A Dynamic Logic of Iterated Belief Change

Bernd van Linder
Utrecht University
Department of Computer Science
P.O. Box 80.089
3508 TB Utrecht
The Netherlands
Email: bernd@cs.ruu.nl

Abstract

In this paper we propose a general framework to study iterated changes of belief. Expansions, contractions and revisions are taken as actions which may be performed by an agent, resulting in a change of its beliefs. The syntax of the framework is given by a multi-modal language, containing modalities to reason about the agent's knowledge — representing its non-defeasible, veridical information — and belief — representing its defeasible, non-veridical, working information — as well as a modality to reason about the results of belief-changing actions. The semantics is defined using Kripke-style possible worlds models. The formalisation of single-step belief changes is proved to be sound and complete with regard to the respective AGM axiomatisations. Iterated changes of belief correspond to sequences of belief-changing actions. When dealing with sequences of actions, we allow the result of a belief-changing action to depend on other factors besides the set of beliefs and the formula that this set is changed with, thereby adopting a dynamic view on iterated belief change. Several desirable properties of iterated belief change are proved to be valid in our formalisation. Furthermore the validity of various recently proposed postulates for iterated belief revision is checked.

1. Introduction

The possibility to change currently held beliefs upon acquiring new information is a typical element of intelligent behaviour. The probably best known and most prominent formal approach towards belief change is the so called AGM framework as proposed by Alchourrón, Gärdenfors and Makinson [1, 5]. In the AGM framework rationality postulates are proposed for three kinds of belief changes. The first of these is the *expansion* through which some formula is added to a set of beliefs regardless of whether the resulting set is consistent. Through a *contraction* some formula is retracted from a belief set, and *revisions* add some formula to a set of beliefs but in order to maintain consistency of the resulting belief set it might be necessary to remove some of the old formulae in the set. Recent research indicates that the AGM framework, which explicitly deals with single-step belief changes only, is not completely suitable for iterated changes of belief. Various modifications of, and alternatives to, this framework have been proposed [2, 3, 4, 10, 14, 16], all aimed at providing an intuitively acceptable account of *iterated belief revision*.

The approach that we propose in this paper is a more general one than the approaches mentioned above. Instead of focusing on iterated belief revisions, we provide a formalisation of *arbitrary sequences of belief changes*, which may consist of expansions, contractions and revisions. Syntactically our system consists of a highly expressive multi-modal language, which combines notions from epistemic, doxastic and dynamic logic and allows for a concise and intelligible representation of all kinds of postulates for (iterated) belief change. For the semantics we employ Kripke-style possible worlds models. To adequately deal with iterated changes of belief we extend the models that are used for single-step belief changes. Whereas the original models formalise *belief sets*, these extended models correspond to *belief systems*, consisting of a set of beliefs and a method for changing the belief set upon execution of belief-changing actions. The resulting system may be seen as providing a *dynamic* logic of iterated belief change, in which an agent continually changes its beliefs by performing belief-changing actions.

The fragment of our framework that deals with non-iterated, single-step belief changes, is in spirit similar to the approach proposed by Segerberg [15], though in effect different, the most notable difference being the fact that all AGM postulates are validated in our framework (see §2.1) whereas this is not the case in Segerberg's. The semantics that we present for the non-iterated fragment, resembles the spheres semantics proposed by Grove [6] to model belief revision. However, whereas Grove uses the representation of the original AGM framework, the system that we define allows for a concise formulation of the AGM postulates in a modal language. Nevertheless, the similarity on the semantic level makes that our approach towards iterated belief change could easily be adapted to make the spheres approach suitable to deal with iterated revision.

Having a well-defined semantics to model belief change has the advantage that it is possible to *prove validities* — thereby *checking* postulates instead of just *proposing* these — that characterise (iterated) changes of belief. This possibility of checking and validating postulates is used to investigate various axiomatisations of iterated belief revision that have been proposed recently [3, 10, 14].

The rest of the paper is organised as follows. In §2 we look into the basic framework for single-step belief changes and prove that it is sound and complete with respect to the AGM axiomatisation. In §3 the formal approach towards iterated belief change is presented. In §4 we provide some validities that describe sequences of belief changes, and we look into various other postulates that have been proposed recently. In §5 we round off.

2. Actions that change your mind

The general setting of our framework is that of an agent[1], which *knows* some formulae, *believes* other formulae and may change its set of beliefs by performing special, *belief-changing* actions. We adopt a *Platonic* view on knowledge, i.e. the knowledge of the agent represents the *veridical* information that it is born or built with, and that neither grows nor shrinks. The beliefs of the agent comprise its knowledge and represent its *working information*. This working information may grow, sometimes not justifiedly, in which case it should be possible to retract certain beliefs. Thus beliefs are *defeasible* and *non-veridical*. In representing the agent's knowledge and belief we follow the approach common in epistemic and doxastic logic [8]: the formula $\mathbf{K}\varphi$ denotes the fact that the agent knows φ, and $\mathbf{B}\varphi$ that it believes φ. As the result of the execution of belief-changing actions, the agent expands, contracts and revises its beliefs. In the spirit of the AGM framework, we consider expansions and contractions to be fundamental, and define revisions in terms of these. Formulae $[\alpha]\varphi$ formalise that φ is brought about as the result of doing α[2].

2.1. Definition. For a denumerable set Π of propositional symbols, the language $\mathcal{L}(\Pi)$ is defined to be the smallest superset of Π such that

- if $\varphi, \psi \in \mathcal{L}, \alpha \in Ac$ then $\neg\varphi, \varphi \wedge \psi, \mathbf{K}\varphi, \mathbf{B}\varphi, [\alpha]\varphi \in \mathcal{L}$

where Ac is the smallest superset of $\{\texttt{expand}\ \varphi, \texttt{contract}\ \varphi \mid \varphi$ is purely propositional$\}$ closed under sequential composition, denoted by ;.

The purely propositional, non-modal, fragment of $\mathcal{L}(\Pi)$ is denoted by $\mathcal{L}_0(\Pi)$. The constructs $\vee, \rightarrow, \leftrightarrow, \top$ and $\bot$ are defined in the usual way; $\mathbf{M}\varphi$ is defined to be $\neg\mathbf{K}\neg\varphi$ and $\texttt{revise}\ \varphi$ is $\texttt{contract}\ \neg\varphi; \texttt{expand}\ \varphi$. When the set Π of propositional symbols is understood, which we assume to be the case in the rest of this paper unless explicitly stated otherwise, we write $\mathcal{L}$ and $\mathcal{L}_0$ rather than $\mathcal{L}(\Pi)$ and $\mathcal{L}_0(\Pi)$.

The definition of the `revise` action is a straightforward implementation of the Levi identity [11], which states that revisions can be defined in terms of contractions and expansions, viz. a revision with some formula φ can be brought about as a contraction with $\neg\varphi$ followed by an expansion with φ (see §2.1).

The semantics for $\mathcal{L}$ is defined using Kripke-style possible worlds models. These models are *universal* with regard to knowledge; belief is interpreted using a subset of the set of worlds, viz. the set of *doxastic alternatives*.

2.2. Definition. The class $\mathcal{M}$ of Kripke models contains all tuples $\mathrm{M} = \langle \mathrm{S}, \mathrm{B} \rangle$ such that

- $\mathrm{S} \subseteq 2^{\Pi}$ is a non-empty set of possible worlds, or states.
- $\mathrm{B} \subseteq \mathrm{S}$ is the set of doxastic alternatives.

The binary relation $\models$ between a formula φ and a pair M, s consisting of a model in $\mathcal{M}$ and a state s in M is inductively defined by:

[1] Throughout this paper we assume the agent to be neuter.

[2] Since the framework is such that all actions are both deterministic, i.e. $\langle\alpha\rangle\varphi \rightarrow [\alpha]\varphi$ is valid, and realizable, i.e. $\langle\alpha\rangle\top$ is valid, it suffices to consider only formulae of this kind: $[\alpha]\varphi$ and $\langle\alpha\rangle\varphi$ are equivalent.

$$\begin{array}{lll}
\mathrm{M}, s \models p & \Leftrightarrow p \in s & \text{for } p \in \Pi \\
\mathrm{M}, s \models \neg\varphi & \Leftrightarrow \mathrm{M}, s \not\models \varphi & \\
\mathrm{M}, s \models \varphi \wedge \psi & \Leftrightarrow \mathrm{M}, s \models \varphi \,\&\, \mathrm{M}, s \models \psi & \\
\mathrm{M}, s \models \mathbf{K}\varphi & \Leftrightarrow \forall s' \in \mathrm{S}[\mathrm{M}, s' \models \varphi] & \\
\mathrm{M}, s \models \mathbf{B}\varphi & \Leftrightarrow \forall s' \in \mathrm{B}[\mathrm{M}, s' \models \varphi] & \\
\mathrm{M}, s \models [\alpha]\varphi & \Leftrightarrow \mathtt{r}(\alpha, \mathrm{M}), s \models \varphi &
\end{array}$$

where $\mathtt{r}$ is defined by:

$$\begin{array}{ll}
\mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}) & = \langle \mathrm{S}, \mathrm{B} \setminus \{s \in \mathrm{S} \mid \mathrm{M}, s \models \neg\varphi\}\rangle \\
\mathtt{r}(\mathtt{contract}\ \varphi, \mathrm{M}) & = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\varphi)\rangle \\
\mathtt{r}(\alpha_1; \alpha_2, \mathrm{M}) & = \mathtt{r}(\alpha_2, \mathtt{r}(\alpha_1, \mathrm{M}))
\end{array}$$

where σ is a fixed *selection function* for M (see Definition 2.3).

If $\mathtt{r}(\mathtt{contract}\ \varphi, \mathrm{M})$ is for all $\varphi \in \mathcal{L}_0$ defined using the selection function σ we call $\mathtt{r}$ to be *based on* σ, or σ*-based*. The formula φ is *satisfiable* in the model M iff $\mathrm{M}, s \models \varphi$ for some $s \in \mathrm{M}$; φ is satisfiable in $\mathcal{M}$ iff φ is satisfiable in some $\mathrm{M} \in \mathcal{M}$. The formula φ is *valid* in the model M, denoted by $\mathrm{M} \models \varphi$, iff $\mathrm{M}, s \models \varphi$ for all s in M; φ is valid in $\mathcal{M}$, denoted by $\models \varphi$, iff φ is valid in all $\mathrm{M} \in \mathcal{M}$. For reasons of convenience we define $[\![\varphi]\!]_{\mathrm{M}} = \{s \in \mathrm{S} \mid \mathrm{M}, s \models \varphi\}$ to denote the set of states that satisfy φ. Whenever the model M is clear from the context, we drop the subscript and denote $[\![\varphi]\!]_{\mathrm{M}}$ simply by $[\![\varphi]\!]$.

2.3. Definition. Let some model $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}\rangle$ be given. A total function $\sigma : \mathcal{L}_0 \to \wp(\mathrm{S})$ is a *selection function* for M if and only if it meets the following constraints for all $\varphi, \psi \in \mathcal{L}_0$.

Σ1. $\sigma(\varphi) \subseteq [\![\varphi]\!]$
Σ2. $\sigma(\varphi) = \emptyset$ iff $\mathrm{B} \cap [\![\varphi]\!] \neq \emptyset$ or $[\![\varphi]\!] = \emptyset$
Σ3. if $[\![\varphi]\!] = [\![\psi]\!]$ then $\sigma(\varphi) = \sigma(\psi)$
Σ4. $\sigma(\varphi \vee \psi) \subseteq \sigma(\varphi) \cup \sigma(\psi)$
Σ5. if $\sigma(\varphi \vee \psi) \cap [\![\varphi]\!] \neq \emptyset$ then $\sigma(\varphi) \subseteq \sigma(\varphi \vee \psi)$

2.4. Remark. The definition of selection functions as presented above may be seen as the implementation of a suggestion by Grove [6]. Grove uses a system of *spheres* surrounding a belief set to define belief revision. Grove indicates that with every system of spheres some function can be associated that 'selects the 'closest' worlds' ([6], p. 159) in which some formula holds. This selection of the closest worlds makes that this function is a selection function in the sense of Definition 2.3. Since selection functions are not defined with the idea of selecting worlds that are closest according to some criterion (though in effect they do something very similar), it is in fact not easy to construct a system of spheres when given a selection function. Nevertheless, combining the soundness theorems of §2.1 with the completeness theorems for the spheres approach ensures that it is in principle possible to associate a system of spheres with every selection function.

Interpreting the **K** and **B** operators as in Definition 2.2 leads to a notion of knowledge that satisfies an S5 axiomatisation and a notion of belief that satisfies a K45 axiomatisation. This means in particular that knowledge and belief are normal modal operators satisfying the axioms of positive and negative introspection, and that knowledge is furthermore veridical. Since our formalisation of belief-changing actions presupposes the existence of inconsistent belief sets, our notion of belief does not satisfy the D-axiom $\neg(\mathbf{B}\varphi \wedge \mathbf{B}\neg\varphi)$ (this in contrast with the more common approach to belief of Kraus & Lehmann [9]).

The definition of `r` for the `expand` action is based on the idea that uncertainties of an agent are formalised through its set of doxastic alternatives, and expansions put an end to uncertainties. The expansion of the beliefs of the agent with a formula φ may be implemented by *removing all states supporting* $\neg\varphi$ *from the set of doxastic alternatives.*

The definition of `r` for the `contract` action is based on the idea that apparent beliefs that an agent has are turned into doubts as the result of a contraction. This idea can be formalised by implementing a contraction with some formula φ as *extending* the set of doxastic alternatives such that it *encompasses at least one state not satisfying* φ*.* The problem with contractions defined in this way is that it is not straightforward to decide *which* worlds need to be added, and this is where selection functions come into play. A selection function picks out 'reasonable' states satisfying the negation of the formula that is to be contracted, i.e. the set of selected states should be such that the resulting `contract` action behaves in a reasonable, intuitively acceptably way.

Belief-changing actions are interpreted as *model-transformers* rather than the more common *state-transitions*, which is the usual interpretation for actions in dynamic logic [7]. The reason for this lies in the fact that belief-changing actions do not change the state of the world in which the agent resides, but the *doxastic state* of the agent, which is formalised through the set of doxastic alternatives in the model. Therefore a change in the agent's doxastic state is interpreted as a change in the set of doxastic alternatives, which works out in a change of the model under consideration.

2.5. Convention. The relation $\vdash_{cpl}$ denotes derivability in classical propositional logic. The function $\mathrm{Th} : 2^{\mathcal{L}_0} \to 2^{\mathcal{L}_0}$ is the deductive closure operator associated with $\vdash_{cpl}$. In the sequel, we let $\models_{cpl}$ denote both validity and semantic entailment in classical propositional logic: for $\pi \in 2^{\Pi}$, $\varphi \in \mathcal{L}_0$ and $\Phi \subseteq \mathcal{L}_0$, all of $\pi \models_{cpl} \varphi$, $\pi \models_{cpl} \Phi$ and $\Phi \models_{cpl} \varphi$ have their usual connotation. The soundness and completeness properties of classical propositional logic will be used freely.

The following propositions summarise some general properties that turn out to be useful in the rest of this paper.

2.6. Proposition. *For all* $\alpha, \alpha_1, \alpha_2, \alpha_3 \in Ac$, $\varphi \in \mathcal{L}$, $\psi \in \mathcal{L}_0$ *and models* M *with state* s *we have:*

1. $\mathrm{M}, s \models \psi \Leftrightarrow s \models_{cpl} \psi$
2. $\mathrm{M}, s \models \mathbf{K}\varphi \Leftrightarrow \mathrm{M} \models \varphi \Leftrightarrow \mathrm{M} \models \mathbf{K}\varphi$
3. $\mathrm{M}, s \models \mathbf{B}\varphi \Leftrightarrow \mathrm{M} \models \mathbf{B}\varphi$
4. $\models [\alpha]\neg\varphi \leftrightarrow \neg[\alpha]\varphi$
5. $\models [\alpha_1; \alpha_2]\varphi \leftrightarrow [\alpha_1][\alpha_2]\varphi$
6. $\models [(\alpha_1; \alpha_2); \alpha_3]\varphi \leftrightarrow [\alpha_1; (\alpha_2; \alpha_3)]\varphi$

7. $\models ([\alpha](\varphi \to \psi) \wedge [\alpha]\varphi) \to [\alpha]\psi$
8. $\models \varphi \Rightarrow \models [\alpha]\varphi$

Items two and three imply that whenever the formula $\mathbf{B}\varphi$ is true in a state of a model, the formula $\mathbf{KB}\varphi$ is valid in the model. Furthermore, whenever $\mathbf{M}\varphi$ is true in M, s then $\mathrm{M} \models \mathbf{M}\varphi$ and hence $\mathrm{M} \models \mathbf{KM}\varphi$; an analogous result holds for $\neg\mathbf{B}\neg\varphi$. The last two items of Proposition 2.6 express that the box operator $[\alpha]$ is a normal modal operator, i.e. it satisfies the K-axiom and the rule of necessitation. Hence despite its non-standard, model-transforming interpretation, the operator $[\alpha]$ still behaves normal. There is however one difference between the interpretation of actions as state-transitions and that as model-transformers. For the usual interpretation of an action α as a transition between states, it holds that whenever some formula φ is valid in a model, $[\alpha]\varphi$ is also valid in the model. This property is in fact obvious, given the fact that execution of such an action causes only a transition to a state (in which φ holds), and does not change the model in any way. This property does however not hold for actions that are interpreted as model-transformers. Consider for example the model $\mathrm{M} = \langle \mathrm{S}, \mathrm{B} \rangle$ such that $\mathrm{B} = \{s, s'\}$ and $p \in s, p \notin s'$. Then $\mathrm{M} \models \neg(\mathbf{B}p \vee \mathbf{B}\neg p)$, whereas $\mathrm{M} \not\models [\texttt{expand}\ p]\neg(\mathbf{B}p \vee \mathbf{B}\neg p)$. The last two items of Proposition 2.6 indicate however, that even though $\forall \mathrm{M}(\mathrm{M} \models \varphi \Rightarrow \mathrm{M} \models [\alpha]\varphi)$ does not hold when interpreting actions as model-transformers, $\forall \mathrm{M}(\mathrm{M} \models \varphi) \Rightarrow \forall \mathrm{M}(\mathrm{M} \models [\alpha]\varphi)$ does.

2.7. Proposition. *Let* $\mathrm{M} = \langle \mathrm{S}, \mathrm{B} \rangle$ *be some Kripke model, and let* $\mathtt{r}$ *be* σ*-based for some selection function* σ *for* M. *For all* $\varphi \in \mathcal{L}_0$ *we have:*

$$\mathtt{r}(\mathtt{revise}\ \varphi, \mathrm{M}) = \langle \mathrm{S}, (\mathrm{B} \cup \sigma(\varphi)) \cap [\![\varphi]\!] \rangle$$

2.1. The AGM framework

In the AGM framework it is investigated how rational changes to the set of beliefs of an agent should work out. Besides a unique characterisation of expansions, rationality postulates are proposed that constrain revisions and contractions. These postulates are defined in terms of a belief set $K \subseteq \mathcal{L}_0$ and a formula $\varphi \in \mathcal{L}_0$: $K^+_\varphi, K^-_\varphi, K^*_\varphi$ denote respectively the expansion, contraction and revision of K with φ.

2.8. Definition. A set $\Phi \subseteq \mathcal{L}_0$ is an *AGM belief set* iff $\Phi = \mathrm{Th}(\Phi)$, i.e. Φ is closed under the derivability operator of classical propositional logic. The absurd belief set, consisting of all formulae from $\mathcal{L}_0$, is denoted by $K_\perp$.

2.9. Definition. The expansion of an AGM belief set K with a formula $\varphi \in \mathcal{L}_0$ is defined to be the logical closure of K and φ, i.e. $K^+_\varphi = \mathrm{Th}(K \cup \{\varphi\})$.

2.10. Definition. For K an AGM belief set and $\varphi, \psi \in \mathcal{L}_0$ the following are the Gärdenfors postulates for belief contraction:
(G^-1) K^-_φ is an AGM belief set.
(G^-2) $K^-_\varphi \subseteq K$.

($G^{-}3$) If $\varphi \notin K$ then $K^{-}_{\varphi} = K$.
($G^{-}4$) If $\nvdash_{cpl} \varphi$ then $\varphi \notin K^{-}_{\varphi}$.
($G^{-}5$) If $\varphi \in K$ then $K \subseteq (K^{-}_{\varphi})^{+}_{\varphi}$.
($G^{-}6$) If $\vdash_{cpl} \varphi \leftrightarrow \psi$ then $K^{-}_{\varphi} = K^{-}_{\psi}$.
($G^{-}7$) $K^{-}_{\varphi} \cap K^{-}_{\psi} \subseteq K^{-}_{\varphi \wedge \psi}$.
($G^{-}8$) If $\varphi \notin K^{-}_{\varphi \wedge \psi}$ then $K^{-}_{\varphi \wedge \psi} \subseteq K^{-}_{\varphi}$.

2.11. Definition. For K an AGM belief set and $\varphi, \psi \in \mathcal{L}_0$ the following are the Gärdenfors postulates for belief revision:
($G^{*}1$) K^{*}_{φ} is an AGM belief set.
($G^{*}2$) $\varphi \in K^{*}_{\varphi}$.
($G^{*}3$) $K^{*}_{\varphi} \subseteq K^{+}_{\varphi}$.
($G^{*}4$) If $\neg\varphi \notin K$ then $K^{+}_{\varphi} \subseteq K^{*}_{\varphi}$.
($G^{*}5$) $K^{*}_{\varphi} = K_{\perp}$ if and only if $\vdash_{cpl} \neg\varphi$.
($G^{*}6$) If $\vdash_{cpl} \varphi \leftrightarrow \psi$ then $K^{*}_{\varphi} = K^{*}_{\psi}$.
($G^{*}7$) $K^{*}_{\varphi \wedge \psi} \subseteq (K^{*}_{\varphi})^{+}_{\psi}$.
($G^{*}8$) If $\neg\psi \notin K^{*}_{\varphi}$, then $(K^{*}_{\varphi})^{+}_{\psi} \subseteq K^{*}_{\varphi \wedge \psi}$.

We call a contraction operator $\Leftrightarrow$ (revision operator $*$) an *AGM contraction* (*AGM revision*) iff it satisfies the AGM postulates for contraction (revision).

Given the AGM postulates for revision and contraction, these two operators turn out to be inter-definable, i.e. starting from an AGM revision one can construct an AGM contraction and *vice versa.* The Levi-identity $K^{*}_{\varphi} = (K^{-}_{\neg\varphi})^{+}_{\varphi}$ constructs AGM revisions out of AGM contractions and the Harper-identity $K^{-}_{\varphi} = K \cap K^{*}_{\neg\varphi}$ may be used to construct AGM contractions out of AGM revisions. Whenever a given AGM revision operator $*$ is used to construct an AGM contraction $\Leftrightarrow$ via the Harper-identity, the revision operator resulting from applying the Levi-identity to $\Leftrightarrow$ is exactly the revision operator $*$ that one started with. Also when defining a contraction operator by applying the Harper-identity to an AGM revision $*$ that resulted from an AGM contraction $\Leftrightarrow$ via the Harper identity, the two contraction operators are identical. The complete inter-definability of revisions and contractions was already used in the definition of `r` for the `revise` action, and will also be used in one of the completeness proofs presented below.

The concept of AGM belief sets is easily incorporated in our modal framework.

2.12. Definition. For a Kripke model M, the belief set associated with M, notation B(M), is the set $\{\varphi \in \mathcal{L}_0 \mid \mathrm{M} \models \mathbf{B}\varphi\}$.

2.13. Proposition. *For any model* M*, the set* B(M) *is an AGM belief set.*

Proof: It is easily seen that under the definition of $\models$ as presented in 2.2, B(M) is indeed deductively closed, for any $\mathrm{M} \in \mathcal{M}$. □

Since model-based belief sets are AGM belief sets, from now on we will use these notions interchangeably.

To achieve a complete correspondence between belief-changing actions in our framework and AGM belief changes, a restriction to so called *full models* is necessary. These are models in which every possible valuation on propositional symbols occurs as a state.

2.14. DEFINITION. A Kripke model M is a *full* model if and only if $\mathrm{S} = 2^{\Pi}$.

It is easily seen that in a full model M a formula $\varphi \in \mathcal{L}_0$ is known iff the formula is a theorem in classical propositional logic. For φ is known in M iff it holds at all states in M, i.e. φ is satisfied by every valuation on propositional symbols and hence φ is a theorem of classical propositional logic.

2.15. PROPOSITION. *For every AGM belief set K, a unique full model* M *exists such that $K = \mathrm{B}(\mathrm{M})$. This model, denoted by M_K, is referred to as the K-model.*

PROOF: Let K be some AGM belief set, i.e. K is a deductively closed set of propositional formulae. Define the model $\mathrm{M}_K = \langle \mathrm{S}, \mathrm{B} \rangle$ where $\mathrm{S} = 2^{\Pi}$ and $\mathrm{B} = \{s \in \mathrm{S} \mid s \models_{cpl} K\}$, i.e. the doxastic alternatives in M_K are exactly those propositional valuations that satisfy K. Then we have for all $\varphi \in \mathcal{L}_0$:

$$
\begin{aligned}
& \varphi \in K \\
\Leftrightarrow\ & K \models_{cpl} \varphi && \text{(since } K \text{ is deductively closed)} \\
\Leftrightarrow\ & \forall s \in \mathrm{S}[s \models_{cpl} K \Rightarrow s \models_{cpl} \varphi] \\
\Leftrightarrow\ & \forall s \in \mathrm{B}[s \models_{cpl} \varphi] \\
\Leftrightarrow\ & \forall s \in \mathrm{B}[\mathrm{M}, s \models \varphi] \\
\Leftrightarrow\ & \mathrm{M} \models \mathbf{B}\varphi \\
\Leftrightarrow\ & \varphi \in \mathrm{B}(\mathrm{M})
\end{aligned}
$$

and hence $K = \mathrm{B}(\mathrm{M})$. □

Note that the full model M_K depends on the set Π of propositional symbols: for different sets of propositional symbols, different K-models result. For a fixed set Π however, the model M_K is unique.

For given K-models, our belief-changing actions behave exactly as the rational changes of belief in the AGM framework.

2.16. THEOREM. *For any belief set K and $\varphi \in \mathcal{L}_0$, $K^+_\varphi = \mathrm{B}(\mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}_K))$, where $\mathtt{r}$ is based on an arbitrary selection function σ for M_K.*

PROOF: Let K be a belief set, let $\varphi \in \mathcal{L}_0$ and let $\mathtt{r}$ be based on an arbitrary selection function σ for M. By Definition 2.9 we have that $K^+_\varphi = \mathrm{Th}(K \cup \{\varphi\})$. We show that $\mathrm{Th}(K \cup \{\varphi\})$ equals $\mathrm{B}(\mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}_K))$ by showing that the two sets are contained in each other. Let $\mathrm{M}' = \langle \mathrm{S}, \mathrm{B}' \rangle$ denote $\mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}_K)$.

'$\subseteq$' By definition of $\mathtt{r}$ for $\mathtt{expand}\ \varphi$ it follows that $\mathrm{B}' = \mathrm{B} \cap [\![\varphi]\!]$, hence φ holds at all states from B', and thus $\varphi \in \mathrm{B}(\mathrm{M}')$. Now let ψ be an element of K. By Proposition 2.15 it follows that $\mathrm{M}_K \models \mathbf{B}\psi$, i.e. ψ holds at all elements of B, and hence, since $\psi \in \mathcal{L}_0$, ψ holds at all elements of B'. Then $\psi \in \mathrm{B}(\mathrm{M}')$, and since ψ is arbitrary, $K \subseteq \mathrm{B}(\mathrm{M}')$. Since $\mathrm{B}(\mathrm{M}')$ is deductively closed by Proposition 2.13, it follows that $\mathrm{Th}(K \cup \{\varphi\}) \subseteq \mathrm{B}(\mathrm{M}')$.

'⊇' Let $\psi \in \mathrm{B}(\mathrm{M}')$, i.e. $\mathrm{M}' \models \mathbf{B}\psi$. Since $\mathrm{M}' = \mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M})$ it follows that $\mathrm{B}' = \mathrm{B} \cap [\![\varphi]\!]$. Hence $\mathrm{M}_K, s \models \psi$ for all $s \in \mathrm{B}$ such that $\mathrm{M}_K, s \models \varphi$, and thus $\mathrm{M}_K, s \models (\varphi \to \psi)$ for all $s \in \mathrm{B}$. Then $\mathrm{M}_K \models \mathbf{B}(\varphi \to \psi)$, and by Proposition 2.15, $\varphi \to \psi \in K$. But then $\psi \in \mathrm{Th}(K \cup \{\varphi\})$. □

2.17. Theorem (Soundness and completeness for contractions).

- *Let* M *be a full model, let* σ *be a selection function for* M, *and let* $K = \mathrm{B}(\mathrm{M})$. *If for any* $\varphi \in \mathcal{L}_0$, K^-_φ *is defined to be* $\mathrm{B}(\mathtt{r}(\mathtt{contract}\ \varphi, \mathrm{M}))$ *for* $\mathtt{r}$ *based on* σ, *then the resulting contraction function is an AGM contraction.*
- *For any belief set* K *and AGM contraction* $\Leftrightarrow$, *some selection function* σ *for* M_K *exists, such that for all* $\varphi \in \mathcal{L}_0$, $K^-_\varphi = \mathrm{B}(\mathtt{r}(\mathtt{contract}\ \varphi, \mathrm{M}_K))$, *where* $\mathtt{r}$ *is based on* σ.

Proof: The soundness part is shown elsewhere [12], and not repeated here. With respect to the completeness part, assume that K is some AGM belief set, and let $\Leftrightarrow$ be an AGM contraction. Let furthermore $\mathrm{M}_K = \langle \mathrm{S}, \mathrm{B} \rangle$ be the K-model. First we need some additional definitions.

2.18. Definition. For a given belief set K, $K \perp \varphi$ is the set of all maximal subsets of K not entailing φ, and $\perp(K)$ is $\bigcup_{\varphi \in K, \not\vdash_{cpl} \varphi} K \perp \varphi$.

Following an observation by Grove [6], we state without proof that for a given belief set K a bijection f exists between the elements of $\mathrm{S} \setminus \mathrm{B}$ in M_K and $\perp(K)$, such that $K' \in K \perp \varphi = \{\psi \in \mathcal{L}_0 \mid \forall s \in \mathrm{B} \cup f(K')[s \models_{cpl} \psi]\}$.

Alchourrón, Gärdenfors and Makinson [1] show that any contraction function $\Leftrightarrow$ can be defined by $K^-_\varphi = \cap\{K' \in K \perp \varphi \mid K' \leq K''$, all $K'' \in K \perp \varphi\}$, where $\leq$ is a *transitive* and *connected* relation on $\perp(K)$ such that $K \perp \varphi$ is *smooth* for all $\varphi \in \mathcal{L}_0$ with $\not\vdash_{cpl} \varphi$[3]. Now let $\leq$ be the relation on $\perp(K)$ that defines $\Leftrightarrow$. Based on $\leq$ we define the relation $\preceq$ on $\mathrm{S} \setminus \mathrm{B}$ by: $s \preceq s' \Leftrightarrow f^{-1}(s) \leq f^{-1}(s')$. It is obvious that $\preceq$ inherits the properties of $\leq$. In particular, $\preceq$ is transitive and connected on $\mathrm{S} \setminus \mathrm{B}$. Furthermore, for all $\psi \in \mathcal{L}_0$ such that $\mathrm{B} \cap [\![\psi]\!] = \emptyset$ it holds that $[\![\psi]\!]$ is smooth whenever this set is non-empty. Now using the relation $\preceq$ on $\mathrm{S} \setminus \mathrm{B}$ we define the function $\varsigma : \mathcal{L}_0 \times \wp(\mathrm{S})$ by:

$$\varsigma(\psi) = \begin{cases} \{s \in \mathrm{S} \mid s \in [\![\psi]\!] \,\&\, s \preceq s' \text{ for all } s' \in [\![\psi]\!]\} & \text{if } \mathrm{B} \cap [\![\psi]\!] = \emptyset \\ \emptyset & \text{otherwise} \end{cases}$$

2.19. Lemma. *The function* ς *as defined above is a selection function for* M_K.

Proof: We show that ς meets the demands imposed on selection functions as given in Definition 2.3. Let $\psi, \vartheta \in \mathcal{L}_0$ be arbitrary.

Σ1. This demand is obviously met since $\varsigma(\psi)$ is either empty or yields the minimal elements of $[\![\psi]\!]$, and in both cases $\varsigma(\psi) \subseteq [\![\psi]\!]$.

[3] Recall that a relation R is transitive on S iff sRs' and $s'Rs''$ implies sRs'' and connected if sRs' or $s'Rs$ for all $s, s', s'' \in \mathrm{S}$. A non-empty set $\mathrm{S}' \subseteq \mathrm{S}$ is smooth iff it contains a minimal element, i.e. some $s \in \mathrm{S}'$ exists such that sRs' for all $s' \in \mathrm{S}'$.

Σ2. The 'if' part is obvious: if $\mathrm{B}\cap[\![\psi]\!] \neq \emptyset$, then $\varsigma(\psi)=\emptyset$ by definition, and if $[\![\psi]\!]=\emptyset$ then no minimal elements of $[\![\psi]\!]$ exist and hence also in this case $\varsigma(\psi)=\emptyset$. The 'only if' part trivially holds whenever $\mathrm{B}\cap[\![\psi]\!]\neq\emptyset$. Hence assume that $\mathrm{B}\cap[\![\psi]\!]=\emptyset$ and $\varsigma(\psi)=\emptyset$. In this case the set of minimal elements of $[\![\psi]\!]$ is empty. But since $[\![\psi]\!]$ is smooth whenever $\mathrm{B}\cap[\![\psi]\!]=\emptyset$ and $[\![\psi]\!]\neq\emptyset$, it follows that $[\![\psi]\!]=\emptyset$, which was to be shown.

Σ3. If $[\![\psi]\!]=[\![\vartheta]\!]$, then $\mathrm{B}\cap[\![\psi]\!]=\emptyset$ iff $\mathrm{B}\cap[\![\vartheta]\!]=\emptyset$. Furthermore, the set of minimal elements of $[\![\psi]\!]$ is equal to that of $[\![\vartheta]\!]$, which suffices to conclude that $\varsigma(\psi)=\varsigma(\vartheta)$, and thus demand Σ3 is met.

Σ4. If $\mathrm{B}\cap[\![\psi\vee\vartheta]\!]\neq\emptyset$ then $\varsigma(\psi\vee\vartheta)=\emptyset$ and hence obviously $\varsigma(\psi\vee\vartheta)\subseteq\varsigma(\psi)\cup\varsigma(\vartheta)$. Hence let $\mathrm{B}\cap[\![\psi\vee\vartheta]\!]=\emptyset$ and let $s\in\varsigma(\psi\vee\vartheta)$, i.e. s is a minimal element of the set $[\![\psi\vee\vartheta]\!]$. Then also $\mathrm{B}\cap[\![\psi]\!]=\emptyset$ and $\mathrm{B}\cap[\![\vartheta]\!]=\emptyset$, and furthermore either $s\in[\![\psi]\!]$ or $s\in[\![\vartheta]\!]$, by definition of $\models$ for disjunctions. Assume that $s\in[\![\psi]\!]$; the case for $s\in[\![\vartheta]\!]$ is completely analogous. Then s is a minimal element of $[\![\psi]\!]$. For assume not: then there is some $s'\in[\![\psi]\!]$ such that $s\not\preceq s'$. But then $s'\in[\![\psi\vee\vartheta]\!]$ and $s\not\preceq s'$ contradicts the minimality of s in $[\![\psi\vee\vartheta]\!]$. Hence s is minimal in $[\![\psi]\!]$, and thus $s\in\varsigma(\psi)$.

Σ5. Assume that $\varsigma(\psi\vee\vartheta)\cap[\![\psi]\!]\neq\emptyset$, i.e. let $s\in\varsigma(\psi\vee\vartheta)\cap[\![\psi]\!]$. Assume towards a contradiction that $\varsigma(\psi)\nsubseteq\varsigma(\psi\vee\vartheta)$, i.e. some $s'\in\varsigma(\psi)$ exists such that $s'\notin\varsigma(\psi\vee\vartheta)$. Now $\mathrm{B}\cap[\![\psi\vee\vartheta]\!]=\emptyset$ and hence also $\mathrm{B}\cap[\![\psi]\!]=\emptyset$. If $s'\preceq s$, then from the fact that s is minimal in $[\![\psi\vee\vartheta]\!]$ it follows by transitivity of $\preceq$ that s' is minimal in $[\![\psi\vee\vartheta]\!]$, which contradicts $s'\notin\varsigma(\psi\vee\vartheta)$. Hence $s'\not\preceq s$. But then s is an element of $[\![\psi]\!]$ such that $s'\not\preceq s$. Hence not for all $t\in[\![\psi]\!]$, $s'\preceq t$, which contradicts $s'\in\varsigma(\psi)$. Hence $\varsigma(\psi)\subseteq\varsigma(\psi\vee\vartheta)$, which suffices to conclude that demand Σ5 is met.

Since ς meets the demands Σ1 to Σ5, we conclude that ς is indeed a selection function. □

Now let $\varphi\in\mathcal{L}_0$ be some arbitrary formula. If either $\varphi\notin K$ or $\vdash_{cpl}\varphi$, then $K^-_\varphi=K$. In this case either $\mathrm{B}\cap[\![\neg\varphi]\!]\neq\emptyset$ or $[\![\neg\varphi]\!]=\emptyset$, which implies that $\varsigma(\neg\varphi)=\emptyset$. Thus $\mathtt{r}(\mathtt{contract}\ \varphi,\mathrm{M}_K)=\mathrm{M}_K$, and $\mathrm{B}(\mathtt{r}(\mathtt{contract}\ \varphi,\mathrm{M}_K)=K=K^-_\varphi$, for $\mathtt{r}$ based on ς. So let $\varphi\in K$ such that $\nvdash_{cpl}\varphi$. Then we have:

$$\begin{aligned}
K^-_\varphi &= \cap\{K'\in K\perp\varphi \mid K'\leq K'' \text{ for all } K''\in K\perp\varphi\}\\
&= \cap\{\{\psi\in\mathcal{L}_0 \mid \forall s\in\mathrm{B}\cup f(K')[s\models_{cpl}\psi]\} \mid f(K')\in[\![\neg\varphi]\!],\\
&\qquad\qquad f(K')\preceq s' \text{ for all } s'\in[\![\neg\varphi]\!]\}\\
&= \cap\{\{\psi\in\mathcal{L}_0 \mid \forall s\in\mathrm{B}\cup f(K')[s\models_{cpl}\psi]\} \mid f(K')\in\varsigma(\neg\varphi)\}\\
&= \{\psi\in\mathcal{L}_0 \mid \forall s\in\mathrm{B}\cup\varsigma(\neg\varphi)[\mathrm{M}_K,s\models\psi]\}\\
&= \{\psi\in\mathcal{L}_0 \mid \mathtt{r}(\mathtt{contract}\ \varphi,\mathrm{M}_K)\models\mathbf{B}\psi\}\\
&= \mathrm{B}(\mathtt{r}(\mathtt{contract}\ \varphi,\mathrm{M}_K))
\end{aligned}$$

□

2.20. Theorem (Soundness and completeness for revisions).

- *Let* M *be a full model, let σ be a selection function for* M, *and let* $K=\mathrm{B}(\mathrm{M})$. *If for any $\varphi\in\mathcal{L}_0$, K^*_φ is defined to be* $\mathrm{B}(\mathtt{r}(\mathtt{revise}\ \varphi,\mathrm{M}))$ *for* $\mathtt{r}$ *based on σ, then the resulting revision function is an AGM revision.*

- *For any belief set* K *and AGM revision* $*$, *some selection function* σ *for* M_K *exists, such that for all* $\varphi \in \mathcal{L}_0$, $K^*_\varphi = \mathrm{B}(\mathtt{r}(\mathtt{revise}\ \varphi, \mathrm{M}_K))$, *where* $\mathtt{r}$ *is based on* σ.

Proof: The soundness part is shown elsewhere [12], and not repeated here. With respect to the completeness part, assume that K is some AGM belief set, and let $*$ be an AGM contraction, i.e. $*$ validates the AGM postulates. Let furthermore $\mathrm{M}_K = \langle \mathrm{S}, \mathrm{B} \rangle$ be the K-model. Let the AGM contraction $\Leftrightarrow$be defined from $*$ through the Harper-identity. In this case $*$ is exactly the AGM revision that results from applying the Levi-identity to the AGM contraction $\Leftrightarrow$[5]. From Theorem 2.20 it follows that some selection function σ for M_K exists, such that for all $\varphi \in \mathcal{L}_0$, $K^-_\varphi = \mathrm{B}(\mathtt{r}(\mathtt{contract}\ \varphi, \mathrm{M}_K))$ for $\mathtt{r}$ based on σ. Now let σ be such a selection function. We claim that for all $\varphi \in \mathcal{L}_0$, $K^*_\varphi = \mathtt{r}(\mathtt{revise}\ \varphi, \mathrm{M}_K)$ for $\mathtt{r}$ based on σ. To see this, let M' be $\mathtt{r}(\mathtt{contract}\ \neg\varphi, \mathrm{M}_K)$, for some $\varphi \in \mathcal{L}_0$ arbitrary. Then $\mathtt{r}(\mathtt{revise}\ \varphi, \mathrm{M}_K) = \mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}')$. From Theorem 2.16 it follows that $\mathrm{B}(\mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}')) = (\mathrm{B}(\mathrm{M}'))^+_\varphi$. Now $\mathrm{B}(\mathrm{M}')$ equals $K^-_{\neg\varphi}$, and hence $\mathrm{B}(\mathtt{r}(\mathtt{revise}\ \varphi, \mathrm{M}_K)) = \mathrm{B}(\mathtt{r}(\mathtt{expand}\ \varphi, \mathrm{M}')) = (K^-_{\neg\varphi})^+_\varphi = K^*_\varphi$, which was to be shown. $\Box$

Theorems 2.16 to 2.20 indicate that our system may be viewed as an agent-oriented, *modal* implementation of AGM belief changes, in which dynamic, doxastic and epistemic logic are combined. The expressiveness of our framework allows for a concise representation of the AGM postulates as validities with respect to $\mathcal{M}$. The proofs of the following propositions are not given here; they can be found elsewhere [12].

2.21. Proposition. *For all* $\varphi, \psi \in \mathcal{L}_0$ *we have:*

- $\models [\mathtt{expand}\ \varphi]\mathbf{B}\psi \leftrightarrow \mathbf{B}(\varphi \to \psi)$

2.22. Proposition. *For all* φ, ψ *and* $\vartheta \in \mathcal{L}_0$ *we have:*

1. $\models ([\mathtt{contract}\ \varphi]\mathbf{B}(\psi \to \vartheta) \wedge [\mathtt{contract}\ \varphi]\mathbf{B}\psi) \to [\mathtt{contract}\ \varphi]\mathbf{B}\vartheta$
2. $\models [\mathtt{contract}\ \varphi]\mathbf{B}\vartheta \to \mathbf{B}\vartheta$
3. $\models \neg\mathbf{B}\varphi \to ([\mathtt{contract}\ \varphi]\mathbf{B}\vartheta \leftrightarrow \mathbf{B}\vartheta)$
4. $\models \neg\mathbf{K}\varphi \to [\mathtt{contract}\ \varphi]\neg\mathbf{B}\varphi$
5. $\models \mathbf{B}\varphi \to (\mathbf{B}\vartheta \to [\mathtt{contract}\ \varphi; \mathtt{expand}\ \varphi]\mathbf{B}\vartheta)$
6. $\models \mathbf{K}(\varphi \leftrightarrow \psi) \to ([\mathtt{contract}\ \varphi]\mathbf{B}\vartheta \leftrightarrow [\mathtt{contract}\ \psi]\mathbf{B}\vartheta)$
7. $\models ([\mathtt{contract}\ \varphi]\mathbf{B}\vartheta \wedge [\mathtt{contract}\ \psi]\mathbf{B}\vartheta) \to$
   $[\mathtt{contract}\ (\varphi \wedge \psi)]\mathbf{B}\vartheta$
8. $\models [\mathtt{contract}\ (\varphi \wedge \psi)]\neg\mathbf{B}\varphi \to$
   $([\mathtt{contract}\ (\varphi \wedge \psi)]\mathbf{B}\vartheta \to [\mathtt{contract}\ \varphi]\mathbf{B}\vartheta)$

2.23. Proposition. *For all* $\varphi, \psi, \vartheta \in \mathcal{L}_0$ *we have:*

1. $\models ([\mathtt{revise}\ \varphi]\mathbf{B}(\psi \to \vartheta) \wedge [\mathtt{revise}\ \varphi]\mathbf{B}\psi) \to [\mathtt{revise}\ \varphi]\mathbf{B}\vartheta$
2. $\models [\mathtt{revise}\ \varphi]\mathbf{B}\varphi$
3. $\models [\mathtt{revise}\ \varphi]\mathbf{B}\vartheta \to [\mathtt{expand}\ \varphi]\mathbf{B}\vartheta$
4. $\models \neg\mathbf{B}\neg\varphi \to ([\mathtt{expand}\ \varphi]\mathbf{B}\vartheta \to [\mathtt{revise}\ \varphi]\mathbf{B}\vartheta)$
5. $\models [\mathtt{revise}\ \varphi]\mathbf{B}\bot \leftrightarrow \mathbf{K}\neg\varphi$

6. $\models \mathbf{K}(\varphi \leftrightarrow \psi) \rightarrow ([\mathtt{revise}\ \varphi]\mathbf{B}\vartheta \leftrightarrow [\mathtt{revise}\ \psi]\mathbf{B}\vartheta)$
7. $\models [\mathtt{revise}\ (\varphi \wedge \psi)]\mathbf{B}\vartheta \rightarrow [\mathtt{revise}\ \varphi; \mathtt{expand}\ \psi]\mathbf{B}\vartheta$
8. $\models \neg[\mathtt{revise}\ \varphi]\mathbf{B}\neg\psi \rightarrow$
   $([\mathtt{revise}\ \varphi; \mathtt{expand}\ \psi]\mathbf{B}\vartheta \rightarrow [\mathtt{revise}\ (\varphi \wedge \psi)]\mathbf{B}\vartheta)$

An important observation concerning the AGM postulates is the almost complete absence of postulates guiding iterated changes of belief. The sequences of belief changes that do occur are all such that the final change in the sequence is the most straightforward kind of belief change, viz. an expansion. The investigation of more general iterated belief changes is the subject of the following sections.

3. Actions that keep on changing your mind

The AGM framework — and hence also the dynamic framework presented in the previous section — is a reasonable one when modelling single-step belief changes. However, for non-trivial sequences of belief changes, the AGM framework is on the one hand not restrictive enough, and on the other hand too restrictive. The problem of the AGM framework being not restrictive enough is tackled by a number of proposed additional postulates [3, 4, 14]. With regard to the AGM framework being too restrictive we would like to focus on what we call the *postulate of rehabilitation*:

$$\varphi \notin K \Rightarrow (K^+_\varphi)^-_\varphi = K$$

The postulate of rehabilitation states that an expansion with some formula can always be undone by a contraction with the same formula. In our opinion this postulate is a reasonable and acceptable one, especially from an agent-oriented point of view. Consider for instance the situation of an agent that adopts some formula by default (cf. [13]). Upon recognising a possible inconsistency caused by this adoption, the agent could decide to retract the formula again, thereby expecting to end up with its original belief set. Despite its apparent acceptability this postulate is explicitly not present in the AGM framework. As Gärdenfors states it:

> 'One may wonder what happens if K is first expanded by A and then contracted with respect to A. Do we always get K back? This cannot be true in general, [...] because if $\neg A \in K$, then $K^+_A = K_\perp$, and if H is another belief set such that $\neg A \in H$, then also $H^+_A = K_\perp$, and hence $(K^+_A)^-_A = (H^+_A)^-_A$. It follows that $(K^+_A)^-_A$ cannot always be identical with K.' ([5], pp. 62–63)

The crux of the argument given by Gärdenfors is the conclusion of $(K^+_A)^-_A = (H^+_A)^-_A$ from $K^+_A = H^+_A$. The formalisation that we present, in which the postulate of rehabilitation is validated, attacks the argument of Gärdenfors exactly on this conclusion. We allow the results of belief changes to depend not only on the belief set under consideration, but also on the way this belief set originated. And for two (seemingly) identical belief sets K^+_A and H^+_A it may be the case that their origin is so much different that retracting A from these sets leads to different results.

3.1. Belief systems

To formalise iterated changes of belief in our framework, we introduce the notion of *belief systems*. Intuitively, a belief system consists of a set of beliefs of the agent and a method describing how the belief set changes under the execution of belief-changing actions. As such, these belief systems are highly analogous to the knowledge systems of Williams [16], which consist of a knowledge set and a preference relation which determines how changes of the knowledge set work out. Formally, belief systems are triples $\langle \mathrm{S}, \mathrm{B}, \sigma \rangle$, consisting of a set S of states, a set B of doxastic alternatives and a function σ which is a selection function for $\langle \mathrm{S}, \mathrm{B} \rangle$ in the sense of Definition 2.3. The function σ provides a method to modify the belief set determined by the set B of doxastic alternatives upon execution of belief-changing actions[4]. The combination of B and σ represents a more structured kind of doxastic state of the agent as compared to plain belief sets. Execution of belief-changing actions affects the whole of the doxastic state of the agent, i.e. in general both the set B of doxastic alternatives and the function σ are modified. The result of executing a belief-changing action is a new belief system, of which the set of doxastic alternatives is modified as in §2, and whose selection function provides a method to change the (new) set of beliefs determined by the modified set of doxastic alternatives. As such, our approach is a typical *dynamic* one, according to the classification proposed by Freund & Lehmann [4]. In dynamic approaches to iterated belief change, as for instance proposed by Boutilier [2] in the setting of a conditional logic, and by Williams [16] in a probabilistic setting, the agent starts with a belief set K and a method of changing K. A change of belief does not only modify K but also affects the method of changing the belief set such that the new method is applicable to the new belief set. In static approaches the effect of a change of a belief set K with some formula φ depends on K and φ only, and is not liable to change. Although a static approach to belief change may seem to be more in line with the AGM framework, the major advantage of dynamic approaches is their flexibility. Freund & Lehmann [4, 10] show that even a fairly weak set of postulates for iterated belief revision inevitably trivialises the revision operator in the static case, whereas from a dynamic point of view iterated changes of belief may satisfy all kind of intuitively acceptable properties without trivialising (see §4).

In the following definitions we propose a formalisation of belief systems, and of belief-changing actions, execution of which modifies belief systems rather than belief sets.

3.1. Definition. The class $\mathcal{M}^\star$ consists of all tuples $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle$ where S and B are as in Definition 2.2 and σ is a selection function for $\langle \mathrm{S}, \mathrm{B} \rangle$ in the sense of Definition 2.3.

[4]Note that belief systems are both very expressive and very flexible. Not only may one formalise belief sets through the set B of doxastic alternatives and a method to change this belief set through the function σ, but by restricting the set S of states it is furthermore possible to formalise *integrity constraints*, which are fixed, unassailable demands that should be met throughout all changes of belief.

3.2. Definition. The binary relation $\models_\star$ between a formula φ and a pair M, s consisting of a model $\mathrm{M} \in \mathcal{M}^\star$ and a state s in M, is defined almost exactly as the relation $\models$ in Definition 2.2, the only difference being the replacement of the function $\mathtt{r}$ by $\mathtt{r}^\star$, which is for sequential compositions defined as $\mathtt{r}$ and for expansions and contractions as follows:

$$\begin{array}{ll}
\mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathrm{M}) & = \langle \mathrm{S}, \mathrm{B}', \sigma' \rangle \text{ where } \mathrm{B}' = \mathrm{B} \setminus [\![\neg\varphi]\!] \text{ and} \\
\quad \sigma'(\vartheta) = \sigma(\vartheta) & \text{if } \mathrm{B} \cap [\![\vartheta]\!] = \emptyset \text{ or } \mathrm{B}' \cap [\![\vartheta]\!] \neq \emptyset \\
\quad \sigma'(\vartheta) = \mathrm{B} \cap [\![\vartheta]\!] & \text{otherwise} \\
\mathtt{r}^\star(\mathtt{contract}\ \varphi, \mathrm{M}) & = \langle \mathrm{S}, \mathrm{B}', \sigma' \rangle \text{ where } \mathrm{B}' = \mathrm{B} \cup \sigma(\neg\varphi) \text{ and} \\
\quad \sigma'(\vartheta) = \sigma(\vartheta) & \text{if } \mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset \text{ or } \mathrm{B}' \cap [\![\vartheta]\!] = \emptyset \\
\quad \sigma'(\vartheta) = \emptyset & \text{otherwise}
\end{array}$$

Validity on and satisfiability in a model $\mathrm{M} \in \mathcal{M}^\star$ and the class $\mathcal{M}^\star$ is defined as for $\mathcal{M}$.

3.3. Remark. Note that in the class $\mathcal{M}^\star$ we no longer have a unique K-model for a given belief set K: although the set B of doxastic alternatives is uniquely determined by K, a whole range of selection functions may be used to constitute a K-model in $\mathcal{M}^\star$.

3.4. Proposition. *For all* $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle \in \mathcal{M}^\star$ *and for all* $\varphi \in \mathcal{L}_0$ *we have:*

$$\begin{array}{ll}
\mathtt{r}^\star(\mathtt{revise}\ \varphi, \mathrm{M}) & = \langle \mathrm{S}, \mathrm{B}', \sigma' \rangle \textit{ where } \mathrm{B}' = (\mathrm{B} \cup \sigma(\varphi)) \cap [\![\varphi]\!] \textit{ and} \\
\quad \sigma'(\vartheta) = \mathrm{B} \cap [\![\vartheta]\!] & \textit{if } \mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset \textit{ and } \mathrm{B}' \cap [\![\vartheta]\!] = \emptyset \\
\quad \sigma'(\vartheta) = \sigma(\vartheta) & \textit{if } \mathrm{B} \cap [\![\vartheta]\!] = \emptyset \textit{ and } \mathrm{B}' \cap [\![\vartheta]\!] = \emptyset \\
\quad \sigma'(\vartheta) = \emptyset & \textit{if } \mathrm{B}' \cap [\![\vartheta]\!] \neq \emptyset
\end{array}$$

Proof: Let $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle \in \mathcal{M}^\star$ and $\varphi \in \mathcal{L}_0$ be arbitrary. From Definition 3.2 it follows that

$$\begin{array}{l}
\quad \mathtt{r}^\star(\mathtt{revise}\ \varphi, \mathrm{M}) \\
= \mathtt{r}^\star(\mathtt{contract}\ \neg\varphi; \mathtt{expand}\ \varphi, \mathrm{M}) \\
= \mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathtt{r}^\star(\mathtt{contract}\ \neg\varphi, \mathrm{M})) \\
= \mathtt{r}^\star(\mathtt{expand}\ \varphi, \langle \mathrm{S}, \mathrm{B} \cup \sigma(\varphi), \sigma' \rangle) \\
= \langle \mathrm{S}, (\mathrm{B} \cup \sigma(\varphi)) \cap [\![\varphi]\!], \sigma'' \rangle
\end{array}$$

which leaves to show that σ'' is adequately characterised. Let $\vartheta \in \mathcal{L}_0$ be arbitrary. For reasons of convenience we define:

- $\mathrm{M}_1 = \langle \mathrm{S}, \mathrm{B}_1, \sigma_1 \rangle = \mathtt{r}^\star(\mathtt{contract}\ \neg\varphi, \mathrm{M})$
- $\mathrm{M}_2 = \langle \mathrm{S}, \mathrm{B}_2, \sigma_2 \rangle = \mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathrm{M}_1) = \mathtt{r}^\star(\mathtt{revise}\ \varphi, \mathrm{M})$

We distinguish three cases:

1. $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$ and $\mathrm{B}_2 \cap [\![\vartheta]\!] = \emptyset$. In this case $\mathrm{B}_1 \cap [\![\vartheta]\!] \neq \emptyset$, and thus $\sigma_2(\vartheta) = \mathrm{B}_1 \cap [\![\vartheta]\!] = (\mathrm{B} \cup \sigma(\varphi)) \cap [\![\vartheta]\!]$. Now since $\mathrm{B}_2 \cap [\![\vartheta]\!] = \emptyset$ it follows that $\sigma(\varphi) \cap [\![\vartheta]\!] = \emptyset$. Hence $\sigma_2(\vartheta) = \mathrm{B} \cap [\![\vartheta]\!]$, which was to be shown.
2. $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$ and $\mathrm{B}_2 \cap [\![\vartheta]\!] = \emptyset$. In this case also $\mathrm{B}_1 \cap [\![\vartheta]\!] = \emptyset$. Hence $\sigma_2(\vartheta) = \sigma_1(\vartheta)$ and $\sigma_1(\vartheta) = \sigma(\vartheta)$. Thus $\sigma_2(\vartheta) = \sigma(\vartheta)$ which was to be shown.

3. $B_2 \cap [\![\vartheta]\!] \neq \emptyset$. In this case $\sigma_2(\vartheta) = \sigma_1(\vartheta)$. Also $B_1 \cap [\![\vartheta]\!] \neq \emptyset$. Now if $B \cap [\![\vartheta]\!] \neq \emptyset$, then $\sigma_1(\vartheta) = \sigma(\vartheta)$, and since σ is a selection function, $\sigma(\vartheta) = \emptyset$. If $B \cap [\![\vartheta]\!] = \emptyset$, $\sigma_1(\vartheta) = \emptyset$, which suffices to conclude this case. □

3.5. Remark. Since the function $\mathtt{r}^\star$ is for a given model $\langle S, B, \sigma\rangle$ based on the selection function σ, and is, as far as the set of doxastic alternatives is concerned, defined as the function $\mathtt{r}$, the soundness and completeness results presented in §2.1 transfer directly to the system presented in this section. That is, the `expand`, `contract` and `revise` actions when defined as in 3.2 satisfy the respective AGM characterisations of expansions, contractions and revisions. Moreover, for every belief set K and every AGM contraction $\Leftrightarrow$, some model $M \in \mathcal{M}^\star$ exists such that for all $\varphi \in \mathcal{L}_0$, $K^-_\varphi = B(\mathtt{r}^\star(\mathtt{contract}\ \varphi, M))$, and the same is true for AGM revisions. In this aspect the system defined above is a *conservative extension* of the one presented in the previous section.

The general idea underlying the definition of $\mathtt{r}^\star$ for both `expand` φ and `contract` φ is that the agent possesses a sense of *historical awareness* which makes it prefer the worlds that it most recently considered possible doxastic alternatives. One can imagine the agent being aware of the presence of these possibilities, which make it easier to reconsider these as compared to other worlds that it never before considered possible: unknown, unloved. This historical awareness of the agent is intuitively related to the postulate of rehabilitation: if the agent should have the possibility to undo undesired changes of belief, the previous constellation of the agent's belief should somehow be recorded. The modification of the selection function as proposed in the definition of $\mathtt{r}^\star$ above, takes care of this recording.

When performing an expansion with some formula φ, the selection function in the resulting model equals the original selection function for those formulae of which the negation was already believed in the original model and hence is also believed in the model resulting from the expansion, and the same holds for formulae whose negation is not believed *a posteriori*. In the latter case the negation of the formulae is also disbelieved *a priori*, which implies that the original selection function — and hence also the resulting selection function — yields an empty set of states for these formulae. For a formula whose negation is not believed *a priori*, but is believed as the result of the expansion, the new selection function is defined to comprise exactly those doxastic alternatives of the original model that supported the formula, these being exactly the states that caused the *a priori* disbelief in the negation of the formula.

When performing a contraction with some formula φ, the resulting selection is identical to the original one for formulae whose whose negation is either not believed *a priori*, or still believed *a posteriori*. For these formulae the original selection function either yields the empty set, in which case the resulting selection function should do the same by demand $\Sigma 2$ for selection functions, or picked a set of worlds that is disjoint with the new, larger, set of doxastic alternatives and therefore needs not to be modified. In the case that the negation of a formula becomes disbelieved as the result of performing a contraction, the new selection function yields an empty set of worlds when applied to this formula.

Following the terminology of Williams [16], we refer to the process of changing a belief system as a *transmutation*. The transmutations caused by the execution of a belief-changing action preserve well-definedness of belief systems.

3.6. PROPOSITION. *If* $\mathrm{M} \in \mathcal{M}^\star$ *is a well-defined Kripke model then*

- $\mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathrm{M})$ *is a well-defined Kripke model*
- $\mathtt{r}^\star(\mathtt{contract}\ \varphi, \mathrm{M})$ *is a well-defined Kripke model*

for all $\varphi \in \mathcal{L}_0$.

PROOF: Let $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle$ be a well-defined Kripke model. We show the case for the `expand` action; the case for the `contract` action goes through analogously and is left to the reader. So let $\mathrm{M}' = \langle \mathrm{S}, \mathrm{B}', \sigma' \rangle$ be $\mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathrm{M})$, for some $\varphi \in \mathcal{L}_0$. Since $\mathrm{B}' = \mathrm{B} \setminus [\![\neg\varphi]\!]$, it is obvious that $\mathrm{B}' \subseteq \mathrm{B} \subseteq \mathrm{S}$. This leaves only to show that σ' is a selection function. First observe that σ' is indeed a total function on $\mathcal{L}_0$, which leaves to show that σ' meets the demands given in Definition 2.3. Let $\vartheta, \rho \in \mathcal{L}_0$ be arbitrary.

Σ1. If $\sigma'(\vartheta) = \sigma(\vartheta)$, then since σ is a selection function, it follows that $\sigma'(\vartheta) \subseteq [\![\vartheta]\!]$. If $\sigma'(\vartheta)$ is $\mathrm{B} \cap [\![\vartheta]\!]$ it follows directly that $\sigma'(\vartheta) \subseteq [\![\vartheta]\!]$.

Σ2. For the 'only if' part assume that $\sigma'(\vartheta) = \emptyset$ and $\mathrm{B}' \cap [\![\vartheta]\!] = \emptyset$. Then to show that $[\![\vartheta]\!] = \emptyset$. If $\mathrm{B}' \cap [\![\vartheta]\!] = \emptyset$, then from $\sigma'(\vartheta) = \emptyset$ if follows that also $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$. For otherwise $\sigma'(\vartheta) = \mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$. Hence in this case $\sigma'(\vartheta) = \sigma(\vartheta)$. Thus $\sigma(\vartheta) = \emptyset$ and since σ is a selection function it follows, from $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$, that $[\![\vartheta]\!] = \emptyset$, which was to be shown. For the 'if' part assume that $\mathrm{B}' \cap [\![\vartheta]\!] \neq \emptyset$ or $[\![\vartheta]\!] = \emptyset$. To show that $\sigma'(\vartheta) = \emptyset$. Now if $\mathrm{B}' \cap [\![\vartheta]\!] \neq \emptyset$, then since $\mathrm{B}' \subseteq \mathrm{B}$ also $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$. Hence $\sigma'(\vartheta) = \sigma(\vartheta)$, and by demand Σ2 for σ it follows that $\sigma(\vartheta) = \emptyset$. If $[\![\vartheta]\!] = \emptyset$, also $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$. Hence $\sigma'(\vartheta) = \sigma(\vartheta)$, which by demand Σ2 for σ is equal to $\emptyset$.

Σ3. Assume $[\![\vartheta]\!] = [\![\rho]\!]$. Then $\mathrm{B} \cap [\![\vartheta]\!] = \mathrm{B} \cap [\![\rho]\!]$, $\mathrm{B}' \cap [\![\vartheta]\!] = \mathrm{B}' \cap [\![\rho]\!]$ and $\sigma(\vartheta) = \sigma(\rho)$, which suffices to conclude that $\sigma'(\vartheta) = \sigma'(\rho)$.

Σ4. If $\mathrm{B} \cap [\![\vartheta \vee \rho]\!] = \emptyset$, then also $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$ and $\mathrm{B} \cap [\![\rho]\!] = \emptyset$. Hence in this case $\sigma'(\vartheta \vee \rho) = \sigma(\vartheta \vee \rho)$, $\sigma'(\vartheta) = \sigma(\vartheta)$ and $\sigma'(\rho) = \sigma(\rho)$. Thus $\sigma'(\vartheta \vee \rho) = \sigma(\vartheta \vee \rho) \subseteq \sigma(\vartheta) \cup \sigma(\rho) = \sigma'(\vartheta) \cup \sigma'(\rho)$, which implies that Σ4 is validated. If $\mathrm{B}' \cap [\![\vartheta \vee \rho]\!] \neq \emptyset$, then $\sigma'(\vartheta \vee \rho) = \sigma(\vartheta \vee \rho)$, and since $\mathrm{B}' \subseteq \mathrm{B}$ also $\mathrm{B} \cap [\![\vartheta \vee \rho]\!] \neq \emptyset$, and by Σ2, $\sigma(\vartheta \vee \rho) = \emptyset$. Hence $\sigma'(\vartheta \vee \rho) = \emptyset$ and trivially $\sigma'(\vartheta \vee \rho) \subseteq \sigma'(\vartheta) \cup \sigma'(\rho)$. If $\mathrm{B} \cap [\![\vartheta \vee \rho]\!] \neq \emptyset$ and $\mathrm{B}' \cap [\![\vartheta \vee \rho]\!] = \emptyset$, then also $\mathrm{B}' \cap [\![\vartheta]\!] = \emptyset$ and $\mathrm{B}' \cap [\![\rho]\!] = \emptyset$. Furthermore, either $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$ or $\mathrm{B} \cap [\![\rho]\!] \neq \emptyset$. If both are true, then $\sigma'(\vartheta) = \mathrm{B} \cap [\![\vartheta]\!]$, $\sigma'(\rho) = \mathrm{B} \cap [\![\rho]\!]$, and since $\sigma'(\vartheta \vee \rho) = \mathrm{B} \cap [\![\vartheta \vee \rho]\!] = (\mathrm{B} \cap [\![\vartheta]\!]) \cup (\mathrm{B} \cap [\![\rho]\!])$, demand Σ4 is met. If $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$, then $\mathrm{B} \cap [\![\vartheta \vee \rho]\!] = \mathrm{B} \cap [\![\rho]\!]$, and hence $\sigma'(\vartheta \vee \rho) = \sigma'(\rho)$, in which case Σ4 is validated. The case where $\mathrm{B} \cap [\![\rho]\!] = \emptyset$ is analogous.

Σ5. Assume $\sigma'(\vartheta \vee \rho) \cap [\![\vartheta]\!] \neq \emptyset$. If $\mathrm{B} \cap [\![\vartheta \vee \rho]\!] = \emptyset$, then also $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$. Hence $\sigma'(\vartheta \vee \rho) = \sigma(\vartheta \vee \rho)$ and $\sigma'(\vartheta) = \sigma(\vartheta)$. Then $\sigma(\vartheta \vee \rho) \cap [\![\vartheta]\!] \neq \emptyset$, and since σ is a selection function, $\sigma'(\vartheta) = \sigma(\vartheta) \subseteq \sigma(\vartheta \vee \rho) = \sigma'(\vartheta \vee \rho)$. If $\mathrm{B}' \cap [\![\vartheta \vee \rho]\!] \neq \emptyset$, then $\sigma'(\vartheta \vee \rho) = \sigma(\vartheta \vee \rho) = \emptyset$, and hence this case is not applicable. If $\mathrm{B} \cap [\![\vartheta \vee \rho]\!] \neq \emptyset$ and $\mathrm{B}' \cap [\![\vartheta \vee \rho]\!] = \emptyset$, then

$\sigma'(\vartheta \vee \rho) = \mathrm{B} \cap [\![\vartheta \vee \rho]\!] = (\mathrm{B} \cap [\![\vartheta]\!]) \cup (\mathrm{B} \cap [\![\rho]\!])$. Since $\sigma'(\vartheta \vee \rho) \cap [\![\vartheta]\!] \neq \emptyset$, it follows that $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$, and since $\mathrm{B}' \cap [\![\vartheta \vee \rho]\!] = \emptyset$ and hence $\mathrm{B}' \cap [\![\vartheta]\!] = \emptyset$, it follows that $\sigma'(\vartheta) = \mathrm{B} \cap [\![\vartheta]\!]$. Thus $\sigma'(\vartheta) \subseteq \sigma'(\vartheta \vee \rho)$ and hence demand $\Sigma 5$ is met. □

3.7. Corollary. *If* $\mathrm{M} \in \mathcal{M}^\star$ *is a well-defined Kripke model then, for all* $\varphi \in \mathcal{L}_0$, $\mathtt{r}^\star(\mathtt{revise}\ \varphi, \mathrm{M})$ *is a well-defined Kripke model.*

When performing an expansion with a formula that is already believed, or a contraction with a disbelieved formula, the generalised law of inertia states that the underlying belief system should not change, i.e. not only should the belief set of the agent be unaffected, but performance of such an intuitively superfluous action should also not influence the way that future changes of belief work out. For the system defined above, it is indeed the case that these degenerated belief-changing actions correspond to void transmutations.

3.8. Proposition. *For all* $\mathrm{M} \in \mathcal{M}^\star$ *and all* $\varphi \in \mathcal{L}_0$ *it holds that:*

- $\mathrm{M} \models_\star \mathbf{B}\varphi \Rightarrow \mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathrm{M}) = \mathrm{M}$
- $\mathrm{M} \models_\star \neg\mathbf{B}\varphi \Rightarrow \mathtt{r}^\star(\mathtt{contract}\ \varphi, \mathrm{M}) = \mathrm{M}$

Proof: We show the case for the `contract` action; the other case is shown in a similar way. So let $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle$ be some model from $\mathcal{M}^\star$, and assume that $\mathrm{M} \models_\star \neg\mathbf{B}\varphi$. Let $\mathrm{M}' = \langle \mathrm{S}, \mathrm{B}', \sigma' \rangle$ be $\mathtt{r}^\star(\mathtt{contract}\ \varphi, \mathrm{M})$. In order to prove that $\mathrm{M}' = \mathrm{M}$ we show that $\mathrm{B}' = \mathrm{B}$ and $\sigma' = \sigma$. From $\mathrm{M} \models_\star \neg\mathbf{B}\varphi$ it follows that $\mathrm{B} \cap [\![\neg\varphi]\!] \neq \emptyset$, and hence, by demand $\Sigma 2$ for selection functions, $\sigma(\neg\varphi) = \emptyset$. Then $\mathrm{B}' = \mathrm{B} \cup \sigma(\neg\varphi) = \mathrm{B}$. This leaves to show that $\sigma' = \sigma$. Let $\vartheta \in \mathcal{L}_0$ be arbitrary. In the case that either $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$ or $\mathrm{B}' \cap [\![\vartheta]\!] \neq \emptyset$, we have direct that $\sigma'(\vartheta) = \sigma(\vartheta)$. But since $\mathrm{B}' = \mathrm{B}$ the case where $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$ and $\mathrm{B}' \cap [\![\vartheta]\!] = \emptyset$ is impossible. Hence $\sigma' = \sigma$, and thus $\mathrm{M}' = \mathrm{M}$. □

3.9. Corollary. *For all* $\mathrm{M} \in \mathcal{M}^\star$ *and all* $\varphi \in \mathcal{L}_0$ *it holds that:*

- $\mathrm{M} \models_\star \neg\mathbf{B}\neg\varphi \Rightarrow \mathtt{r}^\star(\mathtt{revise}\ \varphi, \mathrm{M}) = \mathtt{r}^\star(\mathtt{expand}\ \varphi, \mathrm{M})$
- $\mathrm{M} \models_\star \mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\varphi \Rightarrow \mathtt{r}^\star(\mathtt{revise}\ \varphi, \mathrm{M}) = \mathrm{M}$

4. Postulates for iterated belief change

The AGM postulates, and in fact any postulate describing single-step belief changes, satisfy the general pattern that they contain some occurrences of formulae $[\alpha]\mathbf{B}\varphi$ where $\alpha \in Ac$ and $\varphi \in \mathcal{L}_0$, when phrased in terms of our framework; exactly these formulae describe the *changes to belief sets* that follow execution of belief-changing actions. As soon as one considers *belief systems* instead of just belief sets, other postulates than those constraining belief sets may be worth looking at. For in this case it is not only interesting how the agent's belief set changes as the result of the execution of a belief-changing action, but also how the selection function is modified. While changes of belief sets are described by formulae $[\alpha]\mathbf{B}\varphi$ with $\varphi \in \mathcal{L}_0$, changes of belief systems correspond to formulae of the form $[\alpha]\rho$, where $\alpha \in Ac$ and $\rho \in \mathcal{L}$. The formula ρ is an arbitrary formula, and may in particular contain references to future changes of belief. The following example illustrates the differences in expressive power of the two kinds of formulae.

4.1. EXAMPLE. Let $\mathcal{L} = \mathcal{L}(\{p, q\})$ and consider the full model $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle$, where $\mathrm{B} = \{\{q\}, \emptyset\}$ and $\sigma(\vartheta) = \emptyset$ whenever $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$ and $\mathrm{S} \cap [\![\vartheta]\!]$ otherwise. The function σ is the so called *All-is-Good* function which was previously shown to be a selection function [12]. Now consider revisions with the formulae p and $p \wedge \neg q$ respectively. In the model M the formula $[\texttt{revise}\ p; \texttt{revise}\ p \wedge \neg q]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise}\ p \wedge \neg q]\mathbf{B}\vartheta$ is valid for all $\vartheta \in \mathcal{L}_0$. Hence execution of the actions $\texttt{revise}\ p; \texttt{revise}\ p \wedge \neg q$ and $\texttt{revise}\ p \wedge \neg q$ result in identical changes to the agent's belief set. However, it is not the case that these actions cause the same transmutation of the belief system M. For it is true that $\mathrm{M} \models_\star [\texttt{revise}\ p; \texttt{revise}\ p \wedge \neg q][\texttt{revise}\ q](\mathbf{B}(p \wedge q) \wedge \neg \mathbf{B}\bot)$ whereas $\mathrm{M} \models_\star [\texttt{revise}\ p \wedge \neg q][\texttt{revise}\ q](\mathbf{B}(\neg p \wedge q) \wedge \neg \mathbf{B}\bot)$.

Example 4.1 shows that one may distinguish between the *direct effects* and the *side-effects* that follow execution of a belief-changing action. The direct effects of an action cause a change in the set of beliefs of the agent whereas the side-effects affect the way future changes of belief work out. In terms of models, direct effects reside in the modification of the set of doxastic alternatives and side-effects are visible in the adaption of the selection function. In terms of postulates, direct effects affect formulae $\mathbf{B}\varphi$ with $\varphi \in \mathcal{L}_0$ and side-effects affect general formulae ρ[5]. Now two actions with the same direct effects may have different side-effects. Hence even though they change the belief set of the agent in the same way, they constrain future changes of beliefs in different ways.

Proposition 3.8 and Corollary 3.9 given in the previous section may be interpreted as postulates for belief systems.

4.2. PROPOSITION. *For all* $\varphi \in \mathcal{L}_0$, $\rho \in \mathcal{L}$ *we have:*

- $\models_\star \mathbf{B}\varphi \rightarrow ([\texttt{expand}\ \varphi]\rho \leftrightarrow \rho)$
- $\models_\star \neg\mathbf{B}\varphi \rightarrow ([\texttt{contract}\ \varphi]\rho \leftrightarrow \rho)$
- $\models_\star \neg\mathbf{B}\neg\varphi \rightarrow ([\texttt{revise}\ \varphi]\rho \leftrightarrow [\texttt{expand}\ \varphi]\rho)$
- $\models_\star \neg\mathbf{B}\neg\varphi \wedge \mathbf{B}\varphi \rightarrow ([\texttt{revise}\ \varphi]\rho \leftrightarrow \rho)$

Proposition 4.2 is in fact weaker than the corresponding proposition and corollary of §3, in the sense that it is not stated that two actions result in the same *model* upon execution, but in two models that satisfy exactly the same set of formulae. However, not only lies our main interest in the formulae that are satisfied by a model, and not directly in the shape of the model, but it is furthermore the case that proofs of propositions like 4.2 will in general consist

[5]The difference between direct effects and side-effects may also be used to characterise static and dynamic approaches to iterated belief change. Static approaches are those where changing the belief set does not have any side-effects, whereas in dynamic approaches side-effects may (and do) occur.

of proving the two resulting models to be equal, just as the proof of postulates on the belief set level consists of proving the equality of two sets of doxastic alternatives. This is also the reason why postulates at the general formula level deal with equivalences only, and do in general not contain implications. For in order to prove equivalence of two formulae $[\alpha_1]\rho$ and $[\alpha_2]\rho$ for all ρ, it suffices to prove that the models resulting from execution of α_1 and α_2 are identical, whereas it is not clear how it should be proved that $[\alpha_1]\rho$ implies $[\alpha_2]\rho$ for all $\rho \in \mathcal{L}$.

As argued at the beginning of §3, the semantics for iterated belief change should validate the postulate of rehabilitation. It turns out that the postulate is indeed valid, and even in the strongest possible sense, i.e. it is a property of *belief systems* and not just of *belief sets*.

4.3. Proposition. *For all $\varphi \in \mathcal{L}_0$ and $\rho \in \mathcal{L}$ we have:*

- $\models_\star \neg\mathbf{B}\varphi \to ([\texttt{expand } \varphi; \texttt{contract } \varphi]\rho \leftrightarrow \rho)$

Proof: Let M be a model such that $\mathrm{M} \models \neg\mathbf{B}\varphi$ for some $\varphi \in \mathcal{L}_0$ and let $\mathrm{M}_1 = \langle \mathrm{S}, \mathrm{B}_1, \sigma_1 \rangle = \mathtt{r}^\star(\texttt{expand } \varphi, \mathrm{M})$ and $\mathrm{M}_2 = \langle \mathrm{S}, \mathrm{B}_2, \sigma_2 \rangle = \mathtt{r}^\star(\texttt{contract } \varphi, \mathrm{M}_1)$. We show that $\mathrm{M}_2 = \mathrm{M}$, i.e. $\mathrm{B}_2 = \mathrm{B}$ and $\sigma_2 = \sigma$. Since $\mathrm{M} \models \neg\mathbf{B}\varphi$, $\mathrm{B} \cap [\![\neg\varphi]\!] \neq \emptyset$, and since $\mathrm{B}_1 \cap [\![\neg\varphi]\!] = \emptyset$ it follows that $\sigma_1(\neg\varphi) = \mathrm{B} \cap [\![\varphi]\!]$. Hence $\mathrm{B}_2 = \mathrm{B}_1 \cup \sigma_1(\varphi) = (\mathrm{B} \cap [\![\varphi]\!]) \cup (\mathrm{B} \cap [\![\neg\varphi]\!]) = \mathrm{B}$. This leaves to show that $\sigma_2 = \sigma$. So let $\vartheta \in \mathcal{L}_0$ be arbitrary. If $\mathrm{B} \cap [\![\vartheta]\!] \neq \emptyset$, then also $\mathrm{B}_2 \cap [\![\vartheta]\!] \neq \emptyset$. Hence $\sigma(\vartheta) = \emptyset$ and $\sigma_2(\vartheta) = \emptyset$. So assume $\mathrm{B} \cap [\![\vartheta]\!] = \emptyset$. Then since $\mathrm{B}_1 \subseteq \mathrm{B}$ and $\mathrm{B}_2 = \mathrm{B}$, also $\mathrm{B}_1 \cap [\![\vartheta]\!] = \emptyset$ and $\mathrm{B}_2 \cap [\![\vartheta]\!] = \emptyset$. In this case $\sigma_2(\vartheta) = \sigma_1(\vartheta) = \sigma(\vartheta)$. Thus in both cases $\sigma_2(\vartheta) = \sigma(\vartheta)$, and since ϑ is arbitrary this suffices to conclude that $\sigma_2 = \sigma$. □

Combining Proposition 4.3 with clause 6 of Proposition 2.6 leads to the following corollary.

4.4. Corollary. *For all $\varphi \in \mathcal{L}_0$, $\rho \in \mathcal{L}$ we have:*

- $\models_\star \neg\mathbf{B}\varphi \to ([\texttt{revise } \varphi; \texttt{contract } \varphi]\rho \leftrightarrow [\texttt{contract } \neg\varphi]\rho)$

The property expressed in Corollary 4.4 is quite intuitive if one considers a revision to consist of a 'negative part' (a contraction), followed by a 'positive' part (an expansion): by the postulate of rehabilitation this latter positive part is undone by the consecutive contraction, leaving only the negative first part to be effective.

The following proposition formalises various properties of sequences of belief changes, mostly at the level of belief sets. To make a concise representation possible, we denote theoremhood through the knowledge operator. Since a formula is a theorem in classic propositional logic iff it is known in a full model and full models constitute a special subclass of $\mathcal{M}^\star$, this does not affect the applicability or generality of the validities that we prove here.

4.5. Proposition. *For all $\varphi, \psi, \vartheta \in \mathcal{L}_0$ we have:*

1. $\models_\star ([\texttt{contract } \varphi]\mathbf{B}\vartheta \wedge [\texttt{contract } \psi]\mathbf{B}\vartheta) \to [\texttt{contract } \varphi; \texttt{contract } \psi]\mathbf{B}\vartheta$
2. $\models_\star \mathbf{K}(\varphi \to \psi) \to ([\texttt{contract } \varphi; \texttt{contract } \psi]\mathbf{B}\vartheta \to [\texttt{contract } \psi]\mathbf{B}\vartheta)$

3. $\models_\star \mathbf{K}(\psi \to \varphi) \to ([\texttt{contract}\ \varphi;\texttt{revise}\ \psi]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise}\ \psi]\mathbf{B}\vartheta)$
4. $\mathbf{K}(\psi \to \varphi) \to ([\texttt{contract}\ \varphi;\texttt{revise}\ \psi]\rho \leftrightarrow [\texttt{revise}\ \psi]\rho)$ *is not for all* $\varphi, \psi \in \mathcal{L}_0$, $\rho \in \mathcal{L}$ *valid*

Proof: We show all four items. Let M be some model with state s, and let $\varphi, \psi, \vartheta \in \mathcal{L}_0$ be arbitrary.

1. Assume that $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi]\mathbf{B}\vartheta \wedge [\texttt{contract}\ \psi]\mathbf{B}\vartheta$. The only interesting case is where $\mathrm{M}, s \models_\star \mathbf{B}\varphi \wedge \mathbf{B}\psi$, so assume this to be true. In this case $\mathrm{M}_1 = \mathtt{r}^\star(\texttt{contract}\ \varphi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\varphi), \sigma_1 \rangle$, and $\mathrm{M}_2 = \mathtt{r}^\star(\texttt{contract}\ \psi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\psi), \sigma_2 \rangle$. If $\mathrm{M}_1, s \models_\star \neg\mathbf{B}\psi$, it follows that $\mathtt{r}^\star(\texttt{contract}\ \varphi; \texttt{contract}\ \psi, \mathrm{M}) = \mathrm{M}_1$, and then $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi]\mathbf{B}\vartheta$ implies $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi; \texttt{contract}\ \psi]\mathbf{B}\vartheta$. If $\mathrm{M}_1, s \models_\star \mathbf{B}\psi$, then $\mathtt{r}^\star(\texttt{contract}\ \varphi; \texttt{contract}\ \psi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\varphi) \cup \sigma_1(\neg\psi), \sigma' \rangle$. Since $(\mathrm{B} \cup \sigma(\neg\varphi)) \cap [\![\neg\psi]\!] = \emptyset$, it follows by the first clause of Definition 3.2 for contractions that $\sigma_1(\neg\psi) = \sigma(\neg\psi)$. Now since $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi]\mathbf{B}\vartheta$, ϑ holds at all worlds from $\mathrm{B} \cup \sigma(\neg\varphi)$. Since $\mathrm{M}, s \models_\star [\texttt{contract}\ \psi]\mathbf{B}\vartheta$, ϑ holds at all worlds from $\mathrm{B} \cup \sigma(\neg\psi)$. Hence ϑ holds at all worlds from $\mathrm{B} \cup \sigma(\neg\varphi) \cup \sigma(\neg\psi)$, and thus $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi; \texttt{contract}\ \psi]\mathbf{B}\vartheta$, which was to be shown.
2. Assume that $\mathrm{M}, s \models_\star \mathbf{K}(\varphi \to \psi)$, and, since this is the only interesting case, assume also that $\mathrm{M}, s \models_\star \mathbf{B}\varphi \wedge \mathbf{B}\psi$. Let $\mathrm{M}_1 = \mathtt{r}^\star(\texttt{contract}\ \varphi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\varphi), \sigma_1 \rangle$ and $\mathrm{M}_2 = \mathtt{r}^\star(\texttt{contract}\ \psi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\psi), \sigma_2 \rangle$. We distinguish two case:
   - $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi]\mathbf{B}\psi$. Then $\mathtt{r}^\star(\texttt{contract}\ \varphi; \texttt{contract}\ \psi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\varphi) \cup \sigma_1(\neg\psi), \sigma' \rangle$. Since $(\mathrm{B} \cup \sigma(\neg\varphi)) \cap [\![\neg\psi]\!] = \emptyset$ we have by clause 1 of Definition 3.2 for contractions that $\sigma_1(\neg\psi) = \sigma(\neg\psi)$. Thus $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi; \texttt{contract}\ \psi]\mathbf{B}\vartheta$ implies that ϑ holds at all worlds from $\mathrm{B} \cup \sigma(\neg\varphi) \cup \sigma(\neg\psi)$. In particular this implies that ϑ holds at all worlds from $\mathrm{B} \cup \sigma(\neg\psi)$, and hence $\mathrm{M}, s \models_\star [\texttt{contract}\ \psi]\mathbf{B}\vartheta$.
   - $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi]\neg\mathbf{B}\psi$. Then $\mathtt{r}(\texttt{contract}\ \varphi; \texttt{contract}\ \psi, \mathrm{M}) = \langle \mathrm{S}, \mathrm{B} \cup \sigma(\neg\varphi), \sigma' \rangle$. Since $\mathrm{M}, s \models_\star \mathbf{K}(\varphi \to \psi)$, also $\mathrm{M}, s \models_\star \mathbf{K}((\neg\psi \vee \neg\varphi) \leftrightarrow \neg\varphi)$. Now since $\mathrm{M}, s \models_\star \mathbf{B}\psi$ and $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi]\neg\mathbf{B}\psi$ it follows that $\sigma(\neg\varphi) \cap [\![\neg\psi]\!] \neq \emptyset$. From $\Sigma 3$ it follows that $\sigma(\neg\psi \vee \neg\varphi) \cap [\![\neg\psi]\!] \neq \emptyset$, and by $\Sigma 5$, $\sigma(\neg\psi) \subseteq \sigma(\neg\psi \vee \neg\varphi) = \sigma(\neg\varphi)$. Hence $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi; \texttt{contract}\ \psi]\mathbf{B}\vartheta$ implies that ϑ holds at all worlds in $\mathrm{B} \cup \sigma(\neg\psi)$, and hence $\mathrm{M}, s \models_\star [\texttt{contract}\ \psi]\mathbf{B}\vartheta$.

   Since in both cases it holds that $\mathrm{M}, s \models_\star [\texttt{contract}\ \varphi; \texttt{contract}\ \psi]\mathbf{B}\vartheta \to [\texttt{contract}\ \psi]\mathbf{B}\vartheta$, we conclude that this clause is indeed true.
3. Assume that $\mathrm{M}, s \models_\star \mathbf{K}(\psi \to \varphi)$. If $\mathrm{M}, s \models_\star \neg\mathbf{B}\varphi$ the property holds trivially. Hence assume $\mathrm{M}, s \models_\star \mathbf{B}\varphi$. Let $\mathrm{M}_1 = \mathtt{r}^\star(\texttt{contract}\ \varphi, \mathrm{M})$. We distinguish two cases:
   - $\mathrm{M}, s \models_\star \mathbf{B}\neg\psi$. From $\mathrm{M}, s \models_\star \mathbf{K}(\psi \to \varphi)$ it follows that $\sigma(\neg\varphi) \subseteq [\![\neg\psi]\!]$. Hence $\neg\psi$ holds at all worlds from $\mathrm{B} \cup \sigma(\neg\varphi)$. By clause 1 of Definition 3.2 for contractions it follows that $\sigma_1(\psi) = \sigma(\psi)$. Hence $\mathtt{r}^\star(\texttt{revise}\ \psi, \mathrm{M}_1) = \langle \mathrm{S}, \sigma(\psi), \sigma' \rangle$. Also, $\mathtt{r}^\star(\texttt{revise}\ \psi, \mathrm{M}) =$

$\langle \mathrm{S}, \sigma(\psi), \sigma'' \rangle$. Thus $\mathrm{M}, s \models_\star$ [contract φ; revise ψ]$\mathbf{B}\vartheta$ iff $\mathrm{M}, s \models_\star$ [revise ψ]$\mathbf{B}\vartheta$.

- $\mathrm{M}, s \models_\star \neg\mathbf{B}\neg\psi$. Then also $\mathrm{M}_1, s \models_\star \neg\mathbf{B}\neg\psi$, and $\mathtt{r}^\star$(revise ψ, M_1) $=$ $\langle \mathrm{S}, (\mathrm{B} \cup \sigma(\neg\varphi)) \cap [\![\psi]\!], \sigma' \rangle$. Since $\mathrm{M}, s \models_\star \mathbf{K}(\psi \to \varphi)$, $\sigma(\neg\varphi) \subseteq [\![\neg\psi]\!]$, and hence $(\mathrm{B} \cup \sigma(\neg\varphi)) \cap [\![\psi]\!] = \mathrm{B} \cap [\![\psi]\!]$. Also $\mathtt{r}^\star$(revise ψ, M) $=$ $\langle \mathrm{S}, \mathrm{B} \cap [\![\psi]\!], \sigma'' \rangle$, and thus $\mathrm{M}, s \models_\star$ [contract φ; revise ψ]$\mathbf{B}\vartheta$ iff $\mathrm{M}, s \models_\star$ [revise ψ]$\mathbf{B}\vartheta$.

Since in both cases $\mathrm{M}, s \models_\star$ [contract φ; revise ψ]$\mathbf{B}\vartheta \leftrightarrow$ [revise ψ]$\mathbf{B}\vartheta$, we conclude that clause 3 of Proposition 4.5 is true.

4. This case is shown by the model M of Example 4.1. For it holds that $\mathrm{M} \models_\star$ [contract $(\neg p \vee \neg q)$; revise $\neg p$][revise p]$\mathbf{B}(p \wedge q)$ and $\mathrm{M} \not\models_\star$ [revise $\neg p$][revise p]$\mathbf{B}(p \wedge q)$. □

The first item of Proposition 4.5 states that whenever two single contractions both result in the agent believing a formula ϑ, then the agent also believes ϑ after the two contractions have been performed sequentially (in any order). Intuitively this is clear: if the agent does not believe the second formula that is to be retracted after the contraction of the first one, the second contraction amounts to a void action. Otherwise the belief in the second formula was unaffected by the first contraction and the second contraction is performed as if the first one did not happen. In both cases it is clear that ϑ is believed after the sequence of the two contractions. The second item states that whenever a formula ϑ is believed after two contractions such that the second one retracts a formula that is known to be implied by the first retracted formula, then ϑ is also believed after a single contraction with the second formula. Two important remarks need to be made concerning this validity, the first of these being that the property is not as obvious as it might seem at first sight. For one could think of a situation in which the formula ψ is believed before, and no longer believed after, the contraction with φ. In this case the second contraction in the sequence amounts to a void action, but it is not directly clear that the single contraction with ψ necessarily results in a belief set which contains the one that results from the contraction with φ, even if ψ is known to be implied by φ. The second remark is that the antecedent $\mathbf{K}(\varphi \to \psi)$ is necessary for validity: the consequent is in itself not valid. The third item states that a contraction with a formula φ that precedes a revision with a 'stronger' formula ψ has no effect at all on the belief set level: one might as well directly perform the revision. However, item 4 states that this property does not hold for belief systems. In particular, it is not the case that contract φ; revise ψ and revise ψ constrain future changes of belief in the same way whenever ψ is known to imply φ.

Iterated revisions can be completely characterised with respect to their direct effects.

4.6. PROPOSITION. *For all $\varphi, \psi, \vartheta \in \mathcal{L}_0$ we have:*

- $\models_\star$ [revise φ]$\mathbf{B}\neg\psi \to$ ([revise φ; revise ψ]$\mathbf{B}\vartheta \leftrightarrow$ [revise ψ]$\mathbf{B}\vartheta$)
- $\models_\star$ [revise φ]$\neg\mathbf{B}\neg\psi \to$ ([revise φ; revise ψ]$\mathbf{B}\vartheta \leftrightarrow$ [revise $(\varphi \wedge \psi)$]$\mathbf{B}\vartheta$)

Proof: Let $M = \langle S, B, \sigma \rangle$ be a Kripke model with $s \in S$, and let $\varphi, \psi, \vartheta \in \mathcal{L}$ be arbitrary. We show both clauses.

- Assume that $M, s \models_\star [\texttt{revise}\ \varphi]\mathbf{B}\neg\psi$. For reasons of convenience we introduce the following models:
  - $M_{11} = \mathtt{r}^\star(\texttt{contract}\ \neg\varphi, M)$
  - $M_1 = \mathtt{r}^\star(\texttt{revise}\ \varphi, M) = \mathtt{r}^\star(\texttt{expand}\ \varphi, M_{11})$
  - $M_2 = \mathtt{r}^\star(\texttt{revise}\ \psi, M_1) = \mathtt{r}^\star(\texttt{revise}\ \varphi; \texttt{revise}\ \psi, M)$
  - $M_3 = \mathtt{r}^\star(\texttt{revise}\ \psi, M)$

  For $i = 1, 2, 3, 11$ we assume M_i to be the tuple $\langle S, B_i, \sigma_i \rangle$. Note that since $M, s \models_\star [\texttt{revise}\ \varphi]\mathbf{B}\neg\psi$, $B_2 = \sigma_1(\psi)$. Hence it suffices to show that B_3 is equal to $\sigma_1(\varphi)$: for $M, s \models_\star [\texttt{revise}\ \varphi; \texttt{revise}\ \psi]\mathbf{B}\vartheta$ iff ϑ holds at all worlds from $\sigma_1(\psi)$ and $M, s \models_\star [\texttt{revise}\ \psi]\mathbf{B}\vartheta$ iff ϑ holds at all worlds from B_3. We distinguish four cases:

  1. $M, s \models_\star \mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\varphi$. In this case the property trivially holds since $M_1 = M$ and $M_2 = M_3$.
  2. $M, s \models_\star \neg\mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\varphi$. Then $M_{11} = M$ and $M_1 = \mathtt{r}^\star(\texttt{expand}\ \varphi, M)$. We distinguish two cases:
     - $M, s \models_\star \neg\mathbf{B}\neg\psi$. In this case $B \cap [\![\psi]\!] \neq \emptyset$, and since $B_1 \cap [\![\psi]\!] = \emptyset$, it follows that $\sigma_1(\psi) = B \cap [\![\psi]\!]$. Also, if $M, s \models_\star \neg\mathbf{B}\neg\psi$ then $B_3 = B \cap [\![\psi]\!]$. Hence in this case $\sigma_1(\varphi) = B_3$.
     - $M, s \models_\star \mathbf{B}\neg\psi$. In this case $B \cap [\![\psi]\!] = \emptyset$ and hence $\sigma_1(\psi) = \sigma(\psi)$. Furthermore, if $M, s \models_\star \mathbf{B}\neg\psi$ also $B_3 = \sigma(\psi)$. Hence also in this case $\sigma_1(\psi) = B_3$.

     Since in both cases $\sigma_1(\psi) = B_3$, we conclude that whenever $M, s \models_\star \neg\mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\varphi$, $\sigma_1(\varphi) = B_3$.
  3. $M, s \models_\star \neg\mathbf{B}\varphi \wedge \mathbf{B}\neg\varphi$. In this case $B_{11} = B \cup \sigma(\varphi)$ and $B_1 = \sigma(\varphi)$. Again we distinguish two cases:
     - $M, s \models_\star \neg\mathbf{B}\neg\psi$. Then since $B \cap [\![\psi]\!] \neq \emptyset$, also $B_{11} \cap [\![\psi]\!] \neq \emptyset$. Since $M, s \models_\star [\texttt{revise}\ \varphi]\mathbf{B}\neg\psi$, $B_1 \cap [\![\psi]\!] = \emptyset$. Hence by Definition 3.2 for expansions, $\sigma_1(\psi) = (B \cup \sigma(\varphi)) \cap [\![\psi]\!] = B \cap [\![\psi]\!]$. Also, if $M, s \models_\star \neg\mathbf{B}\neg\psi$, $B_3 = B \cap [\![\psi]\!]$, and hence $\sigma_1(\psi) = B_3$.
     - $M, s \models_\star \mathbf{B}\neg\psi$. Then $B \cap [\![\psi]\!] = \emptyset$. Since $M, s \models_\star [\texttt{revise}\ \varphi]\mathbf{B}\neg\psi$ also $\sigma(\varphi) \cap [\![\psi]\!] = \emptyset$. Hence $B_{11} \cap [\![\psi]\!] = (B \cup \sigma(\varphi)) \cap [\![\psi]\!] = \emptyset$. Then by Definition 3.2 for contractions, $\sigma_{11}(\psi) = \sigma(\psi)$. Since $B_{11} \cap [\![\psi]\!] = \emptyset$, $\sigma_1(\psi) = \sigma_{11}(\psi) = \sigma(\psi)$. From $M, s \models_\star \mathbf{B}\neg\psi$ it follows that $B_3 = \sigma(\psi)$, and hence $\sigma_1(\psi) = B_3$.

     Since in both cases $\sigma_1(\psi) = B_3$, we conclude that whenever $M, s \models_\star \neg\mathbf{B}\varphi \wedge \mathbf{B}\neg\varphi$, $\sigma_1(\varphi) = B_3$.
  4. $M, s \models_\star \mathbf{B}\varphi \wedge \mathbf{B}\neg\varphi$. Then $B = \emptyset$ and both B_{11} and B_1 are equal to $\sigma(\varphi)$. Since $M, s \models_\star [\texttt{revise}\ \varphi]\mathbf{B}\neg\psi$, $\sigma(\varphi) \cap [\![\psi]\!] = \emptyset$. Hence $B_{11} \cap [\![\psi]\!] = \emptyset$ and thus $\sigma_{11}(\psi) = \sigma(\psi)$. Since $B_{11} \cap [\![\psi]\!] = \emptyset$, $\sigma_1(\psi) = \sigma_{11}(\psi) = \sigma(\psi)$. Furthermore, since $M, s \models_\star \mathbf{B}\neg\psi$, $B_3 = \sigma(\psi)$, which suffices to conclude that $\sigma_1(\psi) = B_3$.

  In all four cases, $\sigma_1(\psi) = B_3$. Hence $M, s \models_\star [\texttt{revise}\ \varphi; \texttt{revise}\ \psi]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise}\ \psi]\mathbf{B}\vartheta$, which concludes the proof of clause 1 of Proposition 4.6.
- Assume that $M, s \models_\star [\texttt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$. We distinguish three cases.

1. $M, s \models_\star \mathbf{B}\neg\varphi$. In this case $\mathtt{r}^\star(\mathtt{revise}\ \varphi, M) = \langle S, \sigma(\varphi), \sigma' \rangle$, and since $M, s \models_\star [\mathtt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$, $\mathtt{r}^\star(\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi, M) = \langle S, \sigma(\varphi) \cap [\![\psi]\!], \sigma'' \rangle$. On the other hand $\mathtt{r}^\star(\mathtt{revise}\ \varphi \wedge \psi, M) = \langle S, \sigma(\varphi \wedge \psi), \sigma'' \rangle$ since $M, s \models \mathbf{B}\neg(\varphi \wedge \psi)$. We show that $\sigma(\varphi) \cap [\![\psi]\!] = \sigma(\varphi \wedge \psi)$ by showing that the two sets are contained in each other.
   - '$\subseteq$' By $\Sigma 3$, $\sigma(\varphi) = \sigma((\varphi \wedge \psi) \vee (\varphi \wedge \neg\psi))$, which is by $\Sigma 4$ contained in $\sigma(\varphi \wedge \psi) \cup \sigma(\varphi \wedge \neg\psi)$. Since by $\Sigma 1$, $\sigma(\varphi \wedge \neg\psi) \subseteq [\![\varphi \wedge \neg\psi]\!]$, it follows that $\sigma(\varphi) \cap [\![\psi]\!] \subseteq \sigma(\varphi \wedge \psi)$.
   - '$\supseteq$' Since $M, s \models_\star [\mathtt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$ it follows that $\sigma(\varphi) \cap [\![\psi]\!] \neq \emptyset$. Now by $\Sigma 3$, $\sigma(\varphi) = \sigma((\varphi \wedge \psi) \vee \varphi)$, by $\Sigma 1$, $\sigma(\varphi) \subseteq [\![\varphi]\!]$, and hence $\sigma((\varphi \wedge \psi) \vee \varphi) \cap [\![\varphi \wedge \psi]\!] \neq \emptyset$. Then by $\Sigma 5$, $\sigma(\varphi \wedge \psi) \subseteq \sigma(\varphi)$. Since by $\Sigma 1$, $\sigma(\varphi \wedge \psi) \subseteq [\![\varphi \wedge \psi]\!]$, it follows that $\sigma(\varphi \wedge \psi) \subseteq \sigma(\varphi) \cap [\![\psi]\!]$.

   Now $M, s \models_\star [\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\mathbf{B}\vartheta$ iff ϑ holds at all worlds from $\sigma(\varphi) \cap [\![\psi]\!]$ iff ϑ holds at all worlds from $\sigma(\varphi \wedge \psi)$ iff $M, s \models_\star [\mathtt{revise}\ \varphi \wedge \psi]\mathbf{B}\vartheta$.
2. $M, s \models_\star \mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\varphi$. By Corollary 3.9(2), $\mathtt{r}^\star(\mathtt{revise}\ \varphi, M) = M$. Since $M, s \models_\star [\mathtt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$ also $M, s \models_\star \neg\mathbf{B}\neg\psi$. Now $M, s \models_\star \mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\psi$ implies $M, s \models_\star \neg\mathbf{B}\neg(\varphi \wedge \psi)$ and thus $\mathtt{r}^\star(\mathtt{revise}\ \varphi \wedge \psi, M) = \langle S, B \cap [\![\varphi \wedge \psi]\!], \sigma' \rangle$. Now $\mathtt{r}^\star(\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi, M) = \langle S, B \cap [\![\psi]\!], \sigma'' \rangle$. Since $B \subseteq [\![\varphi]\!]$ it follows that $B \cap [\![\varphi \wedge \psi]\!] = B \cap [\![\psi]\!]$. Hence $M, s \models_\star [\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\mathbf{B}\vartheta$ iff ϑ holds at all worlds from $B \cap [\![\psi]\!]$ iff ϑ holds at all worlds from $B \cap [\![\varphi \wedge \psi]\!]$ iff $M, s \models_\star [\mathtt{revise}\ \varphi \wedge \psi]\mathbf{B}\vartheta$.
3. $M, s \models_\star \neg\mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\varphi$. Then from Corollary 3.9(1) it follows that $\mathtt{r}^\star(\mathtt{revise}\ \varphi, M) = \langle S, B \cap [\![\varphi]\!], \sigma' \rangle$. Since $M, s \models_\star [\mathtt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$, $\mathtt{r}^\star(\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi, M) = \langle S, B \cap [\![\varphi]\!] \cap [\![\psi]\!], \sigma'' \rangle$. From $M, s \models_\star \neg\mathbf{B}\neg\varphi \wedge [\mathtt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$, it follows by Proposition 2.23(4) and Proposition 2.21 that $M, s \models_\star \neg\mathbf{B}(\varphi \to \neg\psi)$, i.e. $M, s \models_\star \neg\mathbf{B}\neg(\varphi \wedge \psi)$. Hence $\mathtt{r}^\star(\mathtt{revise}\ \varphi \wedge \psi, M) = \langle S, B \cap [\![\varphi \wedge \psi]\!], \sigma'' \rangle$. Now $B \cap [\![\varphi]\!] \cap [\![\psi]\!] = B \cap [\![\varphi \wedge \psi]\!]$, and thus $M, s \models_\star [\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\mathbf{B}\vartheta$ iff $M, s \models_\star [\mathtt{revise}\ \varphi \wedge \psi]\mathbf{B}\vartheta$.

Since in all three cases, $M, s \models_\star [\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\mathbf{B}\vartheta \leftrightarrow [\mathtt{revise}\ \varphi \wedge \psi]\mathbf{B}\vartheta$, we conclude that clause 2 of Proposition 4.6 indeed holds. □

The first item of Proposition 4.6 states that if two consecutive revisions are such that the second revision asserts a formula that is inconsistent with the belief set resulting from the first revision, then the second revision prevails. That is, the second revision alone results in the same set of beliefs as the sequence of the two revisions. The second item states that whenever the second revision of a sequence of two revisions asserts a formula that is consistent with the set of beliefs that results from the first revision, then a revision with the conjunction of the two asserted formulae results in the same set of beliefs as the sequence of the two revisions. In both cases the sequence of the two revisions may be reduced to a single revision as far as the beliefs of the agent are concerned. However, this reduction does not go through at the transmutation level.

4.7. PROPOSITION. *The following are not for all $\varphi, \psi \in \mathcal{L}_0$ and $\rho \in \mathcal{L}$ valid:*

- $[\mathtt{revise}\ \varphi]\mathbf{B}\neg\psi \to ([\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\rho \leftrightarrow [\mathtt{revise}\ \psi]\rho)$

- $[\texttt{revise } \varphi]\neg\mathbf{B}\neg\psi \rightarrow ([\texttt{revise } \varphi;\texttt{revise } \psi]\rho \leftrightarrow [\texttt{revise } (\varphi \wedge \psi)]\rho)$

Proof: Both items are shown by the model given in Example 4.1. The invalidity of the first item follows from the observation that although $\mathrm{M} \models [\texttt{revise } p]\mathbf{B}\neg\neg p$, $\mathrm{M} \models [\texttt{revise } p;\texttt{revise } \neg p][\texttt{revise } q]\mathbf{B}(p \wedge q)$ and $\mathrm{M} \not\models [\texttt{revise } \neg p][\texttt{revise } q]\mathbf{B}(p \wedge q)$. The second item follows from combining $\mathrm{M} \models [\texttt{revise } p]\neg\mathbf{B}\neg(p \wedge \neg q)$ with the observations made in Example 4.1. □

A discussion on the importance of the results of Proposition 4.7 is postponed to the end of §4.1.

4.1. Postulates for iterated belief revision

Here we look into several extensions of, and alternatives to, the AGM axiomatisation, all aimed at providing an intuitive account of iterated belief revision. We are mainly interested in the possible validity of the various postulates; for thorough discussions on the desirability and intuitive acceptability of the postulates we refer to the papers in which they were proposed. The notion of validity is in this section to be understood as validity in the class $\mathcal{M}^\star$ of models. Whenever possible, we will interpret the proposed postulates both at the belief set level and at the level of belief systems.

Darwiche & Pearl [3] propose four additional postulates that, in combination with the AGM postulates, should take care of intuitively acceptable iterated revision. Rephrased in our framework, they are the following.

4.8. Definition. For $\varphi, \psi, \vartheta \in \mathcal{L}_0$, the following are the postulates of Darwiche & Pearl:

DP1. $\mathbf{K}(\psi \rightarrow \varphi) \rightarrow ([\texttt{revise } \varphi;\texttt{revise } \psi]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise } \psi]\mathbf{B}\vartheta)$
DP2. $\mathbf{K}(\psi \rightarrow \neg\varphi) \rightarrow ([\texttt{revise } \varphi;\texttt{revise } \psi]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise } \psi]\mathbf{B}\vartheta)$
DP3. $[\texttt{revise } \psi]\mathbf{B}\varphi \rightarrow [\texttt{revise } \varphi;\texttt{revise } \psi]\mathbf{B}\varphi$
DP4. $[\texttt{revise } \psi]\neg\mathbf{B}\neg\varphi \rightarrow [\texttt{revise } \varphi;\texttt{revise } \psi]\neg\mathbf{B}\neg\varphi$

Freund & Lehmann argue that 'the postulate [DP2], contrary to the claims of Darwiche and Pearl is inconsistent with the AGM axioms' ([4],p. 8), i.e. there is no revision that satisfies both the AGM postulates and DP2. Although this inconsistency goes through when adopting the static view to belief change, it does not in our dynamic account.

4.9. Proposition. *The postulates of Darwiche & Pearl are valid for all φ, ψ and ϑ in $\mathcal{L}_0$.*

Proof: Postulate DP1 is seen to be valid by combining both clauses of Proposition 4.6 with clause 6 of Proposition 2.23. DP2 follows from the first clause of Proposition 4.6, the second clause of Proposition 2.23 and the fact that knowledge on propositional formulae persists. For postulate DP3 assume that M is some Kripke model with state s, and let $\varphi, \psi \in \mathcal{L}_0$ be arbitrary. Assume that $\mathrm{M}, s \models_\star [\texttt{revise } \psi]\mathbf{B}\varphi$. Now if $\mathrm{M}, s \models_\star [\texttt{revise } \varphi]\mathbf{B}\neg\psi$, then by the first clause of Proposition 4.6 it follows that that $\mathrm{M}, s \models_\star [\texttt{revise } \varphi;\texttt{revise } \psi]\mathbf{B}\varphi$.

If $M, s \models_\star [\texttt{revise } \varphi]\neg\mathbf{B}\neg\psi$ then by clause 2 of Proposition 4.6, $M, s \models_\star [\texttt{revise } \varphi; \texttt{revise } \psi]\mathbf{B}\varphi$ iff $M, s \models_\star [\texttt{revise } \varphi \wedge \psi]\mathbf{B}\varphi$. From clauses 7 and 8 of Proposition 2.23 it follows that $M, s \models_\star [\texttt{revise } \varphi \wedge \psi]\mathbf{B}\varphi$ iff $M, s \models_\star [\texttt{revise } \varphi; \texttt{expand } \psi]\mathbf{B}\varphi$. Now since $M, s \models_\star [\texttt{revise } \varphi]\mathbf{B}\varphi$ it follows by Proposition 2.21 that also $M, s \models_\star [\texttt{revise } \varphi; \texttt{expand } \psi]\mathbf{B}\varphi$, which suffices to conclude that DP3 is indeed valid. The only interesting case for DP4 is where $M, s \models_\star [\texttt{revise } \varphi]\neg\mathbf{B}\neg\psi$. In this case $M, s \models_\star [\texttt{revise } \varphi; \texttt{revise } \psi]\neg\mathbf{B}\neg\varphi$ iff $M, s \models_\star [\texttt{revise } \varphi \wedge \psi]\neg\mathbf{B}\neg\varphi$. Now since $M, s \models_\star [\texttt{revise } \psi]\neg\mathbf{B}\neg\varphi$, it follows that $M, s \not\models_\star \mathbf{K}(\neg\varphi \wedge \neg\psi)$. Then by clauses 2 and 5 of Proposition 2.23 it follows that $M, s \models_\star [\texttt{revise } \varphi \wedge \psi](\mathbf{B}(\varphi \wedge \psi) \wedge \neg\mathbf{B}\bot)$. But then also $M, s \models_\star [\texttt{revise } \varphi \wedge \psi]\neg\mathbf{B}\neg\varphi$, and thus $M, s \models_\star [\texttt{revise } \varphi; \texttt{revise } \psi]\neg\mathbf{B}\neg\varphi$. □

In our opinion, the statement that DP2 is inconsistent, *contrary to the claims of Darwiche & Pearl*, is not correct. The soundness proof that Darwiche & Pearl provide is with regard to a *dynamic* semantics for belief revision, and it follows from Remark 3.5 and Proposition 4.9 that a dynamic semantics might well be capable of validating both the AGM axiomatisation *and* postulate DP2.

Although Darwiche & Pearl deal with changes to belief sets only, two of their postulates do make sense — and are in fact not valid — at the level of belief systems.

4.10. Proposition. *The following are not for all $\varphi, \psi \in \mathcal{L}_0$, $\rho \in \mathcal{L}$ valid:*
DP1'. $\mathbf{K}(\psi \to \varphi) \to ([\texttt{revise } \varphi; \texttt{revise } \psi]\rho \leftrightarrow [\texttt{revise } \psi]\rho)$
DP2'. $\mathbf{K}(\psi \to \neg\varphi) \to ([\texttt{revise } \varphi; \texttt{revise } \psi]\rho \leftrightarrow [\texttt{revise } \psi]\rho)$

Proof: Both items are shown in exactly the same way as the corresponding items of Proposition 4.7. □

Nayak, Foo, Pagnucco and Sattar [14] proposed a set of postulates that is meant to serve as an alternative to the set proposed by Darwiche & Pearl. They propose to add the following postulates to the AGM axiomatisation.

4.11. Definition. For $\varphi, \psi, \vartheta \in \mathcal{L}_0$, the following are the postulates of Nayak *et al.*:
N0. $\mathbf{B}\bot \to ([\texttt{revise } \varphi]\mathbf{B}\vartheta \leftrightarrow \mathbf{K}(\varphi \to \vartheta))$
N7. $\mathbf{M}(\varphi \wedge \psi) \to ([\texttt{revise } \varphi; \texttt{revise } \psi]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise } (\varphi \wedge \psi)]\mathbf{B}\vartheta)$
N8. $\mathbf{K}\neg(\varphi \wedge \psi) \wedge \mathbf{M}\varphi \to ([\texttt{revise } \varphi; \texttt{revise } \psi]\mathbf{B}\vartheta \leftrightarrow [\texttt{revise } \psi]\mathbf{B}\vartheta)$

After showing that these postulates are inconsistent with the AGM axiomatisation when regarded from a static perspective, Nayak *et al.* argue for a dynamic account of belief change. However, it turns out that not all postulates are valid for our dynamic semantics.

4.12. Proposition. *The first two postulates of Nayak* et al. *are not for all $\varphi, \psi, \vartheta \in \mathcal{L}_0$ valid; postulate N8 is valid for all $\varphi, \psi, \vartheta \in \mathcal{L}_0$.*

Proof: Postulate N8 is a weakened version of DP2, in which the antecedent $\mathbf{K}(\psi \to \neg\varphi)$ is strengthened to $\mathbf{K}(\psi \to \neg\varphi) \wedge \mathbf{M}\varphi$. To see that N0 is not valid, take a belief set K such that some $\varphi, \vartheta \in \mathcal{L}_0$ exist with $\{\varphi, \vartheta\} \subseteq K$ and $\not\vdash_{cpl} (\varphi \to \vartheta)$. Let $\mathrm{M}_K = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle$ be a K-model. If $\mathrm{M}' = \mathtt{r}^\star(\mathtt{expand}\ \neg\varphi, \mathrm{M}_K)$, then $\mathrm{M}' = \langle \mathrm{S}, \emptyset, \sigma' \rangle$, where $\sigma'(\varphi) = \mathrm{B} \cap [\![\varphi]\!] = \mathrm{B}$. Now for arbitrary $s \in \mathrm{S}$ we have that $\mathrm{M}', s \models_\star \mathbf{B}\bot$, $\mathrm{M}', s \models_\star [\mathtt{revise}\ \varphi]\mathbf{B}\vartheta$ but $\mathrm{M}', s \not\models_\star \mathbf{K}(\varphi \to \vartheta)$, since $\not\vdash_{cpl} (\varphi \to \vartheta)$. Hence $\mathrm{M}', s \not\models_\star \mathbf{B}\bot \to ([\mathtt{revise}\ \varphi]\mathbf{B}\vartheta \leftrightarrow \mathbf{K}(\varphi \to \vartheta))$. Postulate N7 is shown to be invalid as follows. Take two different propositional symbols p and q, and define the full model $\mathrm{M} = \langle \mathrm{S}, \mathrm{B}, \sigma \rangle$ where $\mathrm{B} = [\![p \wedge \neg q]\!]$ and σ is the All-is-Good function for $\langle \mathrm{S}, \mathrm{B} \rangle$. Since M is full and $p, q \in \Pi$ are arbitrary, $\mathrm{M} \models_\star \mathbf{M}(p \wedge q)$. Now $\mathtt{r}^\star(\mathtt{revise}\ p; \mathtt{revise}\ q, \mathrm{M}) = \langle \mathrm{S}, [\![q]\!], \sigma' \rangle$ and $\mathtt{r}^\star(\mathtt{revise}\ p \wedge q, \mathrm{M}) = \langle \mathrm{S}, [\![p \wedge q]\!], \sigma'' \rangle$. Since $q \not\models_{cpl} p$, $[\![q]\!] \cap [\![\neg p]\!] \neq \emptyset$, and hence $\mathrm{M}, s \not\models_\star [\mathtt{revise}\ p; \mathtt{revise}\ q]\mathbf{B}p$, while $\mathrm{M}, s \models_\star [\mathtt{revise}\ p \wedge q]\mathbf{B}p$. Thus postulate N7 is not for all $\varphi, \psi \in \mathcal{L}_0$ valid. □

That postulate N0 is not valid, is not really surprising. For this postulate is essentially a *static* one which neglects the possibility that — even for inconsistent belief sets — changes of belief depend on other factors than just the set that is changed and the formula that the set is changed with. As such, this postulate clashes with the postulate of rehabilitation, and the example given in support of the latter postulate may be read as one against postulate N0. With respect to the non-validity of N7 note that this postulate may be interpreted as an unjustified strengthening of the second clause of Proposition 4.6: it is not sufficient that φ and ψ are *a priori* consistent, ψ needs to be consistent *after* the agent's beliefs have been revised with φ. The postulate N8': $\mathbf{K}\neg(\varphi \wedge \psi) \wedge \mathbf{M}\varphi \to ([\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\rho \leftrightarrow [\mathtt{revise}\ \psi]\rho)$ is not for all $\rho \in \mathcal{L}$ valid; this can be shown by a similar argument as the one that shows the invalidity of DP2'.

Having shown that both the postulates of Darwiche & Pearl and those of Nayak *et al.* are not suitable when adopting a static view, Lehmann [10] proposed an alternative account of iterated belief revision. Instead of the usual AGM belief sets, Lehmann considers sequences of revisions with consistent formulae as the fundamental notion. To replace the AGM postulates Lehmann proposes a new set of postulates that is suitable and usable for both static and dynamic approaches to belief revision. Formulated in the language $\mathcal{L}$, Lehmann's postulates are the following.

4.13. Definition. For $\varphi, \psi, \vartheta \in \mathcal{L}_0$ and $\rho \in \mathcal{L}$ the following are the postulates proposed by Lehmann:

- L1. $[\mathtt{revise}\ \varphi]\mathbf{B}\varphi$
- L2. $[\mathtt{revise}\ \varphi]\mathbf{B}\vartheta \to \mathbf{B}(\varphi \to \vartheta)$
- L3. $\mathbf{B}\varphi \wedge \neg\mathbf{B}\bot \to ([\mathtt{revise}\ \varphi]\rho \leftrightarrow \rho)$
- L4. $\mathbf{K}(\psi \to \varphi) \to ([\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\rho \leftrightarrow [\mathtt{revise}\ \psi]\rho)$
- L5. $[\mathtt{revise}\ \varphi]\neg\mathbf{B}\neg\psi \to ([\mathtt{revise}\ \varphi; \mathtt{revise}\ \psi]\rho \leftrightarrow$ $[\mathtt{revise}\ \varphi; \mathtt{revise}\ (\varphi \wedge \psi)]\rho)$
- L6. $[\mathtt{revise}\ \neg\varphi; \mathtt{revise}\ \varphi]\mathbf{B}\vartheta \to [\mathtt{revise}\ \varphi]\mathbf{B}\vartheta$

Postulates L1, L2 and L6 act on the belief set level, whereas the other postulates act on the level of belief systems. In its original form as given by

Lehmann, the prerequisite of L3 is $\mathbf{B}\varphi$. Since Lehmann assumes revisions of consistent belief sets with consistent formulae only, he does not have to take inconsistent belief sets into account. Belief sets may be inconsistent in our framework, which explains the additional $\neg\mathbf{B}\bot$ in the prerequisite of L3.

4.14. PROPOSITION. *Lehmann's postulates, with the exception of L4, are valid for all* φ, ψ, ϑ *in* $\mathcal{L}_0$ *and* $\rho \in \mathcal{L}$.

PROOF: The validity of the postulates L1, L2 and L3 follows directly from the validity of the AGM postulates. Postulate L6 follows from the first clause of Proposition 4.6. The invalidity of L4 follows directly from that of DP1'. The validity of L5 is seen as follows. Let $\mathrm{M} \models [\texttt{revise}\ \varphi]\neg\mathbf{B}\neg\psi$, and let $\mathrm{M}' = \langle \mathrm{S}, \mathrm{B}', \sigma' \rangle = \mathtt{r}^\star(\texttt{revise}\ \varphi, \mathrm{M})$. Then $\mathrm{M}' \models \mathbf{B}\varphi \wedge \neg\mathbf{B}\neg\psi$, and thus $\mathrm{M}' \models \neg\mathbf{B}\neg(\varphi \wedge \psi)$. Then $\mathtt{r}^\star(\texttt{revise}\ \psi, \mathrm{M}') = \mathtt{r}^\star(\texttt{expand}\ \psi, \mathrm{M}') = \langle \mathrm{S}, \mathrm{B}' \cap [\![\psi]\!], \sigma'_1 \rangle$ and $\mathtt{r}^\star(\texttt{revise}\ \varphi \wedge \psi, \mathrm{M}') = \mathtt{r}^\star(\texttt{expand}\ \varphi \wedge \psi, \mathrm{M}') = \langle \mathrm{S}, \mathrm{B}' \cap [\![\varphi \wedge \psi]\!], \sigma'_2 \rangle$. Now since $\mathrm{M}' \models \mathbf{B}\varphi$ it follows that $\mathrm{B}' \cap [\![\psi]\!] = \mathrm{B}' \cap [\![\varphi \wedge \psi]\!]$. From Proposition 3.4 it follows that $\sigma'_1(\vartheta) = \sigma'_2(\vartheta)$ for all $\vartheta \in \mathcal{L}_0$. Hence $\mathtt{r}^\star(\texttt{expand}\ \psi, \mathrm{M}') = \mathtt{r}^\star(\texttt{expand}\ \varphi \wedge \psi, \mathrm{M}')$, which suffices to conclude Proposition 4.14. □

Lehmann's postulate L4 is the weakest one in the set consisting of both items of Proposition 4.7 and the postulates DP1', DP2' and N8'. Moreover, L4 is the only genuine postulate for *belief systems*; the other mentioned postulates are generalisations of postulates that are originally aimed at describing and constraining iterated belief revision at the level of belief sets. The non-validity of L4 in our semantics is therefore considered to be more conspicuous than that of the other postulates. The reason for the non-validity lies in the incompatibility of L4 with the postulate of rehabilitation. Informally this incompatibility is due to the *historical awareness* associated with the latter postulate, which clashes with the flavour of *amnesia* associated with L4. For L4 states that the agent forgets that there ever was a revision with a formula φ after it subsequently performs a revision with some stronger formula ψ, whereas the historical awareness approach allows the revision with φ to have had some side-effects that persist under execution of the revision with ψ, thereby possibly influencing the future course of events. Phrased differently, Lehmann's agents do not remember, our agents do not forget. The incompatibility of these points of view can also be shown formally: extending the AGM axiomatisation with the postulate of rehabilitation precludes Lehmann's postulate L4 in some — fairly natural — situations.

4.15. PROPOSITION. *For all* $\varphi, \psi \in \mathcal{L}_0$ *we have:*

$$\begin{aligned}\models_\star\ & \mathbf{B}\neg\varphi \wedge \neg\mathbf{B}\bot \wedge \neg\mathbf{K}\neg\varphi \wedge [\texttt{revise}\ \varphi]\neg\mathbf{B}(\varphi \wedge \psi) \rightarrow \\ & ([\texttt{revise}\ \varphi; \texttt{revise}\ \varphi \wedge \psi][\texttt{contract}\ \varphi \wedge \psi]\mathbf{B}\varphi \wedge \\ & [\texttt{revise}\ \varphi \wedge \psi][\texttt{contract}\ \varphi \wedge \psi]\neg\mathbf{B}\varphi)\end{aligned}$$

PROOF: We start with a general Lemma that is of use in the proof of the proposition.

4.16. Lemma. *For all* $\varphi, \psi \in \mathcal{L}_0$ *we have:*
$\models_\star \psi \to \varphi \Rightarrow \models_\star \mathbf{B}\varphi \to [\texttt{contract}\ \neg\psi]\mathbf{B}\varphi$

Proof: The proof is straightforward: execution of the $\texttt{contract}\ \neg\psi$ action adds only worlds satisfying ψ — and hence satisfying φ — to the set of doxastic alternatives. Since all the elements from this set satisfied φ *a priori*, it follows that all elements from the new set of doxastic alternatives satisfy φ, and hence $\mathbf{B}\varphi$ holds after the contraction. □

Let $\mathrm{M} \in \mathcal{M}^\star$ be such that $\mathrm{M} \models_\star \mathbf{B}\neg\varphi \wedge \neg\mathbf{B}\bot \wedge \neg\mathbf{K}\neg\varphi \wedge [\texttt{revise}\ \varphi]\neg\mathbf{B}(\varphi\wedge\psi)$. Such an M exists by Proposition 4.17 given below. By Corollary 4.4 it follows that $\models_\star \neg\mathbf{B}(\varphi\wedge\psi) \to ([\texttt{revise}\ \varphi\wedge\psi; \texttt{contract}\ \varphi\wedge\psi]\rho \leftrightarrow [\texttt{contract}\ \neg\varphi\vee\neg\psi]\rho)$, for all $\rho \in \mathcal{L}$ (‡). Since $\mathrm{M} \models_\star [\texttt{revise}\ \varphi]\neg\mathbf{B}(\varphi \wedge \psi)$ it follows that $\mathrm{M} \models_\star [\texttt{revise}\ \varphi]([\texttt{revise}\ \varphi\wedge\psi; \texttt{contract}\ \varphi\wedge\psi]\rho \leftrightarrow [\texttt{contract}\ \neg\varphi\vee\neg\psi]\rho)$, and hence $\mathrm{M} \models_\star [\texttt{revise}\ \varphi][\texttt{revise}\ \varphi\wedge\psi; \texttt{contract}\ \varphi\wedge\psi]\rho \leftrightarrow [\texttt{revise}\ \varphi][\texttt{contract}\ \neg\varphi \vee \neg\psi]\rho$, for all $\rho \in \mathcal{L}$. By Lemma 4.16 we have that $\models_\star \mathbf{B}\varphi \to [\texttt{contract}\ \neg\varphi \vee \neg\psi]\mathbf{B}\varphi$. By Remark 3.5 we have that $\models_\star [\texttt{revise}\ \varphi]\mathbf{B}\varphi$, and therefore $\models_\star [\texttt{revise}\ \varphi][\texttt{contract}\ \neg\varphi \vee \neg\psi]\mathbf{B}\varphi$. Hence $\mathrm{M} \models_\star [\texttt{revise}\ \varphi][\texttt{revise}\ \varphi \wedge \psi; \texttt{contract}\ \varphi\wedge\psi]\mathbf{B}\varphi$. Now from $\mathrm{M} \models_\star \mathbf{B}\neg\varphi\wedge\neg\mathbf{B}\bot$ it follows that $\mathrm{M} \models_\star \neg\mathbf{B}\varphi$. Hence also $\mathrm{M} \models_\star [\texttt{contract}\ \neg\varphi \vee \neg\psi]\neg\mathbf{B}\varphi$, and thus by ‡, $\mathrm{M} \models_\star [\texttt{revise}\ \varphi \wedge \psi; \texttt{contract}\ \varphi \wedge \psi]\neg\mathbf{B}\varphi$, i.e. $\mathrm{M} \models_\star [\texttt{revise}\ \varphi \wedge \psi][\texttt{contract}\ \varphi \wedge \psi]\neg\mathbf{B}\varphi$. Thus $\mathrm{M} \models_\star [\texttt{revise}\ \varphi; \texttt{revise}\ \varphi \wedge \psi][\texttt{contract}\ \varphi \wedge \psi]\mathbf{B}\varphi$ and $\mathrm{M} \models_\star [\texttt{revise}\ \varphi \wedge \psi][\texttt{contract}\ \varphi \wedge \psi]\neg\mathbf{B}\varphi$, which was to be shown. □

4.17. Proposition. *For some* $\varphi, \psi \in \mathcal{L}_0$*, the formula* $\mathbf{B}\neg\varphi \wedge \neg\mathbf{B}\bot \wedge \neg\mathbf{K}\neg\varphi \wedge [\texttt{revise}\ \varphi]\neg\mathbf{B}(\varphi \wedge \psi)$ *is satisfiable.*

Proof: Consider the model M defined in Example 4.1. For this model it holds that $\mathrm{M} \models \mathbf{B}\neg p \wedge \neg\mathbf{B}\bot \wedge \neg\mathbf{K}\neg p \wedge [\texttt{revise}\ p]\neg\mathbf{B}(p \wedge q)$. □

4.18. Corollary. *For every contraction operator* $\Leftrightarrow$ *that satisfies the AGM axiomatisation and the postulate of rehabilitation, it holds for all* $\varphi, \psi \in \mathcal{L}_0$ *and* $K \subseteq \mathcal{L}_0$ *that if* $\neg\varphi \in K, \bot \notin K, \varphi \wedge \psi \notin K^*_\varphi, \not\vdash_{cpl} \neg\varphi$ *then* $\varphi \in ((K^*_\varphi)^*_{\varphi\wedge\psi})^-_{\varphi\wedge\psi}$ *and* $\varphi \notin (K^*_{\varphi\wedge\psi})^-_{\varphi\wedge\psi}$*, where* $*$ *is the revision operator defined out of* $\Leftrightarrow$ *via the Levi-identity.*

4.19. Remark. Proposition 4.15 can be used to provide (counter)examples for all the non-validities that we previously encountered. Extending the antecedent of the implication given in Proposition 4.15 with the conjunct $[\texttt{revise}\ \varphi]\mathbf{B}(\neg\varphi\vee\neg\psi)$ provides a counterexample to Proposition 4.7(1) and DP1', whereas extending it with the conjunct $[\texttt{revise}\ \varphi]\neg\mathbf{B}(\neg\varphi \vee \neg\psi)$ yields a counterexample to Proposition 4.7(2) and DP2'. Extending the implication with $\mathbf{K}\neg(\varphi \wedge \psi)$ leads to a counterexample for N8'.

Proposition 4.15 shows that the approach based on historical awareness is to a considerable extent incompatible with the forgetful attitude formalised by postulate L4. One has to choose either for Lehmann's postulate L4 or for the postulate of rehabilitation, without having the possibility of adopting a middle course. In view of the area of application of our approach, viz. the formalisation of rational (information) agents that may adopt beliefs by default, the arguments in favour of the postulate of rehabilitation in our opinion outweigh those that support L4.

5. Discussion

In this paper we propose a formal framework to study iterated changes of belief in a modal (dynamic-epistemic-doxastic) context. We formalised expansions, contractions and revisions as actions, which, when performed by an agent, result in a change in its beliefs. The expressiveness of our (multi-modal) language allows for a concise representation of all kinds of postulates for (iterated) belief change. Furthermore, the formalisation of single-step belief changes is proved to be sound and complete with regard to the respective AGM axiomatisations. To interpret sequences of belief-changing actions we extend the models used for single-step belief changes, such that they represent belief systems, which consist not only of the agent's belief set but also contain a method for changing this set upon the execution of belief-changing actions. Execution of a belief-changing action has the direct effect that the agent's belief set is modified; as a side-effect a method for changing the modified set of beliefs is yielded. Several intuitively acceptable properties of iterated belief change are proved to be valid in our formalisation. Furthermore the validity of various recently proposed postulates for iterated belief revision is checked, both when interpreted as postulates for belief sets and as postulates for belief systems.

Acknowledgements Thanks are due to Wiebe van der Hoek and John-Jules Meyer for their helpful comments on a draft version of this paper. This research is partially supported by Esprit BRWG project No.8319 'ModelAge' and the Vrije Universiteit Amsterdam.

References

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

[2] C. Boutilier. Revision sequences and nested conditionals. In *Proceedings of IJCAI'93*, pages 519–525, 1993.

[3] A. Darwiche and J. Pearl. On the logic of iterated belief revision. In *Proceedings of the Fifth Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 5–23, Pacific Grove, CA, 1994. Morgan Kaufmann.

[4] M. Freund and D. Lehmann. Belief revision and rational inference. Technical Report TR 94-16, Leibnitz Center for Research in Computer Science, Institute of Computer Science, Hebrew University of Jerusalem, 1994.

[5] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. The MIT Press, Cambridge, Massachusetts and London, England, 1988.

[6] A. Grove. Two modelings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[7] D. Harel. Dynamic logic. In D.M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume 2, chapter 10, pages 497–604. D. Reidel, Dordrecht, 1984.

[8] J. Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, NY, 1962.

[9] S. Kraus and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174, 1988.

[10] D. Lehmann. Belief revision, revised. In *Proceedings of IJCAI'95*, 1995.

[11] I. Levi. Direct inference. *The Journal of Philosophy*, 74:5–29, 1977.

[12] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer. Actions that make you change your mind. In A. Laux and H. Wansing, editors, *Knowledge and Belief in Philosophy and Artificial Intelligence*, pages 105–146. Akademie-Verlag, 1995.

[13] B. van Linder, W. van der Hoek, and J.-J. Ch. Meyer. The dynamics of default reasoning (extended abstract). In C. Froidevaux and J. Kohlas, editors, *Symbolic and Quantitative Approaches to Reasoning and Uncertainty (Proceedings of ECSQARU'95)*, volume 861 of *Lecture Notes in Computer Science (subseries LNAI)*, pages 277–284. Springer-Verlag, 1995.

[14] A. Nayak, N. Foo, M. Pagnucco, and A. Sattar. Changing conditional beliefs unconditionally. Manuscript, 1995.

[15] K. Segerberg. Trying to understand the logic of theory change in terms of modal logic. In N. Foo, editor, *Record of the Workshop on Logic and Action*, 1994.

[16] 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 (KR'94)*, pages 619–629. Morgan Kaufmann, 1994.