Revision by Translation
(short version)

Dov Gabbay, Odinaldo Rodrigues and Alessandra Russo

Abstract

In this paper, we show that it is possible to accomplish belief revision in any logic which is translatable to classical logic. We give the example of the propositional modal logic K and show that a belief operation in K defined in terms of $K's$ translation to classical logic verifies the AGM postulates.

Contents

1 Introduction

This paper will present a method for revision of theories in logics other than classical logic. The idea is to translate the other logic into first-order classical logic, perform the revision there and then translate back. The general schema looks as follows.

Let $*_a$ be a revision process in classical logic. Typically, given a classical logic theory Δ an input formula ψ[1], the operation $*_a$ gives us a new theory $\Gamma = \Delta *_a \psi$, corresponding to the result of the revision of Δ by ψ. Ideally, $*_a$ has some desirable properties, for instance, the well known AGM postulates for belief revision (see Section 2).

We would like to *export* this machinery to other logics. For example, given a theory Δ of some logic L and an input L-sentence ψ, can we define a revision operation $*_L$ such that $\Delta *_L \psi$ is a revised L theory and $*_L$ satisfies the AGM postulates? Can we make use of the revision operator of classical logic?

This paper presents such a method. The idea is to translate the object logic L into classical logic, perform the AGM revision there and translate the results back. Suppose that τ denotes a translation function from L into classical logic and T^τ is a classical logic theory encoding the basic properties of the logic L. If the axiomatization given by T^τ is sound and complete, we have that for all Δ and α of the logic L

$$\Delta \vdash_L \alpha \text{ iff } T^\tau \cup \Delta^\tau \vdash \alpha^\tau \tag{1}$$

Therefore, we can define a revision operator $*_L$ in the logic L as follows:

Definition 1 [Belief revision in L] Let $*_a$ be a revision operator for classical logic, and let τ, T^τ be as above. We define

$$\Delta *_L \psi = \{\beta \mid \Delta^\tau *_a (\psi^\tau \wedge T^\tau) \vdash \beta^\tau\}$$

The motivation for this definition is as follows. Δ^τ is the translation of the original L-logic theory Δ. Δ is to be revised by ψ, which in classical logic is translated as ψ^τ. We revise instead Δ^τ by ψ^τ. However, in classical logic the properties of the object logic (T^τ) have to be added as well, since it describes how the object logic works, and we want it to be preserved in the revision process (i.e. we want the resulting revised theory to satisfy T^τ), so we revise by $\psi^\tau \wedge T^\tau$[2].

Of course, the details have to be worked out. The difficulties mainly have to do with the notion of inconsistency in L. L may have theories Δ which are considered L-inconsistent while their translation Δ^τ is classically consistent[3].

[1] There is no special need that the input is a single formula ψ. It can be a theory Ψ. The AGM postulates work for input theories as well.

[2] Our method is restricted to logics L which have a translation τ which can be characterized by a classical theory T^τ. If the translation is, for example, semantically based, this means that the semantics of L can be expressed by a first-order theory T^τ, as is the case in many modal logics.

[3] In paraconsistent logics, for example, $p \wedge \neg p$ is considered inconsistent but we do not have $p \wedge \neg p \vdash q$. Equation (1) of the translation still holds, i.e., for all α, $p \wedge \neg p \vdash \alpha$ iff $(p \wedge \neg p)^\tau \cup T^\tau \vdash \alpha^\tau$, but in classical logic $(p \wedge \neg p)^\tau \cup T^\tau$ is consistent.

Thus, we may have a situtation in L where Δ is L-consistent, the input ψ is L-consistent, but $\Delta \cup \{\psi\}$ is L-inconsistent and requires L-revision. However, when we translate into classical logic we set $\Delta^\tau \cup \{\psi^\tau\} \cup T^\tau$ and this theory is classically consistent and so classical revision will be an expansion. We therefore need to write some additional axioms say Acc(for acceptability) in classical logic that will make $\Delta^\tau \cup \{\psi^\tau\} \cup T^\tau \cup Acc$ classically inconsistent, whenever $\Delta \cup \{\psi\}$ is L-inconsistent, and thus trigger a real revision process in classical logic.

There is a problem, however, with this approach. Classical revision of $\Delta^\tau \cup \{\psi^\tau\} \cup T^\tau \cup Acc$ may give us a theory Δ_c of classical logic such that the reverse translation $\Delta_L = \{\alpha \mid \Delta_c \vdash \alpha^\tau\}$ is not a theory we are happy with in L. Put in other words, when we look at the relation between $\Delta \cup \{\psi\}$ and Δ_L we are not happy in L to consider Δ_L as the L-revision of $\Delta \cup \{\psi\}$. The reason that such a situation may arise has to do with the fact that the notion of inconsistency in classical logic is too strong. We now explain why: if K is a consistent theory in classical logic and $K \cup \{\psi\}$ is inconsistent in classical logic, then $\mathrm{Cn}(K \cup \{\psi\})$ is the set of all wffs. Our revision intuition wants to take a consistent subset K' of $\mathrm{Cn}(K \cup \{\psi\})$.

In the logic L with K_L and ψ_L and with a different notion of inconsistency, the theory $\mathrm{Cn}_L(K_L \cup \{\psi_L\})$ may not be the set of all wffs of L. We still want to get a consistent subset K'_L of $\mathrm{Cn}_L(K_L \cup \{\psi_L\})$ as the revision. Our strategy of revision by translation may give us a revised theory via translation which is not a subset of $\mathrm{Cn}_L(K_L \cup \{\psi_L\})$ because in classical logic $\mathrm{Cn}(K_L^\tau \cup \{\psi_L^\tau\} \cup T^\tau \cup Acc)$ is too large (i.e., all wffs) and gives the revision process too much freedom. One way to solve this difficulty is to tighten up the revision process in classical logic.

The structure of the paper is as follows: in Section 2, we provide a quick introduction to the theory of Belief Revision. We analyse the meaning of these postulates for both classical and non-classical logics. This is followed in Section 3, by the application of the idea of revision by translation to the modal logic K. We finish the paper with some conclusions and comments in Section 4.

2 Belief Revision

The term Belief Revision is used to describe the kind of information change in which an agent reasoning about his beliefs about the world is forced to adjust them in face of new (possibly contradictory) information. One important assumption in the process is that the world is taken as a static entity. Even though changes in the world itself are not considered, the agent reasons about his knowledge about the world, which may be incorrect or incomplete. Therefore, Belief Revision is an intrinsically non-monotonic form of reasoning.

When the set of beliefs held by an agent is closed under the consequence relation of some formal language, it is usually called a *belief set*. Some variants of the standard belief revision approach also consider the case when the focus is done on a finite set of beliefs, called the *belief base*. These variants are usually called *base revision*. If Cn is the consequence relation of a given logic and K is a belief set, then it is assumed that $K = \mathrm{Cn}(K)$. Similarly, if for a belief φ and a belief set K, $\varphi \in K$, we say that φ is *accepted* in K.

The whole framework of Belief Revision is governed by some desiderata of the operations on belief sets, called *the AGM postulates for belief revision*. The term "AGM" stands for the initials of the main proposers of the theory, namely, Alchourrón, Gärdenfors and Makinson. According to the AGM theory [8], there are three main types of belief change:

- Expansion, when new information is consistent with the current belief set. All is necessary to do is to close the union of the previous belief set together with the new sentence under the consequence relation.
- Contraction, when the agent is forced to retract some beliefs. Notice that, since the belief set is closed under the consequence relation, in order to retract a belief φ, it is also necessary to remove other beliefs that imply φ.
- Revision, which is the acceptance of new information contradicting the current belief set and the subsequent process of restoring the consistency of that belief set whenever the new information is not itself contradictory.

Thus, the interesting cases are contractions and revisions. In fact, there are corresponding identities to translate between the two processes: the

Levi Identity defines revisions in terms of contractions and the *Harper Identity* defines contractions in terms of revisions. We will concentrate on the revision part here.

The general task of the revision process is to determine what is *rational* to support after a new contradictory belief is accepted. As we mentioned before, some general postulates describe ideal properties of the operation. One of these properties is sometimes referred to as the principle of *informational economy* [8, page 49]:

> "... when we change our beliefs, we want to retain as much as possible of our old beliefs – information is not in general gratuitous, and unnecessary losses of information are therefore to be avoided."

One of the main references to the general theory of belief revision is the book "Knowledge in Flux", by Peter Gärdenfors [8]. Other references include, for instance, [2, 3, 1, 9].

We now present the postulates for the revision operation as given in [8], pages 54–56. The following conventions are assumed: K is a set of formulae of the language representing the current belief set and A (B) is a formula representing the new piece of information. We use the symbol $*_a$ to denote an AGM belief revision operator. Thus, $K *_a A$ stands for the revision of K by A. The symbol $K_\perp$ denotes the inconsistent belief set, and is equivalent to the consequences of all formulae in the language.

AGM postulates for Belief Revision (in classical logic)

(K*1) $K *_a A$ is a belief set

This postulate requires that the result of the revision operation is also a belief set. One can perceive this as the requirement that the revised set be also closed under the consequence relation.

(K*2) $\quad A \in K *_a A$

(K*2) is known as the *success postulate* and corresponds to Dalal's *principle of primacy of the update* [4].

It basically says that the revision process should be successful in the sense that the new belief is effectively accepted in the revised belief state.

(K*3) $\quad K *_a A \subseteq \mathrm{Cn}(K \cup \{A\})$

(K*4) $\quad$ If $\neg A \notin K$, then $\mathrm{Cn}(K \cup \{A\}) \subseteq K *_a A$

(K*5) $\quad K *_a A = K_\perp$ only if A is unsatisfiable

To understand what the above three postulates (K*3)–(K*5) say, we need to consider two cases. Let $K_1 = K *_a A$.

Case 1: $K \cup \{A\}$ is consistent in classical logic.

In this case, AGM says that we want $K_1 = K *_a A$ to be equal to the closure of $K \cup \{A\}$:

- postulate (K*3) says that $K *_a A \subseteq \mathrm{Cn}(K \cup \{A\})$.
- postulate (K*4) says that $\mathrm{Cn}(K \cup \{A\}) \subseteq K *_a A$.
- postulate (K*5) is not applicable, since $K *_a A$ is consistent.

Case 2: $K \cup \{A\}$ is inconsistent.

In this case, let us see what the postulates (K*3)–(K*5) say about K_1:

- postulate (K*3) says nothing about K_1. If $K \cup \{A\}$ is classically inconsistent, then any theory whatsoever is a subset of $\mathrm{Cn}(K \cup \{A\})$, because this theory is the set of all formulae.
- postulate (K*4) says nothing. Since $K \cup \{A\}$ is inconsistent in classical logic, we have $\neg A \in K$ (since K is a closed theory), so (K*4) is satisfied, because it is an implication whose antecedent is false.
- To understand what postulate (K*5) says in our case, we distinguish two subcases:

(2.1) A is consistent.

(2.2) A is inconsistent.

Postulate (K*5) says nothing about $K_1 = K *_a A$ in case (2.2) above, it however requires K_1 to be consistent, whenever A is a consistent – case (2.2).

The above case analysis shows that the AGM postulates (K*3)–(K*5) have something to say only when $K \cup \{A\}$ is consistent, or if not when A is consistent. The particular way of writing these postulates as above makes use of technical properties of classical logic (the way inconsistent theories prove everything).

When we check the AGM postulates for logics other than classical, we may have a different notion of consistency and so we are free to interpret what we want the revision to do in the case of inconsistency according to what is reasonable in the object (non-classical) logic). AGM for classical logic gives us no clue beyond (K*5) as to what to require when $(K *_a A) \cup \{B\}$ is inconsistent.

Summary of (K*3)–(K*4)

Postulates (K*3)–(K*4) effectively mean the following:

($K^*_{3,4}$) If A is consistent with K, then $K{*_a}A = \mathrm{Cn}(K \cup \{A\})$.

If K is finite, we can take it as a formula and the postulate above corresponds to one of the rules in Katsuno and Mendelzon's rephrasing of the AGM postulates for finite knowledge bases ([10], page 187):

(R2) If $K \wedge A$ is satisfiable, then $K{*_a}A \leftrightarrow K \wedge A$.

For non-classical logics, where the notion of consistency is different, we need check only ($K^*_{3,4}$) and (K*5).

(K*6) If $A \equiv B$, then $K{*_a}A \equiv K{*_a}B$

(K*6) specifies that the revision process should be independent of the syntactic form of the sentences involved. It is called the *principle of irrelevance of syntax* by many authors, see for instance, [4].

(K*7) $K{*_a}(A \wedge B) \subseteq \mathrm{Cn}((K{*_a}A) \cup \{B\})$

(K*8) If $\neg B \notin K{*_a}A$, then $\mathrm{Cn}(K{*_a}A \cup \{B\}) \subseteq K{*_a}(A \wedge B)$

To understand what postulates (K*7)–(K*8) are saying, we again have to make a case analysis. The postulates have to do with the relationship of inputing (A, B) as a sequence (first revising by A, then expanding by B), as compared with revising by $\{A, B\}$ at the same time (i.e, revising by $A \wedge B$). It is well known that AGM does not say enough about sequences of revisions and their properties. These postulates are the bare minimum (see, for instance, [5, 6]).

We distinguish the following cases:

Case 1: A is consistent with K.

In this case, $K_1 = K{*_a}A$ is equal (by previous postulates) to $\mathrm{Cn}(K \cup \{A\})$.

(1.1) B is consistent with K_1. In this case, the antecendet of (K*8) holds and (K*7) and (K*8) together effectively say that $\mathrm{Cn}((K{*_a}A) \cup \{B\}) = K{*_a}(A \wedge B)$.

We can use previous postulates to say more than AGM says in this case, namely, that

$$(K{*_a}A){*_a}B = \mathrm{Cn}(K{*_a}A \cup \{B\}).$$

(1.2) B is inconsistent with $K_1 = K{*_a}A$, but B itself is consistent.

In this case, $\mathrm{Cn}(K{*_a}A \cup \{B\})$ is the set of all wffs.

- (K*7) holds because the right hand side of the inclusion is the set of all wffs and any other set of formulae is included in this set.
- (K*8) holds vacuously, since the antecedent of the implication is false.

(1.3) B is itself inconsistent.

- (K*7) requires that $K{*_a}(A \wedge B) \subset \mathrm{Cn}((K{*_a}A) \cup \{B\})$ and
- (K*8) holds vacuously.

The postulates say nothing new in this case, since the sets on either side of the inclusion in (K*7) are equal to the set of all wffs of the language and (K*8) is not applicable.

Case 2: A is not consistent with K, but A is itself consistent.

In this case, K_1 can be any consistent theory (by previous postulates), such that $A \in K_1$.

(2.1) B is consistent with K_1.

(2.2) B is inconsistent with K_1, but B itself is consistent.

(2.3) B is itself inconsistent.

These three cases follow, respectively, the same reasoning of cases (1.1), (1.2) and (1.3) above.
of all wffs.

Case 3: A is itself inconsistent.

In this case, $K *_a A$ is the set of all wffs of the language. Whether or not B is consistent is irrelevant in the postulates in this case. $\mathrm{Cn}(K *_a A \cup \{B\})$ is equal to the set of all wffs and as for case (1.2) above

- (K*7) holds because any set of wff is included in $\mathrm{Cn}(K *_a A \cup \{B\})$.
- (K*8) holds vacuously, since the antecedent of the implication is false.

Summary of (K*7)–(K*8)

Postulates (K*7)–(K*8) do not tell us anything new (beyond what we can deduce from earlier postulates), except in the case where B is consistent with $K *_a A$ (case 1.1), when (K*7) and (K*8) together are equivalent to the postulate below:

$(\mathrm{K}^*_{7,8})$ $\quad \mathrm{Cn}((K *_a A) \cup \{B\}) = K *_a (A \wedge B)$, when B is consistent with $K *_a A$

Therefore, for non-classical logics, we are committed only to $(\mathrm{K}^*_{7,8})$. Other cases involving inconsistency can have properties dictated by the local logic requirements.

(K*7) and (K*8) are the most interesting and controversial postulates. They capture in general terms the requirement that revisions are performed with a *minimal change* to the previous belief set. In order to understand them, recall that in a revision of K by A, one is interested in keeping as much as possible of the *informational content* of K and yet accept A. In semantical terms, this can be seen as looking for the models[4] of A that are somehow most *similar* to the models of the previous belief state K. The postulates do not constrain the operation well enough to give a precise meaning to the term similar, and this is how it should be, since they represent only general principles.

(K*7) says that if an interpretation I is among the models of A which are most similar to the models of K and it happens that I is also among the models of B, then I should also be among the models of $A \wedge B$ which are most similar to models of K.

[4]We consider models of a formula A, interpretations (or valuations) of the language which make A *true*.

Similarly, to understand the intuitive meaning of (K*8) consider the following situation: suppose that $(K *_a A) \wedge B$ is satisfiable. It follows that some models of A which are closest to models of K are also models of B. These models are obviously in $\mathrm{mod}(A \wedge B)$, since by (K*1), $\mathrm{mod}(K *_a A) \subseteq \mathrm{mod}(A)$. Now, every model in $\mathrm{mod}(A \wedge B)$ which is closest to models of K, must also be a model of $(K *_a A) \wedge B$.

This situation is depicted in Figure 1, where interpretations are represented around K according to their degree of similarity. The closer to $\mathrm{mod}(K)$ the more similar to K an interpretation is (the exact nature of this similarity notion is irrelevant to the understanding of the postulates).

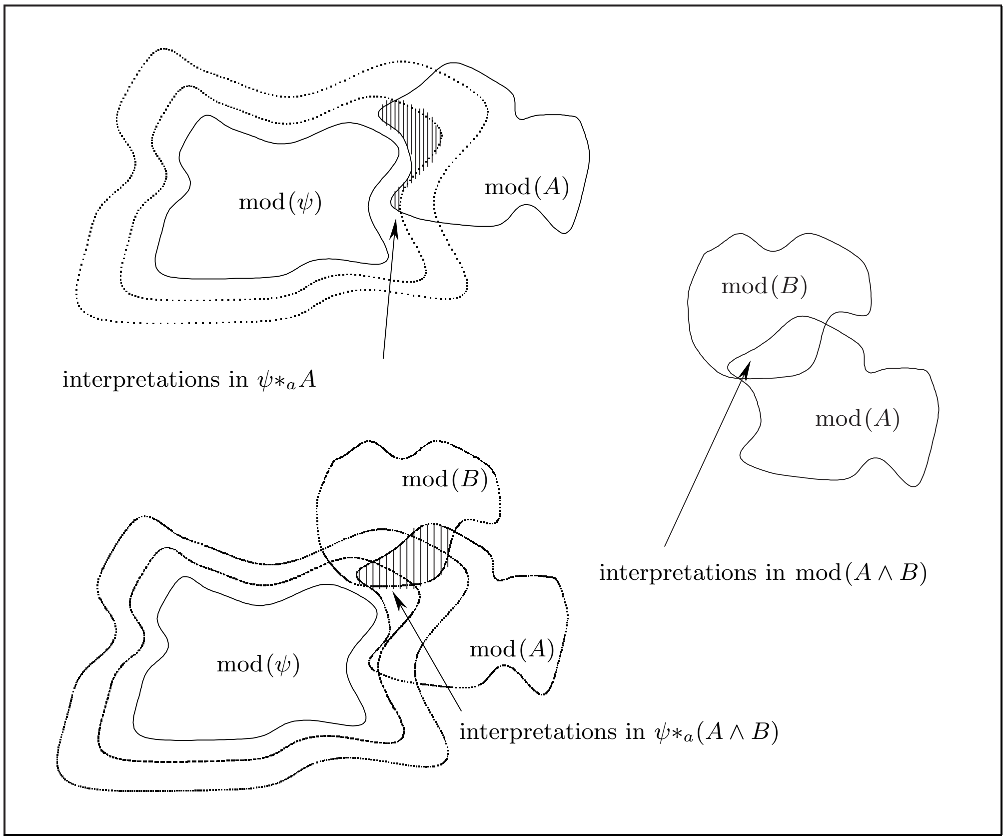

Figure 1: Illustrating postulate (K*8).

3 Revising in the modal logic K

We consider the case of the propositional modal logic K. The first thing we need to do is to provide the translation method from formulae and theories of K into formulae and theories of classical logic. This is done via the translation scheme described as follows.

Translation of modal K into classical logic

We need a binary predicate R in classical logic to represent the accessibility relation and unary predicates $P_1, P_2, P_3, \ldots$, for each propositional symbol p_i of K. We will use the subscript k whenever we wish to emphasize that we mean an operation (relation) in k and differentiate it from its classical logic counterpart (which will not be subscripted).

The idea is to encode the information of satisfiability of modal formulae by worlds into the variable of each unary predicate. In general, for a given world w and formula β the translation method can be stated as follows, where $\beta^\tau(w)$ represents $w \Vdash_k \beta$.

$$\begin{array}{rcl} p_i^\tau(w) & = & P_i(w) \\ (\neg\beta)^\tau(w) & = & \neg(\beta^\tau(w)) \\ (\beta \wedge \gamma)^\tau(w) & = & \beta^\tau(w) \wedge \gamma^\tau(w) \\ (\beta \to \gamma)^\tau(w) & = & \beta^\tau(w) \to \gamma^\tau(w) \\ (\Box\beta)^\tau(w) & = & \forall y(wRy \to \beta^\tau(y)) \end{array}$$

Finally, for a modal theory Δ, we define

$$\Delta^\tau(w) = \{\beta^\tau(w) \mid \beta \in \Delta\}.$$

We have, where w_0 is a completely new constant to Δ and β, and represents the actual world, that:

$\Delta \vdash_k \beta$ iff in every Kripke model with actual world w_0, we have $w_0 \Vdash_k \Delta$ implies $w_0 \Vdash_k \beta$ iff in classical logic we have that $\Delta^\tau(w_0) \vdash \beta^\tau(w_0)$.
(Correspondence)

$$\Delta \vdash_k \beta \text{ iff } T^\tau \cup \Delta^\tau(w_0) \vdash \beta^\tau(w_0) \tag{2}$$

The theory T^τ in the case of the logic K is empty (i.e., Truth)[5].

If Δ is finite we can let $\delta = \bigwedge \Delta$ and we have

$$\delta \vdash_k \beta \text{ iff } \vdash \forall x(\delta^\tau(x) \to \beta^\tau(x))$$

We can define a revision operator $*_k$ for K, as outlined before (we will omit the reference to the actual world w_0 in the rest of this section).

Definition 2 [Belief revision in K]

$$\Delta *_k \psi = \{\alpha \mid \Delta^\tau *_a (\psi^\tau \wedge T^\tau) \vdash \alpha^\tau\}$$

We can now speak more specifically of properties of $*_k$:

[5] The logic K imposes no properties on R. Had we been translating $S4$, we would have $T^\tau = \{\forall x(xRx) \wedge \forall x \forall y \forall z(xRy \wedge yRz \to xRz)\}$. Our notion also allows for non-normal logics, e.g., if w_0 is the actual world, we can allow reflexivity in w_0 by setting $T^\tau = \{w_0 R w_0\}$.

Properties of $*_k$:

1. $\Delta *_k \psi$ is closed under $\vdash$.

   This can be easily shown.

2. $\Delta *_k \psi \vdash_k \psi$.

   By (K*2), $\psi^\tau \wedge T^\tau \in \Delta^\tau *_a (\psi^\tau \wedge T^\tau)$. Since $\Delta^\tau *_a (\psi^\tau \wedge T^\tau)$ is closed under $\vdash$, $\Delta^\tau *_a (\psi^\tau \wedge T^\tau) \vdash \psi^\tau$ and hence $\psi \in \Delta *_k \psi$, by (K*1), $\Delta *_k \psi \vdash_k \psi$.

3. If ψ is (modally) consistent with Δ, then $\Delta *_k \psi = \mathrm{Cn}_k(\Delta \cup \{\psi\})$.

   We first show that if ψ is modally consistent with Δ, then $\Delta^\tau(k)$ is classically consistent with $\psi^\tau(k) \wedge T^\tau$. This holds because $\Delta \cup \{\psi\}$ has a Kripke model which will give rise to a classical model of the translation. Therefore, $\Theta = \Delta^\tau(k) *_a (\psi^\tau(k) \wedge T^\tau)$ is the classical provability closure of $\Delta^\tau(k) \cup \{\psi^\tau(k) \wedge T^\tau\}$.

   We now have to show that if $\alpha^\tau(k) \in \theta$ then $\Delta *_k \psi \vdash \alpha$.

   Lemma 1 *Let Δ be a closed theory. Let Δ^τ be its translation and let $\mathrm{Cn}(\Delta^\tau)$ be its T^τ-closure in classical logic. Let β be such that $\beta^\tau \in \mathrm{Cn}(\Delta^\tau)$. It follows that $\beta \in \Delta$.*

   Proof: If $\beta \notin \Delta$, then there exists a Kripke model of $\Delta \cup \{\neg\beta\}$. This gives a classical model of $\Delta^\tau \cup \{T^\tau\} \cup \{\neg\beta^\tau\}$, and so $\beta^\tau \notin \mathrm{Cn}(\Delta^\tau)$. □

   Lemma 2 *Let Δ^τ be a closed classical theory such that $\Delta^\tau \vdash T^\tau$ and let $\Delta = \{\beta \mid \beta^\tau \in \Delta^\tau\}$. Then if $\Delta \vdash \alpha$, then $\alpha^\tau \in \Delta^\tau$.*

   Proof: If $\alpha^\tau \notin \Delta^\tau$, there exists a model of $\Delta^\tau \cup \{\neg\beta^\tau\}$. This can be viewed as a kripke model of $\Delta \cup \{\neg\beta\}$. □

4. $\Delta *_k \psi$ is modally inconsistent, only if ψ is modally contradictory.

   If $\Delta *_k \psi$ is inconsistent, then so is $\Delta^\tau *_a (\psi^\tau \wedge T^\tau)$, since $\Delta^\tau *_a (\psi^\tau \wedge T^\tau)$ is closed under $\vdash$. By (K*5), $\psi^\tau \wedge T^\tau$ is contradictory. But by the correspondence, $\psi^\tau \wedge T^\tau \vdash \perp$ iff $\psi \vdash_k \perp$.

5. If $\psi \equiv_k \varphi$, then $\Delta *_k \psi \equiv \Delta *_k \varphi$.

   By correspondence ((2), page 9) and since $\psi \vdash_k \varphi$ and $\varphi \vdash_k \psi$, it follows that $T^\tau \cup \psi^\tau \equiv T^\tau \cup \varphi^\tau$. Therefore, by (K*6), $\Delta *_a (T^\tau \cup \psi^\tau) \equiv \Delta *_a (T^\tau \cup \varphi^\tau)$ and hence $\Delta *_k \psi \equiv \Delta *_k \varphi$.

6. $\Delta *_k (\psi \wedge \varphi) \subseteq \mathrm{Cn}_k((\Delta *_k \psi) \cup \{\varphi\})$.

   Suppose that $\Delta *_k (\psi \wedge \varphi) \vdash_k \alpha$, for some α. By the definition of $*_k$, $\Delta^\tau *_a (\psi^\tau \wedge \varphi^\tau \wedge T^\tau) \vdash \alpha^\tau$. By (K*7), it follows that $\mathrm{Cn}(\Delta^\tau *_a (\psi^\tau \wedge T^\tau) \cup \{\varphi^\tau\} \vdash \alpha^\tau$. Notice that for every $\gamma^\tau \in \Delta^\tau *_a (\psi^\tau \wedge T^\tau)$, there is a corresponding γ in $\Delta *_k \psi$ (by the definition of $*_k$) and similarly for φ^τ. By correspondence, $\mathrm{Cn}_k((\Delta *_k \psi) \cup \{\varphi\}) \vdash \alpha$.

7. If φ is modally consistent with $\Delta *_k \psi$, then $\mathrm{Cn}((\Delta *_k \psi) \cup \{\varphi\}) \subseteq \Delta *_k (\psi \wedge \varphi)$.

   If φ is modally consistent with $\Delta *_k \psi$, then φ^τ is modally consistent with $\Delta *_a (\psi^\tau \wedge T^\tau)$ and then by (K*8), $\mathrm{Cn}(\Delta *_a (\psi^\tau \wedge T^\tau) \cup \{\varphi^\tau\}) \subseteq \mathrm{Cn}(\Delta^\tau *_a ((\psi \wedge \varphi)^\tau \wedge T^\tau))$. But, $\{\alpha \mid \Delta^\tau *_a (\psi^\tau \wedge T^\tau) \vdash \alpha^\tau\} \cup \{\varphi\} \vdash_k \beta$ iff $\Delta^\tau *_a (\psi^\tau \wedge T^\tau) \cup \{\varphi^\tau\} \vdash \beta^\tau$.

What we have just proven is that $*_k$ verifies all eight conditions of an AGM operation. Item 3, is actually a proof for postulates (K*3) and (K*4).

The actual process of revision briefly discussed in the example above may be more complex than this. For instance, the object language might have a consistency notion other than that of classical logic or even none at all.

4 Conclusions

We have presented a way of exporting an AGM revision process in classical logic to other non-classical logics by translating into classical logic. There are considerable benefits to such a revision by translation over any direct revision in the non-classical logic itself.

1. It is a standard for many non-classical logics to be translated into classical logic. Such translations are done for a variety of reasons:
   - to give the logic a meaning
   - to give semantics to the logic
   - to compare it with other logics
   - to get decidability/undecidability results
   - to make use of automated deduction of classical logic

   Adding to this culture a revision capability makes sense.
2. Revision theory is very well developed in classical logic. There are various notions and fine tuning involved and translation into classical logic will not only open a wealth of distinctions for the source logic but also enrich classical logic revision itself with new ideas and problems arising from non-classical logics.
3. From the point of view of classical logic, what we are doing is a *relative revision*. This concept can be defined as follows. Given a theory T of classical logic (for example a theory of linearly ordered Abelian groups[6]), we can talk about a set of sentences Δ being acceptable (i.e., $\mathrm{Cn}(T \cup \Delta)$ being acceptable)[7]. We can talk about revision relative to T of Δ by ψ (denoted by $\Delta *_T \psi$) yielding a new acceptable theory $\Delta' = \Delta *_T \psi$. The case where T arises in connection with a translation from another logic is only one instance of this general relative revision.

We need to distinguish two cases in our studies of relative revision

(a) The notion of acceptability can be easily handled in classical logic.

This case is simpler to handle and includes the translation from modal logics seen in Section 3.

b) The concept of acceptability is not directly expressible. Here we have problems to overcome (see the Belnap's logic translations, in the long version of this paper [7]).

[6]We are preparing here for the future, when we examine revision in fuzzy logic.

[7]Let $\mathcal{M}$ be a class of models of T. Δ is T-consistent if $\Delta \cup T$ has a model. Δ is $\mathcal{M}$-acceptable if $\Delta \cup T$ has a model in $\mathcal{M}$. Given Δ and ψ, we want the result of the revision of Δ by ψ to be $\mathcal{M}$-acceptable.

We leave the investigations of relative revision for a future paper.

5 Acknowledgments

Dov Gabbay, Odinaldo Rodrigues and Alessandra Russo were supported by the UK EPSRC projects "Temporal Logics of Knowledge and Belief" (GR/K57268), "Data Driven Algorithms" (GR/L91818) and "Managing Inconsistency in Software Engineering" (GR/L55964), respectively.

References

[1] C. A. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50:510–530, 1985.

[2] C. A. Alchourrón and D. Makinson. On the logic of theory change: Contraction functions and their associated revision functions. *Theoria*, 48:14–37, 1982.

[3] C. A. Alchourrón and D. Makinson. The logic of theory change: Safe contraction. *Studia Logica*, 44:405–422, 1985.

[4] M. Dalal. Updates in propositional databases. Technical report, Department of Computer Science, Rutgers University, February 1988. Technical Report DCS-TR-222.

[5] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. Technical Report R-202, Cognitive Science Laboratory, Computer Science Department, University of California, Los Angeles, CA 90024, November 1996.

[6] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

[7] Dov Gabbay, Odinaldo Rodrigues, and Alessandra Russo. *"Information, Uncertainty, Fusion"*, chapter Revision by Translation. "Kluwer Scientific Publishers", 1999. "To appear".

[8] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. A Bradford Book - The MIT Press, Cambridge, Massachusetts - London, England, 1988.

[9] Peter Gärdenfors, editor. *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1992.

[10] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. *Belief Revision*, pages 183–203, 1992.