Definitional Constraints

Norman Y. Foo[1] and **Abhaya Nayak**[1] and **Maurice Pagnucco**[1]

Abstract. Certain constraints that should be interpreted as definitions or quasi-definitions can cause prediction problems for logics of action. Standard ways to avoid them rely on notions of priority in the minimization of fluent change. We offer an alternative approach based on belief revision. Revision is performed on the reduct of the system theory after isolating definitions, and the revised theory is then expanded by restoring the definitions. The basis for this is a weak monotonicity result for AGM contraction. Examples from the literature are used to illustrate the approach.

1 Introduction

A substantial part of knowledge representation research is devoted to formalisms for reasoning about, and procedures for computing, the effects of actions. While much has been achieved, the notorious frame problem is still resistant to a complete solution. In particular, every logic proposed for reasoning about state change seems to provoke a toy example in which the logic fails to make the "commonsense" prediction. Sandewall's response [13] is to provide a new methodology altogether — first classify dynamic systems using ontological characteristics like inertia, functional dependencies, determinism, etc., and then show that some action logics are adequate for them. Indeed, Lin and Shoham's recent work [11] may be seen in this light. The basic thrust of the ontological approach is to acknowledge at the outset that certain kinds of information or knowledge have *semantic* content that may be extra-logical. One such acknowledgement is that system constraints come in a variety of semantic guises, among which are (i) truly systemic ones like "a block can only be in one place at a time", (ii) causal ones like "if something lives it must absorb energy", and (iii) definitional ones like "bachelors are human males who are not married". (These are semantic distinctions that cannot be deduced from the syntax of constraints, which may simply be clauses). It should not be surprising that logical conflation of such diverse semantics could pose difficulties for state updates.

In this paper we show how some problems identified with constraints of types (ii) and (iii) above can be avoided using particular kinds of belief revision. Systemic constraints of type (i), when present, need not receive special treatment.

These problems are well-known and so are their proposed solutions in terms of traditional logics like circumscription [9]. We propose here an alternative approach based on belief revision [5]. One reason for doing so is that it exposes clearly the ancillary nature of definitions while reducing the state change problem to belief change. This is in conformity with Sandewall's ontological paradigm in that we acknowledge certain constraints are definitional in character. Other reasons are (a) belief revision (update) as captured by the AGM [5] (respectively the KM [8]) postulates are now widely accepted as rational ways to change theories, and (b) computational versions of AGM are available [15, 3] and easy to use. Definitional constraints are of interest because they appear frequently in knowledge representation as a result of abstraction while causal constraints often appear as quasi-definitions (see below). While the effect of our approach is not different from priority-based schemes, the identification of definitions and quasi-definitions automatically selects the appropriate underlying theory for revision (or update).

The paper is structured as follows. Section 2 reviews a number of popular state update proposals. Section 3 identifies, in general form, the difficulties that definitional and quasi-definitional constraints cause and outlines a common approach to overcome them. This section also cites well-known examples from the literature that illustrate these difficulties. Section 4 presents our alternative approach based on belief revision and justifies it on the basis of the commutativity of belief contraction and quasi-definitional and definitional expansion. Section 5 re-examines the problem examples in section 3 from the viewpoint of section 4 and shows that those difficulties are thereby obviated. Section 6 discuses limitations and future work.

2 Some Approaches to State Updates

All state update procedures employ some notion of minimization. This reflects the inert nature of the dynamic system wherein "nothing is changed if it does not have to". The words within quotes are subject to various interpretations, and each such interpretation gives rise to a different notion of minimal change. For instance, circumscription [9] and belief revision using the AGM postulates [5] are two such notions. Depending on what is being circumscribed we get different measures of minimality. Belief revision however uses postulates on changes to belief sets that encode the agent's knowledge of the present state, and depending on an extra-logical ordering called *entrenchment* we get different dynamics for beliefs. In this paper we discuss the proposals in the possible worlds approach (PWA) of Ginsberg and Smith [6], the possible models approach (PMA) of Winslett [16], Dalal's (DAL) [2] update, and the persistence set approach (PSA) of Zhang and Foo [17].

Let us start with reviews of the PWA, PMA, DAL, and PSA. All these approaches deal with the question, "Given that an agent believes itself to be in a particular state S of the world, and some constraints C that the agent believes no world will violate, what reasonable inferences can the agent

[1] Knowledge Systems Group, Basser Department of Computer Science, Madsen Building, F09, University of Sydney, NSW 2006, Australia.
Email: {norman | abhaya | morri}@cs.su.oz.au.

ECAI 96. *12th European Conference on Artificial Intelligence*
Edited by W. Wahlster
Published in 1996 by John Wiley & Sons, Ltd.

make about the new state of the world when an action A takes place?"

The constraints C are first-order formulas, and states are captured as Herbrand models of C. A state description is a formula or a set of literals (*fluents*[13]) that are true in that state. It can be partial or indefinite in the sense that it can be satisfied by several states; a complete description is a ground-complete set of literals. The "distance" between complete states S_1 and S_2, $\text{diff}(S_1, S_2)$ is simply the symmetric difference of the literal sets S_1 and S_2.

It is traditional in the logic of actions to specify by a formula $\text{post}(A)$ (the *post-condition* of the action A) the immediate consequences of the action A. If the action A has side-effects (called the ramifications), they are obtained as logical consequences of $\text{post}(A)$ and the constraints. All candidates for the next state have to obey the constraints as well as satisfy the formula $\text{post}(A)$. All approaches select as next states a subset that is in some sense "minimally different" from the set of initial states. Each has its own notion of minimality. Here we will merely outline these notions. In the PWA, minimality refers to the difference in the *theories* that describe the (sets of) states. If two theories are admissible as descriptions of $\text{post}(A)$ and the constraints, the one with a smaller change from the initial theory is selected. The notion of minimal theory change is required to satisfy rationality conditions like the AGM or KM postulates. (Familiarity with these postulates is not required until section 4.)

In the PMA, minimality is figured as follows. Let S be an initial state which satisfies the initial state description (it can be partial, so there may be more than one such S). Suppose T is a possible next state (so it satisfies $\text{post}(A)$ and system constraints), then T is selected iff there is no other possible next state T_1 such that $\text{diff}(S, T_1) \subseteq \text{diff}(S, T)$; the collection of all such T's for every initial state S is the set of PMA next states. In DAL, the measure $\text{diff}(S, T)$ is over all possible initial states and candidate next states, and it is the cardinality $\mid \text{diff}(S, T) \mid$ (rather than set containment as in the PMA) that counts. Only those T's that give minimal $\mid \text{diff}(S, T) \mid$ are selected. In the PSA, $\text{diff}(\_, \_)$ does not play a direct role. Instead, if S is an initial state, $\text{post}(A)$ is first used to eliminate from S literals that conflict with it. Any eliminated literal that can be interpreted as an essential logical support for some hitherto surviving literal l in S will then cause l to become insecure, hence deleted. The successive propagation of this insecurity and deletion, when it stabilizes, will yield a residual set U of literals, and any state consistent with U (and satisfying $\text{post}(A)$ and C) is a possible next state.

Relative to the AGM [5] and KM [8] postulates for belief change and updates respectively, it is known that PMA satisfies KM, DAL satisfies AGM, and PSA satisfies neither. KM differs from AGM only when initial state descriptions are incomplete; indeed, it is known that the KM operators can be derived from the AGM operators if intersection and maximal completions are admitted. *So, for complete state descriptions it suffices to use belief revision for the purpose of state update.* PWA satisfies AGM (hence KM) when initial states are unique, but violates KM otherwise.

3 Definitional Formulas and Problems

Recently Foo and Zhang [4] demonstrated that difficulties with state change predictions by all of these approaches are intrinsic to the logical structure of the state description and constraint specifications when *definitional* formulas are not isolated from the logical description of the system. This is true even when the initial state is unique, i.e., when the state description is complete, so it is not missing state information that is the problem. Indeed, in this paper we shall assume that initial states are complete (hence KM update is AGM revision).

Definition 1 *Let T be a collection of formulas. A clause CL is T-literal-minimal if (i) it is not a tautology, (ii) $T \vdash CL$ and (iii) for no proper subclause CLS of CL is it the case that $T \vdash CLS$.*

T is usually the logical closure of the constraints and a state description. Literal-minimality corresponds to the property of being a prime-implicate in Boolean simplification, and is intended to exclude clauses that contain trivial (e.g. tautologies or constraints) subclauses.

Definition 2 *If CL is a constraint clause, CL is literal-minimal if it is not a tautology and no proper subclause of it is also a constraint.*

The next two assumptions ensure that action and constraint specifications are non-trivial.

Assumptions

(i) Every constraint is literal-minimal; (ii) For every action A, its post-condition $\text{post}(A)$ is consistent with every constraint clause.

We consider below formulas of the form $d \leftrightarrow l_1 \vee \ldots \vee l_n$. These can be decomposed into the clause $d \rightarrow l_1 \vee \ldots \vee l_n$ and the set of clauses $\{d \leftarrow l_1, \ldots, d \leftarrow l_n\}$. By the literal-minimality of $d \leftrightarrow l_1 \vee \ldots \vee l_n$, we mean the literal-minimality of each of these decomposed clauses. The intended meaning of $d \leftrightarrow l_1 \vee \ldots \vee l_n$ is that it is a defining formula for d in terms of the literals l_i on the right.

Proposition 1 (from [4]) and its generalization proposition 2 below are used in specific forms to arrive at the anomalies discussed in the Brewka and Hertzberg example below, and also in the four examples that follow it. The generality of these propositions indicate that the problems are pervasive.

Proposition 1 *[4] Suppose a ground constraint clause CL is of the form $d \leftrightarrow l_1 \vee \ldots \vee l_n$ and is literal-minimal, S_0 (the initial state) is $\{d, l_1, \ldots l_k, \neg l_{k+1}, \ldots, \neg l_n\}$, and $post(A)$ is $\neg l_1 \wedge \ldots \wedge \neg l_k$. Then $\neg d \vee l_{k+1} \vee \ldots \vee l_n$ is $post(A) \cup CL$-minimal.*

Corollary 1 *(i) None of the literals $d, l_{k+1}, \ldots, l_n$ is in the persistence set of A starting from S_0, hence by the PSA all of these literals may change in the new state.*

(ii) Either d remains true in the new state and (at least) one of $\neg l_{k+1}, \ldots, \neg l_n$ changes to its positive, or $\neg d$ holds in the new state with $\neg l_{k+1}, \ldots, \neg l_n$ remaining unchanged. Then by the PMA and DAL, d is uncertain in the new state.

The *duals* of the proposition and its corollary hold. In other words, conjunctions on the right hand side of CL cause similar problems, so it suffices to consider solutions for the disjunctive form above. Moreover, the same line of reasoning applies to a more general form of the constraint CON below, where it is in conjunctive normal form, or equivalently, is a set of clauses.

Definition 3 *Let T be a collection of formulas. A conjunction CON of clauses CL_i, $1 \leq i \leq n$, is T-clause-minimal if (i) it is not a tautology, (ii) $T \vdash CL_i$ for each i, $1 \leq i \leq n$, and (iii) CL_i is T-literal-minimal for each i, $1 \leq i \leq n$.*

If T is just the logical closure of CON, we just say clause-minimal instead of T-clause-minimal.

Proposition 2 *Suppose CON is of the form $d \leftrightarrow \bigwedge_{1 \leq i \leq n} C_i$ as above and is clause-minimal. Let C_1 be $l_1 \vee \ldots \vee l_n$. Assume also that the initial state S_0 is $\{d, l_1, \ldots l_k, \neg l_{k+1}, \ldots, \neg l_n\} \cup L$ where L is a set of literals that satisfies $\bigwedge_{2 \leq i \leq n} C_i$; and $post(A)$ is $\neg l_1 \wedge \ldots \wedge \neg l_k$ and is consistent with $\bigwedge_{2 \leq i \leq n} C_i$. Then $\neg d \vee l_{k+1} \vee \ldots \vee l_n$ is $post(A) \cup CON$-minimal.*

Corollary 2 **(i)** *None of the literals $d, l_{k+1}, \ldots, l_n$ is in the persistence set of A starting from S_0, hence by the PSA all of these literals may change in the new state.*
(i) *Either d remains true in the new state and (at least) one of $\neg l_{k+1}, \ldots, \neg l_n$ changes to its positive, or $\neg d$ holds in the new state with $\neg l_{k+1}, \ldots, \neg l_n$ remaining unchanged. Then by the PMA and DAL, d is uncertain in the new state.*

The conclusions in the corollary are counter-intuitive in many interpretations of the literals, as illustrated by a slight modification of the door-window example of Brewka and Hertzberg [1] where we have added an abstraction "bright" to their example, giving the effect of a definitional constraint. Let the constraint C be bright $\leftrightarrow$ door_open $\vee$ window_open. Suppose the initial state S_0 is {bright, door_open, ¬window_open}, and action A is to shut the door so post(A) is ¬door_open. Then $C \cup$post(A) $\vdash$ bright$\vee\neg$window_open, but $C \cup$post(A) $\not\vdash$ bright and $C \cup$ post(A) $\not\vdash$ ¬window_open, so by the PSA both bright and window_open may not persist, an unintuitive result. Similarly, since the states {¬bright, ¬door_open, ¬window_open} and {bright, ¬door_open, window_open} are both minimally different from the initial state {bright, door_open, ¬window_open} because they both have the same number of literals changed, DAL and PMA admit both as resultant states, and one of these is unintuitive.

These difficulties are common, as evidenced by the following two examples from the literature. In each case, the state updates given by the cited approaches cause problems similar to the ones above, and can be verified by appeal to either corollary 1 or 2.

(a) The stuffy room and ducts example of Ginsberg and Smith [6] contains essentially a constraint much like C above. In simplified form we can formulate it as stuffy $\leftrightarrow$ $\exists x \exists y(\mathrm{on}(x, d_1) \wedge \mathrm{on}(y, d_2) \wedge x \neq y)$. Using a closed world assumption, this gives rise to constraints that are dual to the form C. The simplest case that illustrates the problem is an initial state $\{\neg\mathrm{stuffy}, \mathrm{on}(\mathrm{tv}, d_1), \neg\mathrm{on}(\mathrm{speaker}, d_2)\}$ with PMA and DAL admitting two future states {stuffy, $\mathrm{on}(\mathrm{tv}, d_1)$, $\mathrm{on}(\mathrm{speaker}, d_2)$}
and
{¬stuffy, $\neg\mathrm{on}(\mathrm{tv}, d_1)$, on(speaker, d_2)}
as a result of moving the speaker to duct 2. PSA makes stuffy $\wedge$ $\mathrm{on}(\mathrm{tv}, d_1)$ uncertain.

(b) In the switch example of Lifschitz [10], the constraint is of the form on_lamp $\leftrightarrow$ ((sw1 $\rightarrow$ sw2) $\wedge$ (sw2 $\rightarrow$ sw1)), again like C above. Here the simplest case is an initial state {¬on_lamp, sw1, ¬sw2} with PMA and DAL admitting two next states {on_lamp, ¬sw1, ¬sw2} and {¬on_lamp, ¬sw1, sw2} as a result of switching switch 1 to its negative position. PSA makes on_lamp and sw2 uncertain.

Another type of constraint is only *quasi-definitional* in that C takes the form $d \leftarrow l_1 \vee \ldots \vee l_n$. This is illustrated by the examples below.

(a) Myers and Smith's example [12] has, among others, constraints like on_table(x) $\rightarrow$ liftable(x) and on_table(x) $\rightarrow$ safe_to_baby(x). Here the problem is illustrated by the impossibility of distinguishing the proper update based on syntax alone. For instance, assuming there is only one object to reduce atoms to propositional form, starting from the initial state {on_table, safe_to_baby}, and an action of moving it from the table to the floor, the PMA yields ¬on_table, safe_to_baby as the next state while the PSA yields the more intuitive {¬on_table, ¬safe_to_baby} as the next state. However, in considering liftable as an intrinsic property, the intuitive and counter-intuitive roles are reversed for the PMA and PSA.

(b) Introducing into the Yale Shooting Problem constraints such as alive $\rightarrow$ walking or walking $\rightarrow$ alive gives unexpected next states. Using the first constraint starting with the state {alive, walking}, the PWA, PMA and DAL admit {¬alive, walking} as a possible next state for action A with post(A) being ¬alive. The PSA would give the expected result here but it fails for the second constraint. What is considered to be quasi-definitional depends on causality.

There are a number of existing remedies for these problems, and they most commonly hinge on prioritizing some predicates in the minimization process or on entrenching them shallower in belief revision. An exemplar of the former is a very recent approach due to Kartha and Lifschitz [7] which employs the notion of *frame literals*. It is assumed that in state change the frame literals should be minimized before minimizing change in other literals. So-called *relaxation* axioms are introduced to exempt frame literals from this prioritized minimization in the context of, say, nondeterminism. For definitions of the form $d \leftrightarrow l_1 \vee \ldots \vee l_n$ where d is defined in terms of the literals on the right, these literals are part of the frame, so they should not change as easily as the d that is defined. In belief revision, the same effect is achieved by an entrenchment that places $l_1 \vee \ldots \vee l_n$ at least as deep as $d \vee l_1 \vee \ldots \vee l_n$ (which implies, in the AGM framework, that the entrenchments are equal). Whatever the remedy, some *extra-logical* decision has to made about the priorities, or entrenchments of the predicates or literals (and even formulas).

4 Revision of Reducts

Familiarity with the AGM (and KM) postulates is assumed in this section. We examine here the strategy for state updates based on belief revision. The central idea is to first ignore definitions or quasi-definitions in knowledge sets when belief revision is performed, and these should be restored only as ramifications of the underlying revision. The revision is made on the *reduct* of the state description and theory by suppressing definitional formulas. The basic results are two propositions that characterize a weak monotonicity property of AGM contraction.

The first-order language considered will have equality, a finite number of constant and predicate symbols, but no function symbols. Ground literals are the main concern unless otherwise stated. By $\mathcal{L}(T)$ we mean the language of the set of formulas T. We will use upper case letters to stand for literals or formulas. When a sequence of literals are considered, lower cases with subscripts are also used. Lower case Greek is used for conjunctions of literals (e.g. accumulated proof hypotheses). The symbol d is reserved for a defined or quasi-defined atom in the sense explained below. Belief or knowledge sets are logically closed, i.e., theories (but see later remarks on finite bases).

Let H be a belief set, def be a formula of the form $d \rightarrow \phi$ or $d \leftrightarrow \phi$ where ϕ is in $\mathcal{L}(H)$, and d is not in $\mathcal{L}(H)$. (Typically, ϕ has the shape of the right hand sides of CL in proposition 1 or CON in proposition 2.) Let $K = Cn(H \cup \{\text{def}\})$ where Cn denotes logical consequence. As AGM logic is insensitive to syntax, we note that the presence of def permits us to confine attention to the AGM operators using only formulas in $\mathcal{L}(H)$ as every occurrence of d can be replaced by the rhs of def. With these assumptions the following is well-known.

Proposition 3 *[14] K is a conservative extension of H, i.e., for every $B \in \mathcal{L}(H)$, if $B \in K$ then $B \in H$ (the converse is trivial).*

The importance of this result is that if we start with K, as long as we know d the theory H can be obtained simply by suppressing all formulas that have d as a component. H is the *reduct* of K relative to d-formulas and will retain all the systemic constraints. (Practically, however, we specify H and K indirectly with finite *bases*.) The discussion below is confined to a single definition def. The extension to multiple definitions is easy and omitted. Now, in the AGM framework, entrenchments determine belief dynamics. If K is our knowledge set with a definition def added as above to H, and H has an underlying entrenchment $\leq_H$, it would be expected that the corresponding entrenchment $\leq_K$ in K would agree with that in H for formulas in $\mathcal{L}(H)$. A way to see this is that the addition of *abbreviations* or *abstractions* (which are interpretations of the role of d in the definition def) should not alter the relative importance of the formulas in the original knowledge base H. Further, once def is added, it should have the character of an axiom or system constraint. Formally, we have:

Assumptions
(i) $\leq_H \subseteq \leq_K$
(ii) def $\in K_B^-$ for all $B \in \mathcal{L}(H)$.

The notation K_B^- means the contraction of B from K. (Some formulas that imply B are also removed as a consequence.) For the AGM postulates on contraction and entrenchment, see [5].

Assumption (ii) is entailed by a K-entrenchment for def that is strictly greater than any formula in H. That would almost make it a H-tautology except for the fact that def $\notin \mathcal{L}(H)$. Assumption (i) strictly strengthens proposition 3 — it is possible to have K conservatively extend H and yet proposition 4 below can fail to hold if assumption (i) is not satisfied. In this sense, assumption (i) can be regarded as a strong conservative extension postulate.

Lemma 1 *If $B, C \in \mathcal{L}(H)$, then $B \in H_C^-$ iff $B \in K_C^-$.*

This lemma is a direct consequence of assumption (i) and the (C$-$) condition [5] for AGM contraction, viz., $B \in H_C^-$ iff (a) $B \in H$ and (b) either $\vdash C$ or $C <_H B \vee C$.

The first and second basic results below consider contractions relative to a formula in H. Later we will indicate how revisions can be handled, and how both contractions and revisions relative to formulas in K but not in H can be accommodated.

Proposition 4 *With the notation and assumptions as above, and any $A \in \mathcal{L}(H)$, the diagram in figure 1 commutes. (Upward arrows correspond to AGM expansion.)*

$$\begin{tikzcd} K \arrow[r, "-A"] & K_A^- \\ H \arrow[u, "+\text{def}"] \arrow[r, "-A"] & H_A^- \arrow[u, "+\text{def}"'] \end{tikzcd}$$

Figure 1. Commutativity of definitional expansion and contraction.

Proof: There are two cases. In the first case, assume $B \in \mathcal{L}(H)$.

(a) Suppose $B \in Cn(H_A^- \cup \{\text{def}\})$, i.e., we chase the arrows along the bottom to upward right. Then there is some γ in H_A^- such that def $\vdash \gamma \rightarrow B$. By lemma 1, $\gamma \in K_A^-$, and by assumption (ii) def $\in K_A^-$, so $B \in K_A^-$.
(b) On the other hand, chasing the arrows up left and along the top right, suppose $B \in K_A^-$. By lemma 1, $B \in H_A^-$, hence $B \in Cn(H_A^- \cup \{\text{def}\})$.

In the second case, assume $B \notin \mathcal{L}(H)$.

(a) Chasing the arrows from bottom then upward right, assuming $B \in Cn(H_A^- \cup \{\text{def}\})$, the proof is identical to (a) above in showing $B \in K_A^-$.
(b) On the other hand, chasing the arrows up left and top right, suppose $B \in K_A^-$. Let B_1 be the formula that results from B by substituting every occurrence of d by the right hand side of the definition def (which is $d \leftrightarrow \phi$). By AGM insensitivity to syntax, we have $B_1 \in K_A^-$, and $B_1 \in \mathcal{L}(H)$, so the proof now follows that of (b) above to give $B_1 \in Cn(H_A^- \cup \{\text{def}\})$. By AGM syntax insensitivity again, $B \in Cn(H_A^- \cup \{\text{def}\})$.

The second basic result is about quasi-definitions of the form $d \leftarrow \phi$ where ϕ is in $\mathcal{L}(H)$ and $d \notin \mathcal{L}(H)$. To highlight its quasi-definitional nature, we write this formula as qdef, and let $K = Cn(H \cup \{\text{qdef}\})$. The assumptions above and Lemma 1 still hold, but with qdef replacing def. If the proof of proposition 4 is examined, it will be seen that only at part (b) of the case $B \notin \mathcal{L}(H)$ is the particular form of def used. Hence the commutativity of figure 1 holds by interpreting the contraction in K as before except for this case. To recover figure 1 for this remaining case we need to re-interpret the contraction in K as a *foundational* one.

Definition 4 *Let Δ be a set of ground literals in $\mathcal{L}(H)$, and B a literal in K. Δ is a qdef-support for B if (i) $\Delta \cup \{qdef\} \vdash B$, (ii) $\Delta \not\vdash B$, (iii) qdef $\not\vdash B$, and (iv) Δ is set-inclusion minimal with respect to the preceding conditions.*

The particular form of qdef, viz., $d \leftarrow \phi$, where ϕ is in $\mathcal{L}(H)$ and $d \notin \mathcal{L}(H)$, yields the following unsurprising corollary:

Corollary 3 *(i) If $K = Cn(H \cup \{qdef\})$ and qdef has the form above, then the only qdef-support for d is in H, and $\neg d$ has no qdef-support.*
(ii) For $B \in K$ and $B \notin \mathcal{L}(H)$, where B is not d alone, there are $C_1, \ldots, C_n$ in H such that $C_1, \ldots, C_n,\ qdef \vdash B$.

Part (i) is easily established by a model-theoretic argument and part (ii) by induction on the complexity of B with (i) as basis.

Definition 5 *If B is in K but not in $\mathcal{L}(H)$ a contraction K_A^- is qdef-foundational if whenever A is in every qdef-support of B, $B \notin K_A^-$.*

By this definition, every such B is characterized by corollary 3. If B happens to be in $\mathcal{L}(H)$ the contraction reduces to the one for H by virtue of lemma 1. There is nothing to prevent some contractions in H itself to be α-foundational for some literal B. This may be achieved by constructing a δ that expresses the disjunction of all the α-supports for B in H, and specifying equal entrenchments for δ and $B \vee \delta$. This is a consequence of the (C$-$) condition alluded to earlier. The qdef-foundational nature of the contraction in K can be similarly specified by entrenchments that will guarantee the behavior required in the definitions above.

Proposition 5 *In figure 1, if def is replaced by qdef and K_A^- is qdef-foundational, then the diagram commutes.*

Proof: We only need to replace part (b) of the proof for the case $B \notin \mathcal{L}(H)$. So, chasing the arrows up left and top right, suppose $B \in K_A^-$. Because K_A^- is qdef-foundational, by corollary 3 B has a qdef-support in K_A^-, and therefore there are $C_1, \ldots, C_n$ in $K_A^-, C_1, \ldots, C_n$ all in $\mathcal{L}(H)$, such that $C_1, \ldots, C_n, \text{qdef} \vdash B$. By lemma 1, $C_1, \ldots, C_n$ are all also in H_A^-, hence $B \in \text{Cn}(H_A^- \cup \text{qdef})$.

We now address the cases that were postponed in the two basic propositions. First, consider the case in figure 1 where A is not in $\mathcal{L}(H)$. When we have def as the definitional formula, the situation reduces to that already dealt with because we can replace every occurrence of d in A by ϕ where def is $d \leftrightarrow \phi$ and ϕ is in $\mathcal{L}(H)$. In the qdef case where the quasi-definition is $d \leftarrow \phi$ with ϕ in $\mathcal{L}(H)$, contracting with respect to d is, by the AGM entrenchment postulates on provability, equivalent to contracting with respect to ϕ. Contracting with respect to more complex formulas that contain d is a minor extension of this idea.

Nondeterminism can arise from post(A) of the form $\neg d$ where def or qdef has a right hand side of the form $\phi \vee \psi$, where $\phi \vee \psi$ is in $\mathcal{L}(H)$. The entrenchment within H may bias the result in one way or another. In the (modified) Brewka and Hertzberg example, this can arise when the initial state has both the window and door shut, and post(A) is bright. Our proposal remains agnostic about how this is achieved by just revising H to be indefinitely described by door_open $\vee$ window_open, and then inferring bright by expansion. Nondeterminism thus comes naturally with belief revision and does not require any special device like relaxation of frames.

Next, when post(A) is inconsistent with K, *revision* rather than contraction is necessary. Revision, denoted by $*$, is another AGM operator that can be used to encode post(A) when the latter conflicts with K. The revision K_A^* of K by A is expressible as the Levi identity (see [5]) $K_A^* = (K_{\neg A}^-)_A^+$, and since AGM expansion $+$ is monotone, figure 1 will now hold for both def and qdef (with the foundational proviso in the latter case) with contraction $-$ replaced by revision $*$. In other words, propositions 4 and 5 also hold for revision. These cases complete the picture.

5 Examples

To illustrate the utility of the above approach we re-examine some of the examples. This re-examination uses propositions 4 and 5 for the def and qdef cases, and for contraction and revision.

(a) *Modified Brewka and Hertzberg's Door-Window*: Recall that the definitional constraint (def) is bright $\leftrightarrow$ door_open $\vee$ window_open, and the problem arose with $S_0 = \{\text{bright}, \text{door\_open}, \neg\text{window\_open}\}$, and action A is to shut the door so post(A) is $\neg$door_open. The underlying theory H is (the closure of) S_0 itself (bright is not in $\mathcal{L}(H)$), and using the Levi identity, we revise H by post(A) to give $H^*_{\neg\text{door\_open}}$ which is $\{\neg\text{door\_open}, \neg\text{window\_open}\}$. Then expanding by (def =) bright $\leftrightarrow$ door_open $\vee$ window_open will give the revised state $K^*_{\neg\text{door\_open}} = Cn(H^*_{\neg\text{door\_open}} \cup \{\text{def}\})$, which is $\{\neg\text{bright}, \neg\text{door\_open}, \neg\text{window\_open}\}$.

(b) *Lifschitz's Switch*: In the switch example of Lifschitz, the constraint (def) is on_lamp $\leftrightarrow$ ((sw1 $\rightarrow$ sw2) $\wedge$ (sw2 $\rightarrow$ sw1)) and post(A) is $\neg$sw1. With initial state $\{\neg\text{on\_lamp}, \text{sw1}, \neg\text{sw2}\}$, $\{\neg on\_lamp, sw1, \neg sw2\}$, H is (the closure of) $\{\text{sw1}, \neg\text{sw2}\}$ — on_lamp is not in $\mathcal{L}(H)$ — so, using the Levi identity, the revision of H by post(A) gives $\{\neg\text{sw1}, \neg\text{sw2}\}$ in $H^*_{\neg\text{sw1}}$. Then expansion of $H^*_{\neg\text{sw1}}$ by def is $K^*_{\neg\text{sw1}}$ (the new state) also. Hence on_lamp is in the new state.

(c) *Myers and Smith's Cup and Baby*: The quasi-definition qdef here is on_table(cup) $\rightarrow$ safe_to_baby(cup), with initial state {on_table(cup), safe_to_baby(cup)}. Let d (the quasi-defined literal) be safe_to_baby(cup), so H is {on_table(cup)}, and post(A) is $\neg$on_table(cup) say. The literal safe_to_baby(cup) is not in $\mathcal{L}(H)$. $H^-_{\text{on\_table(cup)}}$, the contraction of H by on_table(cup) is empty. Then the new state, the expansion of $H^-_{\text{on\_table(cup)}}$ by qdef, is still empty. Note however, that this approach will not work if the constraint on_table(cup) $\rightarrow$ liftable(cup) is interpreted as a quasi-definition.

(d) *Yale Shooting Problem*: It can be similarly verified that the constraint alive $\rightarrow$ walking can be interpreted as a quasi-constraint, but not walking $\rightarrow$ alive.

6 Discussion and Future Work

To emphasize that incomplete descriptions are not the culprit we have only examined initial state descriptions that are complete. This enabled us to use the AGM framework as the update engine in the reduct of the theories for state change. However, the KM postulates replace the AGM postulates when incomplete descriptions are admitted. That proposition 4 survives for the KM update operator is plausible because of the connection between the KM update and AGM revision operators via maximal extensions and intersection. An application of proposition 4 to the modified Brewka and Hertzberg example when the initial state is given by door_open $\vee$ window_open and post(A) being $\neg$door_open gives $\{\neg\text{door\_open}\}$ as the KM update on the reduct H, and the correct result $\{\neg\text{door\_open}, \text{bright} \leftrightarrow \text{door\_open} \vee \text{window\_open}\}$ as the basis for the new state. In it the status of bright is unknown.

The fact that syntactically analogous constraints may not be interpreted identically as definitions or quasi-definitions is strong motivation for a careful ontological classification of different kinds of *causality* and *dependency* along the lines of Sandewall [13]. Intimately connected with this is their classical engineering notion of *control* which has not been seriously considered in logics of actions. These are semantic properties that account for the difference between constraints of the form $\text{on\_table(cup)} \rightarrow \text{safe\_to\_baby(cup)}$ and $\text{on\_table(cup)} \rightarrow \text{liftable(cup)}$. The consequent fluent in the first is controllable, but not in the second, so the first is causative but not the second, and a plausible design rule is to rule out non-causative constraints as quasi-definitions. These ideas deserve to be investigated closely in conjunction with the Sandewall ontological program.

Our proposal to perform state change on the reduct of a theory does not solve the problem of how to resolve nondeterminism in the reduct. The specification of entrenchment that may resolve this was not discussed. However, we believe this problem is one of system modelling, and is inherent in all logics.

REFERENCES

[1] G. Brewka. Preferred subtheories: An extended logical framework for default reasoning. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 1043–1048, 1989.

[2] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 1988.

[3] S. E. Dixon and W. Wobcke. The implementation of a first-order logic AGM belief revision system. In *Proceedings of the Fifth IEEE International Conference on Tools in Artificial Intelligence*, 1993.

[4] N. Y. Foo and Y. Zhang. Constraints and updates. In C. Zhang, J. Debenham, and D. Lukose, editors, *Proceedings of the Seventh Australian Joint Conference on Artificial Intelligence*, 1994.

[5] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* Bradford Books, MIT Press, Cambridge Massachusetts, 1988.

[6] M. L. Ginsberg and D. E. Smith. Reasoning about action (parts i and ii). In F. Brown, editor, *The Frame Problem in Artificial Intelligence.* Morgan Kaufmann Publishers, San Mateo, CA., 1987.

[7] G. N. Kartha and V. Lifschitz. Actions with indirect effects (preliminary report). In *Proceedings Fourth International Conference on Principles of Knowledge Representation and Reasoning.* Morgan Kaufmann Publishers, San Mateo, CA, 1994.

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

[9] V. Lifschitz. Circumscriptive theories: A logic-based framework for knowledge representation. *Journal of Phil. Logic*, 17:391–441, 1988.

[10] V. Lifschitz. Frames in the space of situations. *Artificial Intelligence*, 46:365–376, 1990.

[11] F. Lin and Y. Shoham. Provably correct theories of action. In *Proceedings of the 9th National Conference on Artificial Intelligence*, 1991.

[12] K. L. Myers and D. E. Smith. The persistence of derived information. *Proc. of the 7th National Conference on Artificial Intelligence*, 1988.

[13] E. Sandewall. *Features and Fluents.* Oxford University Press, 1994.

[14] J. R. Shonfield. *Mathematical Logic.* Addison-Wesley, 1967.

[15] M.-A. Williams. Transmutations of knowledge systems. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning*, pages 619–629. Morgan Kaufmann, 1994.

[16] M. Winslett. Reasoning about actions using a possible models approach. In *Proc Seventh National Artificial Intelligence Conference*, San Mateo, CA., 1988. Morgan Kaufmann Publishers.

[17] Y. Zhang and N. Y. Foo. The persistent set approach to reasoning about actions. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence.* France, 1993.