief state in the same way as a revision operation.

Nebel's framework, as the example below shows, suffers from an identical problem.

Example 2. Consider Nebel's belief base revision (Nebel 1989). Let $B_0 = \{a \rightarrow b\}$. $B_1 = B_0 \simeq b = \{a \rightarrow b, \neg b \vee B_0\}$, where we take B_0 to denote the conjunction of its elements. Revising B_1 with a gives $B_2 = \{a, a \rightarrow b, \neg b \vee B_0\}$. The contraction of b does not persist (and the spurious belief b appears in the deductive closure of B_2) due to the lack of an explicit record of belief change steps.

Several authors have pointed out the absence of any commitment on how the selection function evolves as a result of a belief change step in the AGM framework (Hansson 1991; Nayak 1993; Spohn 1988). An epistemic entrenchment ordering, for instance, must be defined relative to a specific belief set K . While such an ordering is sufficient for uniquely determining the outcome of a revision or contraction operation on K , there is no prescription of how a subsequent belief change operation on, say K_x^* may be performed, because we do not know what the epistemic entrenchment relation associated with K_x^* is. Also, the notion of epistemic entrenchment can contradict the requirement, in the AGM framework, that every belief change operation succeed (as set forth in the *success postulates* for revision and contraction). More entrenched beliefs may be discarded as a result of revising with a less entrenched belief, simply because no connection is made between the entrenchment relation and the new evidence.

Nebel's (1991) *epistemics relevance* ordering of beliefs suffers from an identical problem. Clearly, we need a framework in which belief change steps are recorded and *persist* and that treats priority relations on evidence and on existing beliefs in a uniform way. We present such a framework in the next section.

4. AN ALTERNATIVE FRAMEWORK FOR BELIEF REVISION

Our formalization of belief change is motivated by the following observations:

- A belief state is best represented as a collection of theories. Given that *minimal change* is a guiding principle in belief revision, it could be argued that instead of selecting some outcomes of the belief change and discarding the others (thus losing potentially useful beliefs), all outcomes should be retained, provided that there exists a compact and elegant way of representing these multiple possible outcomes. Such representation languages exist; nonmonotonic formalisms are immediate candidates for compactly representing the possibly inconsistent and incomplete picture of the world that each belief state corresponds to. The observation that real-life agents typically have incomplete (and sometimes inconsistent) knowledge about the world is independent justification for choosing such an approach.
- The persistence postulate suggests that an explicit record of belief change steps should be maintained. We achieve this by maintaining a set of *belief constraints* (these may be of two kinds: constraints specifying beliefs that must necessarily be held, and constraints specifying beliefs that must necessarily not be held), which may be viewed as the integrity constraints of a belief system; every theory constituting a belief state must respect them.
- Beliefs originally held to be true can become tentative as a result of belief change if they are contradicted or brought into question (intuition is that a belief becomes questionable

if it is not in every possible outcome of the belief change step) by new evidence. In syntactically oriented nonmonotonic formalisms, this can be viewed as a process of demotion from a fact to a default.

We demonstrate that a formalization that satisfies most of the common sense requirements for belief change, as given by the AGM postulates as well as by the persistence postulate identified in the previous section, is indeed possible. To prove our point, we use the language of PJ-default logic (Delgrande and Jackson 1991) to represent a belief state. PJ-default logic is a variant of the default logic in which default rules are restricted to be prerequisite-free and semi-normal (i.e., a PJ-default rule is of the form $\frac{\beta}{\gamma}$ such that $\beta \models \gamma$). PJ-default logic improves over Reiter's (1980) default logic by avoiding cases where Reiter's logic is too weak, preventing the derivation of "reasonable" conclusions (such as in the disjunctive default problem) as well as cases where Reiter's logic is too strong, permitting the derivation of unwanted conclusions (for a detailed discussion of these issues, see Delgrande and Jackson 1991). This approach has other useful properties as well, such as semimonotonicity, guaranteed existence of extensions, weak orthogonality of extensions, and a constructive definition for extensions. PJ-default extensions are defined as follows:

Definition 1. Let (W, D) be a prerequisite-free seminormal default theory. Define:

$$\begin{aligned} E_0 &= (E_{J_0}, E_{T_0}) = (Cn(W), Cn(W)) \\ E_{i+1} &= (E_{J_{i+1}}, E_{T_{i+1}}) = (Cn(E_{J_i} \cup \{\beta \wedge \gamma\}), Cn(E_{T_i} \cup \{\beta\})) \end{aligned}$$

where

$$\begin{aligned} i &\geq 0, \\ \frac{(\beta \wedge \gamma)}{\beta} &\in D, \\ \neg(\beta \wedge \gamma) &\notin E_{J_i}. \end{aligned}$$

Then E is a PJ-extension for (W, D) iff

$$E = (E_J, E_T) = (\bigcup_{i=0}^{\infty} E_{J_i}, \bigcup_{i=0}^{\infty} E_{T_i}).$$

In the remaining part of the paper, whenever we refer to an extension, we refer to the E_T part of a PJ-extension.

Ghose and Goebel (1992) have earlier defined a belief change framework in which a belief state is represented as a potentially inconsistent set of sentences, together with a partial order on these sentences. An operator is defined that identifies consistent subsets of this set of sentences, that respect the partial order as well as some set of disbelief constraints. A translation from PJ-default theories to this framework is defined such that the process of identifying PJ-default extensions is shown to be equivalent to the process of identifying consistent subsets of sentences using the operator mentioned above, with a partial order that assigns higher priority to sentences obtained from W (given a PJ-default theory (W, D)) than sentences obtained from D and with the set of disbelief constraints consisting of the conjunction of the negations of the justifications of each PJ-default rule.

Formally, we define a *belief state* to be a set of theories $\{Th_1, Th_2, \dots\}$. The *necessary belief set* BC_{belief} , for a given belief state, is a set of sentences such that, for every $x \in BC_{belief}$ and for all Th_i in that state, $Th_i \models x$. The *necessary disbelief set* $BC_{disbelief}$, for a given belief state, is a set of sentences such that, for every $\neg x \in BC_{disbelief}$ and for all Th_i in that state, $Th_i \not\models x$. Note that $BC_{disbelief}$ contains the *negations* of beliefs that may not be held. A *tentative belief* is a belief that is an element of some, but not all Th_i in that belief state. We

define the *belief constraint set* $BC = BC_{belief} \cup BC_{disbelief}$. Updating the belief constraint set requires identifying maximal parts of a constraint that is consistent with some others. We do not present a formal notion of a part of constraint here, but the following example should make our intuitions clear.

Example 3. Consider retracting both a and b in a single step from a belief state. This corresponds to ensuring that the formula $a \vee b$ is not believed in this state. We therefore add a necessary disbelief constraint $\neg a \wedge \neg b$. Let the next belief change step be a revision with a , which corresponds to adding a necessary belief constraint a . Clearly, it is impossible to enforce a disbelief in both a and b , and belief in a at the same time. The two constraints cannot be satisfied at the same time. One option is to delete the disbelief constraint $\neg a \wedge \neg b$ because it contradicts the newly acquired constraint. However, if we do so, we shall allow the belief b to be held. All that we have been told because both a and b were retracted, is that a is to be believed again. It makes better sense to retain the requirement for disbelief in b . Thus, we should retain the maximal part of the disbelief constraint $\neg a \wedge \neg b$ that is consistent with the new belief constraint a , i.e., the disbelief constraint $\neg b$.

A set-theoretic representation of each constraint such that each conjunct in a constraint is clearly identifiable as a set element would enable us to identify maximal subsets that are consistent with other constraints. Representing belief constraints in clausal form satisfies these requirements. To be able to distinguish between the elements of BC_{belief} and $BC_{disbelief}$, we adopt a syntactic convention such that every disbelief constraint ϕ is written as $-\phi$. Thus, constraints requiring necessary belief in a and $a \rightarrow b$ will be written as $\{\{a\}\}$ and $\{\{-a, b\}\}$, respectively, while a constraint requiring disbelief in $(c \vee d) \wedge e$ will be written as $-\{\{-c, \neg e\}, \{-d, \neg e\}\}$. Belief change will be viewed as the process of adding a new belief constraint (in the case of revision, a necessary belief constraint, and in the case of contraction, a necessary disbelief constraint).

We require that the belief constraint set be totally ordered; we refer to this total order $<$ as the *constraint prioritization* (we write $x_j < x_i$ if constraint x_i has a higher priority than constraint x_j). This is similar to the orderings used in a variety of belief change frameworks, including the AGM epistemic entrenchment ordering, but there are significant differences. Whereas epistemic entrenchment requires that all beliefs be prioritized, we require that only the current set of belief constraints be prioritized. We will see later that the size of the belief constraint set can typically be expected to be much smaller than the size of the theories constituting a belief state. Also, the size of the belief constraint set does not grow with time and may shrink as belief constraints are discarded (as a result, for example, of newer constraints contradicting them). This also represents a more principled approach to prioritizing beliefs, because it does not suffer from the contradictions in belief prioritization in the AGM framework pointed out in the previous section. Note also, that, unlike AGM epistemic entrenchment, this ordering does not uniquely determine which subset of the possible outcomes is finally selected. It only determines what the new set of belief constraints should be. To draw a database analogy, the constraint prioritization ranks only the current integrity constraints, whereas epistemic entrenchment requires that every assertion in the database be ranked. Several obvious heuristics suggest themselves as candidate constraint prioritization policies in the absence of any other information on the relative reliability of the belief constraints. In the case of *revision* (in the sense of Katsuno and Mendelzon 1991, i.e., for belief change in static worlds), a constraint prioritization base on the recency of the belief inputs appears to be appropriate. In the case of *update* (Katsuno and Mendelzon 1991) (i.e., belief change in dynamic worlds), one might choose to accord the highest priority to physical laws at all times, because these are never questioned or discarded.

The process of belief change involves two steps. First, the belief constraint set is updated. Based on the updated belief constraint set, the PJ-default theory representing the previous belief state is modified.

Definition 2. A belief constraint Ω (where $\Omega = \phi$ if Ω requires necessary belief in ϕ and $\Omega = \neg\phi$ if Ω requires necessary disbelief in ϕ) is said to be compatible with a set of belief constraints BC (we write $\Omega \cup BC$ is compatible), if and only if, for all, $x_i \in BC_{belief}$ and all $\neg y_j \in BC_{disbelief}$, $\Omega \wedge (\bigwedge_i x_i) \wedge (\bigwedge_j \neg y_j)$ is satisfiable.

We must be able to identify subsets of an individual belief constraint (viewing each constraint as a set of clauses) that are compatible with a set of belief constraints. The operator \uparrow that identifies such subsets is defined below.

Definition 3. Let bc be a belief constraint and BC be a set of belief constraints. Then

$$bc \uparrow BC = \{x \subseteq bc \mid (x \cup BC \text{ is compatible}) \wedge (\forall x' \text{ such that } x \subset x' \subseteq bc, x' \cup BC \text{ is incompatible})\}.$$

Updating the set of belief constraints involves starting with the constraint of highest priority and working downwards, adding as many constraints (or parts of constraints) of lower priority as can be *compatibly* added. In the definition below, we assume that BC_{old} is the old set of constraints, Ω is the new constraint being added in the current belief change step, and $BC_{inter} = BC_{old} \cup \Omega$. We assume that the constraint prioritization relation $<$ is updated to reflect the ranking of Ω relative to the other constraints. Also, every bc_i belongs to BC_{inter} and $bc_j < bc_k \leftrightarrow k < j$ (i.e., constraints of a higher priority have a smaller subscript).

$$\begin{aligned} BC_{new} &= \{Y \subseteq BC_{inter} \mid Y = \bigcup_{i \geq 1} Y_i, \\ &\forall i \geq 1 ((Y_i = \{bc_i\}) \wedge (bc_i \cup (\bigcup_{j=1}^{i-1} Y_j) \text{ is compatible})) \\ &\vee ((Y_i = \bigcap (bc_i \uparrow (\bigcup_{j=1}^{i-1} Y_j \cup \Theta)) \wedge (bc_i \cup (\bigcup_{j=1}^{i-1} Y_j) \text{ is incompatible})))\} \end{aligned}$$

The new constraint prioritization relation is the subset of the old $<$ defined on the elements of BC_{new} . If a subset of constraint belonging to BC_{old} is retained in BC_{new} , then it retains its position in the new constraint prioritization relation.

In the above definition here, note that in the process of collecting constraints from the highest to the lowest priority, if a constraint turns out to be incompatible with the set of higher priority constraints collected so far, we identify maximal subsets of the constraint (given the representation of a constraint as a set of clauses) that are compatible with the collected constraints, and add the intersection of these subsets to the set of collected constraints. The following examples clarify this.

Example 4. Let $BC_{old} = \{\{\neg a, b\}, \{\neg b\}\}$ and let $<$ be such that $\{\{\neg b\}\} < \{\{\neg a, b\}\}$. Let $\Omega = \{a\}$. Thus $BC_{inter} = \{\{\neg a, b\}, \{\neg b\}, \{a\}\}$ and the updated prioritization relation $<$ is such that $\{\{\neg b\}\} < \{\{\neg a, b\}\} < \{a\}$. Then $BC_{new} = \{\{a\}, \{\neg a, b\}\}$.

Example 5. Let the initial belief state have a single constraint requiring disbelief in $a \vee b$. Thus $BC_{old} = \{-\{\neg a\}, \{\neg b\}\}$ and the relation $<$ is empty. Let $\Omega = \{a\}$ be added at the highest priority level. $BC_{inter} = \{\{a\}, -\{\neg a\}, \{\neg b\}\}$ and the updated $<$ relation is such that $-\{\{\neg a\}, \{\neg b\}\} < \{a\}$. Then $BC_{new} = \{-\{\neg b\}, \{a\}\}$ and the resulting $<$ relation is such that $-\{\{\neg b\}\} < \{a\}$. Thus, while a constraint requiring disbelief in both a and b is not

compatible with a necessary belief constraint in a , a subset that requires disbelief in b only is compatible.

Using PJ-default logic as the belief representation language enables us to represent belief constraints in a PJ-default theory. The set of facts W of a PJ-default theory corresponds to the beliefs that the agent is constrained to necessarily hold, because W will be contained in every consistent belief set corresponding to that belief state (i.e., every PJ-extension). Because every PJ-default rule is of the form $\frac{\beta \wedge \gamma}{\beta}, \neg \gamma$ corresponds to the theory that the agent is constrained to necessarily disbelieve. The process of demoting a belief that has been contradicted, or made questionable (meaning that it is no longer in every possible outcome of that belief change step), as a result of a belief change step to the status of a tentative belief involves removing a formula from W and adding a new PJ-default rule containing this formula as its consequent to D . Discredited beliefs are thus never totally discarded in our framework in anticipation of future situations in which these beliefs could be consistently held again. Belief change in our framework thus involves mapping one PJ-default theory to another. The possible consistent belief sets that may be held in a given belief state correspond to the extensions of the PJ-default theory representing that belief state.

The new PJ-default theory (W', D') is obtained from the previous belief state (W, D) and the new belief constraint set BC_{new} as follows

$$\begin{aligned} W' &= BC_{Belief_{new}} \\ D' &= \left\{ \frac{\delta_i \wedge (\bigwedge BC_{Disbelief_{new}})}{\delta_i} \mid \delta_i \in (W - W') \right\} \cup \\ &\quad \left\{ \frac{\beta_i \wedge (\bigwedge BC_{Disbelief_{new}})}{\beta_i} \mid \frac{\beta_i \wedge \phi_i}{\beta_i} \in D \right\}. \end{aligned}$$

Here $\bigwedge BC_{Disbelief_{new}}$ stands for the conjunction of all the elements of $BC_{Disbelief_{new}}$. Thus, if $BC_{Disbelief_{new}} = \{-\{\neg b\}, -\{\neg a\}\}$, then $\bigwedge BC_{Disbelief_{new}} = \neg b \wedge \neg a$.

Note that the justification for every default rule in PJ-default theories representing belief states in our framework is identical and corresponds to the set of necessary disbelief constraints. We define the notion of *uniform* default theories as follows:

Definition 4. A semi-normal default theory (W, D) is said to be uniform if for any two default rules $\frac{\alpha_i \wedge \beta_i \wedge \gamma_i}{\beta_i}, \frac{\alpha_j \wedge \beta_j \wedge \gamma_j}{\beta_j} \in D$, $\gamma_i = \gamma_j$ (if $|D| = 1$, then the theory is trivially uniform).

Clearly, every PJ-default theory representing a belief state in our framework is *uniform*.

We require that the dummy default rule $\frac{\top}{\top}$ be an element of D for every (W, D) representing a belief state. This is to enable us to record necessary disbelief constraints even if there are no tentative beliefs.

Example 6. Let the initial belief state be given by the uniform PJ-default theory (W, D) where $W = \{a, a \rightarrow b, c, c \rightarrow d\}$ and $D = \{\frac{\top}{\top}\}$. Note that there are no necessary disbelief constraints at this point. Assume that the elements of W were obtained in a single belief change step. Thus, there is a single necessary belief constraint and $BC_{Belief_{old}} = \{\{a\}, \{-a, b\}, \{c\}, \{-c, d\}\}$.

Let us now retract the belief b .

- Belief Constraints Update

$$Let \{\{a\}, \{\neg a, b\}, \{c\}, \{\neg c, d\}\} \prec -\{\{\neg b\}\}.$$

$$BC_{Belief_{new}} = \{\{\{c\}, \{\neg c, d\}\}\}$$

$$BC_{Disbelief_{new}} = \{\{-\{b\}\}\}$$

- Default Theory Update

$$W' = \{c, c \rightarrow d\}$$

$$D' = \left\{ \frac{: a \wedge \Gamma}{a}; \frac{:(a \rightarrow b) \wedge \Gamma}{(a \rightarrow b)}; \frac{:\top \wedge \Gamma}{\top} \right\}, \quad \text{where } \Gamma = \neg b.$$

Note that, as a result of the contraction, a and $a \rightarrow b$ become *tentative beliefs*. Note also that there are two possible belief sets that may be consistently held in the final belief state (corresponding to the two extensions of (W', D')) given by $Cn(\{c, c \rightarrow d, a\})$ and $Cn(\{c, c \rightarrow d, a \rightarrow b\})$ and that the belief b is not held in either of these. The new constraint prioritization is $\{\{c\}, \{\neg c, d\}\} \prec -\{\{\neg b\}\}$.

Example 7. We refer to the initial example in Section 3 motivating the need to record contraction operations in the same way as revision operations. Let the initial belief state be given by

$$(W_1, D_1) = \left(\{b \rightarrow f\}, \left\{ \frac{:\top}{\top} \right\} \right)$$

$$BC_{Belief_1} = \{\{\{\neg b, f\}\}\}$$

The relation \prec_1 is empty.

The belief state obtained by contracting f and assigning the highest priority to the disbelief constraint in f is given by

$$(W_2, D_2) = \left(\{b \rightarrow f\}, \left\{ \frac{:\top \wedge \neg f}{\top} \right\} \right)$$

$$BC_{Belief_2} = \{\{\{\neg b, f\}\}, -\{\{\neg f\}\}\}$$

The relation \prec_2 is such that $\{\{\neg b, f\}\} \prec -\{\{\neg f\}\}$.

The belief state obtained by further revising with b and assigning this belief constraint the highest priority is given by

$$(W_3, D_3) = \left(\{b\}, \left\{ \frac{:(b \rightarrow f) \wedge \neg f}{(b \rightarrow f)}, \frac{:\top \wedge \neg f}{\top} \right\} \right)$$

$$BC_{Belief_3} = \{\{\{\neg b, f\}\}, -\{\{\neg f\}\}, \{\{b\}\}\}$$

The relation \prec_3 is such that $\{\{\neg b, f\}\} \prec -\{\{\neg f\}\} \prec \{\{b\}\}$.

Note that the contraction of f persists because f is in no extension of (W_3, D_3) .

By representing each belief state as a uniform PJ-default theory, we have factored out the question of which theory (extension) to choose as the currently operative set of beliefs from the process of belief change. The user, or agent, could thus employ a variety of techniques to

actually pick the currently operative set of beliefs. If priorities on beliefs are used to select the currently operative set of beliefs (as in AGM epistemic entrenchment), then our framework would permit dynamic prioritization of beliefs. In other words, different orderings could be used at different times to select theories without causing inconsistencies or unreasonable outcomes from the belief change process. The more crucial priority relation, however, is the one that determines the nature of the new belief state. In our case, this is the *constraint prioritization* relation. Unlike the AGM epistemic entrenchment relation and Nebel’s epistemic relevance ordering (which, like our constraint prioritization relation, determines the nature of the new belief state), we provide a clear prescription of how the relation evolves over iterated belief change steps and how the new evidence is integrated into this relation. An additional advantage with our framework is that we permit, in addition to the traditional operations of expansion, revision and contraction, the “undoing” of a contraction step; it simply involves the removal of the relevant necessary disbelief constraint.

4.1. Properties

It has become popular in recent times, to evaluate every new belief change operator against the yardstick of rationality provided by the AGM postulates for belief change, primarily because they seem to be the best formalization of the consensus view on the requirements an ideal belief change operator should satisfy. Our formalization cannot, however, be evaluated using the AGM postulates directly, for the following two reasons. First, the AGM postulates consider transitions between belief states represented as a single deductively closed propositional theory. Our operator maps between belief states represented as collections of theories (the multiple possible extensions of the PJ-default theories). Second, because the AGM postulates consider belief change as a single-step process, it is difficult to evaluate “rationality” over iterated belief change steps. It is possible, however, to articulate a reformulated version of these postulates, and show that our framework satisfies them under certain conditions. We formalize these conditions first.

Definition 5. Let a belief change operation result in the introduction of a new constraint x in a belief state with a constraint set denoted by BC . Let $y \in BC$ be such that y is incompatible with x and there is no $z \in BC$ such that $y < z$ and z is incompatible with x . The belief change operation is *imperative* iff $y < x$.

Thus, an imperative belief change operation introduces a constraint that has a higher priority than the existing constraint of the highest priority that is incompatible with it. Because our framework is general enough to permit any prioritization of the belief constraints, it is possible for a belief change operation to fail (in the case that there exists a belief constraint, with higher priority than the newly introduced constraint, which is incompatible with the newly introduced constraint. If a belief change operation is imperative, the operation is guaranteed to succeed. Because every belief change step in the AGM framework must succeed, our framework satisfies the AGM postulates only in the case of imperative belief change operations.

The second condition involves the prevention of beliefs that were previously suppressed by the existence of a disbelief constraint from reappearing in a belief state when this disbelief constraint is discarded.

Example 8. Consider a belief state given by $W = \{\}$ and $D = \{\frac{\neg a \wedge a \wedge (a \rightarrow b)}{\neg a}, \frac{\neg \top \wedge a \wedge (a \rightarrow b)}{\top}\}$. There is a single, empty, belief set corresponding to $Cn(\top)$, which is the only extension of (W, D) . Let the belief constraint set in this belief state consist of a single disbelief constraint— $\{\{a\}, \{\neg a, b\}\}$. Let the belief b be retracted from this belief state. We get a new belief constraint

set given by a single disbelief constraint— $\{\neg b\}$. The new belief state will be given by $W' = \{\}$ and $D' = \{\frac{: \neg a \wedge \neg b}{\neg a}, \frac{: \top \wedge \neg b}{\top}\}$. This default theory has one extension, $Cn(\neg a)$. Thus we get a belief state with a belief set containing $\neg a$ as result of contracting b from a belief state containing a single, empty, belief set. This clearly violates the AGM contraction postulate, which requires that the contracted belief set should be a subset of the original belief set, yet the behavior is perfectly rational. The tentative belief $\neg a$ reappears in a belief set as a consequence of the removal of the disbelief constraint that caused this tentative belief to be suppressed.

Clearly, only operations, which do not display such behavior, can be related to the AGM framework.

Definition 6. Let a belief change operation introduce a belief constraint x in a belief state (W, D) with a belief constraint set BC . The operation is stable iff there exists no disbelief constraint $y \in BC$. s.t $y \prec x$ and there exists $\frac{\beta \wedge \gamma}{\beta} \in D$ where $y \cup \beta$ is incompatible.

Before we state and prove the results relating our framework to the AGM postulates, we establish a connection between the THEORIST system developed by Poole, Goebel and Aleliunas (1987) and uniform PJ-default theories that simplifies the proofs. The THEORIST framework envisages a knowledge base comprising a set of closed formulas that are necessarily true, called *facts*, and a set of possibly open formulas that are tentatively true, called *hypotheses*. Default reasoning in this framework involves identifying *maximal scenarios* (or extensions), where a scenario consists of the set of facts together with some subset of the set of ground instances of the hypotheses that is consistent with the set of facts. The framework can be augmented with *constraints*, which are closed formulas such that every THEORIST scenario is required to be consistent with the set of constraints. Following Poole et al., we can present the following definition of a maximal scenario.

Definition 7. For a THEORIST system (F, H, C) where F is the set of facts, H is the set of hypotheses and C is the set of constraints, such that every element of F , H and C is a ground formula, a *maximal scenario* is a set $F \cup h$ such that $h \subseteq H$ and $F \cup h \cup C$ is satisfiable.

We present a translation that is simpler than the one presented in Delgrande and Jackson (1991) because we are dealing with uniform PJ-default theories rather than general ones. In the following, $S(F, H, C)$ refers to the set of maximal scenarios of the THEORIST system (F, H, C) . Also, $E(\Delta)$ refers to the set of extensions of the default theory Δ .

Definition 8. Let $(F_{(W,D)}, H_{(W,D)}, C_{(W,D)})$ denote the THEORIST translation of the uniform PJ-default theory (W, D) . Then

$$\begin{aligned} F_{(W,D)} &= W \\ H_{(W,D)} &= \left\{ \beta \mid \frac{\beta \wedge \gamma}{\beta} \in D \right\} \\ C_{(W,D)} &= \left\{ \gamma \mid \frac{\beta \wedge \gamma}{\beta} \in D \right\}. \end{aligned}$$

Theorem 1. Let $(F_{(W,D)}, H_{(W,D)}, C_{(W,D)})$ denote the THEORIST translation of the uniform PJ-default theory (W, D) . Then $S(F_{(W,D)}, H_{(W,D)}, C_{(W,D)}) = E((W, D))$.

Proof. The proof follows directly from the equivalence of the definitions of THEORIST maximal scenario computation and PJ-default extension computation, given that

$E_{J_i} = C_{(W,D)}$ at every step in the computation of extensions for a uniform PJ-default theory (W, D) .

Thus, the facts F correspond to the necessary belief constraints. The constraints C correspond to the necessary disbelief constraints in our system. The hypotheses H correspond to the tentative beliefs. The process of belief change can thus be equivalently defined as computing the new set of belief constraints, replacing F and C accordingly and demoting beliefs from F to H if necessary.

We interpret postulate 1 as one way of articulating the following *principle of categorical matching* stated by Gärdenfors and Rott (1993), which requires that the representation of a belief state after a belief change has taken place should be of the same format as the representation of the belief state before the change. For postulates 2 through 8, we reformulate every condition on knowledge sets to apply to every extension of the PJ-default theory representing a belief state. For postulates 7 and 8 we can actually prove a stronger condition in the case that the antecedent in postulate 8 is satisfied. If the antecedent is not satisfied, there appears to be no obvious way to reformulate postulate 7. ■

Theorem 2. For imperative and stable operations, the contraction operator—satisfies.

1. The principle of categorical matching.
2. $\forall e' \in E((W, D)_A^-)$, there exists some e where $e \in E((W, D))$ s.t. $e' \subseteq e$.
3. If $\forall e : (e \in E((W, D))) \supset (e \not\models A)$, then $E((W, D)_A^-) = E((W, D))$.
4. If $\not\models A$, then $\forall e : (e \in E(W, D)_A^-) \supset (e \not\models A)$.
5. If $\forall e' : (e' \in E((W, D))) \supset (e' \models A)$ then for every $e' \in E((W, D))$, there exists some e where $e \in E(((W, D)_A^-)_A^+)$ s.t. $e' \subseteq e$.
6. If $\models A \leftrightarrow B$ then $E((W, D)_A^-) = E((W, D)_B^-)$.
7. If $\forall e : (e \in E((W, D)_{A \wedge B}^-)) \supset (e \not\models A)$ then $E((W, D)_{A \wedge B}^-) = E((W, D)_A^-)$.

Proof. We prove these results using the THEORIST translations (F, H, C) of (W, D) .

1. Obvious.
2. To prove $\forall e' \in S((F, H, C)_A^-)$, there exists some e where $e \in S((F, H, C))$ s.t. $e' \subseteq e$.
Let $(F, H, C)_A^- = (F', H', C')$. Assume the contrary. Thus, there must exist some $e_x \in S((F', H', C'))$ s.t. there exists no $e \in S((F, H, C))$ where $e_x \subseteq e$. Let $e_x = Cn(F' \cup h')$ where $h' \subseteq H'$. Two cases are possible:
(a) $F = F'$. In this case, $H = H'$. Our assumption holds iff $e_x \cup C$ is unsatisfiable. By stability, this is impossible. Thus, our assumption does not hold.
(b) $F' \subset F$. Thus $H \subset H'$. Let $y = H' - H = F - F'$. If $h' \subseteq H$, then e_x is included in some $e \in S((F, H, C))$, because F' is included in F , h' is included in H , $Cn(F' \cup h')$ is satisfiable (by virtue of being a scenario of (F', H', C')) and $Cn(F' \cup h')$ is consistent with C , as a consequence of consistency with C' and stability. If $h' \not\subseteq H$, then $y'' \subseteq y$, where $y'' = h' - H$. Thus, $Cn(F' \cup h') = Cn(F' \cup y'' \cup h'')$ where $h'' \subseteq H$. Then e_x is included in some $e \in S((F, H, C))$ because $F' \cup y''$ is included in F , h'' is included in H , $Cn(F' \cup y'' \cup h'')$ is satisfiable (by virtue of being a scenario of (F', H', C')), and $Cn(F' \cup y'' \cup h'')$ is consistent with C as a consequence of consistency with C' and stability.
3. To prove that if $\forall e : (e \in S((F, H, C))) \supset (e \not\models A)$, then $S((F, H, C)_A^-) = S((F, H, C))$.

Let $(F, H, C)_A^- = (F', H', C')$. Clearly $F \not\models A$. Thus, $F = F'$, $H = H'$. Then, by stability, because all elements of $S((F, H, C))$ are consistent with C , they will be consistent with C' too.

4. To prove that if $\not\models A$, then $\forall e : (e \in S((F, H, C)_A^-) \supset (e \not\models A))$.
 $C \models \neg A$. Hence proved.
5. To prove that if $\forall e' : (e' \in S((F, H, C))) \supset (e' \models A)$ then for every $e' \in S((F, H, C))$, there exists some e where $e \in S(((F, H, C)_A^-)_A^+)$ s.t. $e' \subseteq e$.
 Let $((F, H, C)_A^-)_A^+ = (F', H', C')$. Let $e_x = Cn(F \cup h)$, where $h \subseteq H$ be some arbitrarily chosen element of $S((F, H, C))$. Because $H \subseteq H'$, $h \subseteq H'$. Two cases are possible:
 - (a) $F \models A$. Then $F' \subseteq F$. Let $y = F - F'$. Clearly $y \subseteq H'$. Thus, $Cn(F \cup h) = Cn(F' \cup y \cup h)$. $Cn(F' \cup y \cup h)$ is satisfiable by virtue of being a scenario of (F, H, C) . Clearly, $Cn(F' \cup y \cup h)$ includes F' . Also, $y \cup h \subseteq H'$. Because e_x is consistent with C , $Cn(F' \cup y \cup h)$ is consistent with C' , by stability. Thus e_x must be included in some $e \in S((F', H', C'))$.
 - (b) $F \not\models A$. Then $F \subset F'$. Clearly, $F' - F = A$. Because $Cn(F \cup h) \models A$, $Cn(F \cup h) = Cn(F \cup A \cup h) = Cn(F' \cup h)$. Clearly $Cn(F \cup A \cup h)$ includes F' . As well, $h \subseteq H'$. $Cn(F \cup A \cup h)$ is satisfiable by virtue of being a scenario of (F, H, C) . Because e_x is consistent with C , $Cn(F \cup A \cup h)$ is consistent with C' , by stability. Thus, e_x must be included in some $e \in S((F', H', C'))$.
6. To prove that if $\models A \leftrightarrow B$ then $S((F, H, C)_A^-) = S((F, H, C)_B^-)$.
 Trivially true.
7. To prove that if $\forall e : (e \in S((F, H, C)_{A \wedge B}^-)) \supset (e \not\models A)$ then $S((F, H, C)_{A \wedge B}^-) = S((F, H, C)_A^-)$.
 Let $(F, H, C)_A^- = (F', H', C')$ and $(F, H, C)_{A \wedge B}^- = (F'', H'', C'')$. By stability, and by the precondition, it is easy to see that every element of $S((F', H', C'))$ is consistent with C'' and every element of $S((F'', H'', C''))$ is consistent with C' . The precondition implies that $F'' \not\models A$. Because $F'' \subseteq F$ and $F'' \not\models A$, $F'' \subseteq F'$. Assume that $F'' \subset F'$. Then there must exist some $x \in F'$ such that $F'' \cup \{x\} \models A \wedge B$. However, this is impossible because $F'' \cup \{x\} \subseteq F'$. Hence $F'' = F'$. Thus $H'' = H'$. ■

5. RELATED WORK

Over a single step, and starting with a deductively closed theory, our framework is identical to the AGM framework in the sense that the outcomes generated by our framework are identical to the choices available to the AGM *selection function*.

Theorem 3. For a uniform PJ-default theory (W, D) , $E((W, D)_A^-) = K \downarrow A$ if $W = K$ where K is a belief set, $D = \{\frac{\top}{\perp}\}$, the contraction operation is imperative and the initial constraint prioritization relation is empty.

Proof. The proof follows directly from the definitions of the removal operation \downarrow , the constraint update operation and imperative belief change operations.

Thus, if we were to start with a deductively closed theory as the set of facts, and an empty set of defaults, then the set of extensions of the default theory obtained after contracting a belief A , would correspond precisely to the set of possible outcomes that the selection function in partial meet contraction. Whereas partial meet contraction requires that a choice is actually made, we do not require any choices, but retain all the multiple possible outcomes compactly represented as a PJ-default theory. ■

The following theorem shows how our approach relates to Nebel's base contraction operator (1989).

Theorem 4. Let (W, D) be a uniform PJ-default theory with $W = B$, where B is a finite belief base, and $D = \{\frac{\perp}{\perp}\}$. Then $Cn(B \simeq A) = Cn((\bigvee E((W, D)^-)) \wedge (B \vee \neg A))$ if the contraction operation is imperative and the initial constraint prioritization relation is empty.

Proof. The proof follows directly from the definitions of the operation \simeq , the constraint update operation and imperative belief change operations.

As with our framework, a belief that becomes suppressed as a result of a contraction operation can be recovered in Nebel's framework when the belief state is revised with the contracted belief. However, our framework permits an explicit operation to undo a contraction, which can also result in beliefs being recovered. Such an operation is not possible in Nebel's framework.

With their commitment to producing a unique outcome for the belief change operation, both the AGM and Nebel frameworks render too many potentially useful beliefs unusable (note that they are not actually discarded, but can be recovered later under certain circumstances); in the AGM framework, this is a consequence of taking the intersection of the selected outcomes, while in Nebel's case, this is a consequence of taking the disjunction of every possible outcome. Our framework retains every possible outcome at all times, and thus does not suffer from this problem.

Brewka (1989) shows how belief revision can be viewed as a simple process of adding new information to theories represented in his *preferred subtheories framework*, which is a generalization of the THEORIST framework (Poole et al. 1987). Brewka's framework (1989) of preferred subtheories differs from THEORIST in two significant ways. First, *facts* are done away with, making every formula in the knowledge base potentially refutable. Second, one is allowed to define a partial order on the formulas in the knowledge base. A preferred subtheory, Brewka's analogue of a THEORIST maximal scenario, is a consistent subset of the knowledge base constructed by starting with formulas with the highest priority (as defined by the partial order) and progressively adding as many formulas of lower priority levels as can be consistently added. As with THEORIST, a knowledge base can have several preferred subtheories. Brewka (1989) shows that a knowledge base of this kind can be revised by simply adding the new formula and augmenting the partial order to incorporate any ordering relationships that might exist between this formula and the existing elements of the knowledge base. Also, if this framework is augmented to include THEORIST-style constraints, and a partial order is defined on the set containing both the formulas representing hypotheses and formulas representing constraints, then contraction is shown to be a simple case of adding a constraint to the knowledge base and augmenting the partial order. The improvements achieved by Brewka's belief change framework over earlier ones are twofold. First, the belief change operator is simple and totally incremental. Second, earlier information is not thrown away, but is retained in an elegant fashion. Nebel (1991) establishes a restricted form of equivalence between nonmonotonic inference and belief change along similar lines.

In the case that the new belief (either a new hypothesis, as in revision, or a new constraint, as in contraction) always has a higher priority, under the partial ordering, than all existing beliefs, Brewka's framework turns out to be similar to ours. As the following example shows, his framework avoids the problem of spurious beliefs in most cases. ■

Example 9. Let the initial knowledge base consist of the set $\{a, a \rightarrow b\}$ of hypotheses with no ordering relationship being defined for the hypotheses. To contract b from this knowledge base, we add the constraint $\neg b$, written as $< \neg b >$ to the knowledge base, together with the

ordering relations $\langle \neg b \rangle \geq a$ and $\langle \neg b \rangle \geq a \rightarrow b$. We get two maximal scenarios, one containing a and the other containing $a \rightarrow b$. Further revision of the knowledge base with a results in the addition of this hypothesis at a higher priority level than all existing elements (hypotheses or constraints) of the knowledge base. There is only one maximal scenario at this point, consisting of a and its logical consequences.

The similarity of Brewka's framework to ours is not surprising, given that we use, like Brewka, nonmonotonic theories that can generate possibly many different consistent sets of beliefs, to represent a belief state. Like Brewka, our approach is incremental, and information is never thrown away. Our choice of nonmonotonic formalism is very similar too, given the results in Delgrande and Jackson (1991) relating PJ-default logic to THEORIST with constraints. However, because Brewka does not explicitly account for the maintenance of belief constraints, his formalization may provide unintuitive results as the following example shows.

Example 10. Consider an initial knowledge base containing only one hypothesis and no constraints $\{a\}$. Let us now contract $a \vee b$ from this knowledge base. This entails the addition of the constraint $\langle \neg(a \vee b) \rangle$ to the knowledge base, and augmenting the partial order such that the new constraint has higher priority than all existing elements of the knowledge base. If one were to revise the knowledge base with b , there would be one maximal scenario containing both a and b . Note, however, that the new evidence obtained since retracting $a \vee b$ from the knowledge base does not warrant renewed belief in a . The problem arises because the presence of b at a higher priority level disables the constraint $\langle \neg(a \vee b) \rangle$.

We improve upon Brewka's work by explicitly accounting for the maintenance of belief constraints. Necessary disbelief constraints are treated as a set of formulas to be explicitly disbelieved. We update this set at every belief change step, by retaining as many constraints, or parts of constraints, as are compatible with more recent constraints. Thus, in the previous example, we would update this theory to account for revision with b by removing b from the set of necessary disbelief constraints, but leaving a intact. The problems with Brewka's framework stems from the fact that it uses syntactic units (the constraints) that are enabled or disabled as whole units and not in terms of the individual components. In fact, his system would behave like ours only if the only constraints permitted are atomic constraints.

Whereas Brewka's system uses a recency-based heuristic to order belief constraints, our framework is more general by permitting arbitrary constraint prioritizations. We differ further from Brewka in that we factor out the use of priorities on beliefs entirely from the belief change process. Whereas Brewka's framework would only generate those maximal scenarios, that respect the existing orderings on the beliefs, our framework would generate all maximal scenarios that satisfy the relevant belief constraints. Our framework would coincide with Brewka's, in this respect, if the constraint prioritization was based on recency.

Nayak et al. (1994), Boutilier (1992) and Williams (1993) address the question of generating a new selection function as result of a belief change step. However, they all use a recency-based heuristic for ranking revisions. More importantly, they do not address the problem of nonpersistence of the effects of contractions.

Sattar and Goebel (1991) have extended the THEORIST system to a consistency-based reason maintenance system that incrementally identifies inconsistent sets of hypothesis instances (nogoods) and stores them along with partial explanations for use in subsequent computations. The computational advantages of this approach can be seen in Goodwin and Sattar (1993). Our framework could be easily implemented using this extended version of THEORIST. Also, it would be interesting to see how the idea of *experiments*, as used in

THEORIST-style systems (Sattar and Goebel 1991) could be used to discriminate between the competing outcomes of a belief change step.

6. CONCLUSION

The primary contribution of this study is the identification of a set of requirements for belief change operators from the viewpoint of their iterated application. A threefold categorization of belief in a belief state is introduced: *necessary beliefs*, *necessary disbeliefs*, and *tentative beliefs*. A precise characterization of belief migration across these classes is provided. The problem of spurious beliefs is identified, and is used to motivate the need both for maintaining multiple theories in the same belief state, as well as for maintaining belief constraints. One compact representation is obtained for the multiple possible outcomes of a belief change operation. By using PJ-default logic to represent a belief state, we are able to obviate the need for recording belief constraints separately; the constraints are clearly identifiable components of a PJ-default theory. By retaining all outcomes at any given point during the process of belief change, we permit the use of different prioritizations of the beliefs at different times to actually select one theory out of the many that may potentially constitute a belief state. Our framework can be viewed as a formalization of the process through which a default theory evolves into a more accurate representation of the world. Belief change drives this process, and in abstract terms, this involves demoting facts known to be true to the status of defaults. By showing how a nonmonotonic formalism such as default logic can be crucial to knowledge revision, this study further explicates the close relationship between these related, yet separate, areas of inquiry.

REFERENCES

- ALCHOURRÓN, C. E., P. GÄRDENFORS, and D. MAKINSON. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, **50**:510–530.
- BOUTILIER, C. 1992. Sequences of revisions: On the semantics of nested conditionals. Technical Report 92-23, Department of Computer Science, University of British Columbia.
- BREWKA, G. 1989. Belief revision in a framework for default reasoning. *In Proceedings of the Konstanz Workshop on Belief Revision*, pp. 206–222.
- DELGRANDE, J. P., and W. K. JACKSON. 1991. Default logic revisited. *In Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning*, pp. 118–127.
- GÄRDENFORS, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, MA.
- GÄRDENFORS, P., and D. MAKINSON. 1988. Revisions of knowledge systems using epistemic entrenchment. *In Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pp. 83–95.
- GÄRDENFORS, P., and H. ROTT. 1992. Belief revision. Technical Report 30, Zentrum Philosophie und Wissenschaftstheorie, Universität Konstanz.
- GHOSE, A. K., and R. G. GOEBEL. 1992. Default reasoning as belief change: A rationale. *In Proceedings of the Second Pacific Rim International Conference on AI*.
- GHOSE, A. K., P. O. HADJINIAN, A. SATTAR, J.-H. YOU, and R. G. GOEBEL. 1993. Iterated belief change: A preliminary report. *In Proceedings of the Sixth Australian Joint Conference on AI*.
- GOODWIN, S., and A. SATTAR. 1993. On computing preferred explanations. *In Proceedings of Australian Joint Conference on Artificial Intelligence*, Melbourne, Victoria, November, World Scientific Publishing Co., pp. 45–52.

- HANSSON, S. O. 1991. *Belief Base Dynamics*. Uppsala. Kluwer Academic.
- HARPER, W. L. 1975. Rational belief change, popper functions and counterfactuals. *Synthese*, **30**:221–262.
- KATSUNO, H., and A. O. MENDELZON. 1991. On the difference between updating a knowledge base and revising it. *In* Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning.
- NAYAK, A. C. 1993. Studies in Belief Change. Ph.D. Thesis, Department of Philosophy, University of Rochester.
- NAYAK, A. C., N. Y. FOO, M. PAGNUCCO, and A. SATTAR. 1994. Entrenchment kinematics 101. *In* Proceedings of the Seven Australian Joint Conference on Artificial Intelligence, pp. 301–311.
- NEBEL, B. 1989. A knowledge level analysis of belief revision. *In* Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning, pp. 301–311.
- NEBEL, B. 1991. Belief revision and default reasoning: Syntax-based approaches. *In* Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning, pp. 417–428.
- POOLE, D., R. GOEBEL, and R. ALELIUNAS. 1987. Theorist: A logical reasoning system for defaults and diagnosis. *In* The Knowledge Frontier: Essays in the Representation of Knowledge. *Edited by* N. J. Cercone and G. McCalla. Springer Verlag, Berlin, pp. 331–352.
- REITER, R. 1980. A logic for default reasoning. *Artificial Intelligence*, **13**(1 and 2):81–132.
- SATTAR, A. 1991. Efficiency Aspects of Hypothetical Reasoning. Ph.D. Thesis, University of Alberta.
- SATTAR, A., and R. GOEBEL. 1991. Consistency-motivated Reason Maintenance in Hypothetical Reasoning. Technical Report TR-91-12, Department of Computing Science, Edmonton, Alberta.
- SATTAR, A., and R. GOEBEL. 1991. Using crucial literals to select better theories. *Computational Intelligence*, **7**(1):11–22.
- SPOHN, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. *In* Causation in Decision, Belief Change and Statistics, vol. 2. *Edited by* W. L. Harper and B. Skyrms, pp. 105–134.
- WILLIAMS, M. A. 1993. Transmutations for theory bases. *In* Proceedings of the sixth Australian Joint Conference on AI, pp. 83–89.