1 Actions That Make You Change Your Mind

B. van Linder, W. van der Hoek and J.-J. Ch. Meyer[1]

ABSTRACT In this paper we study the dynamics of belief from an agent-oriented, semantics-based point of view. In a formal framework used to specify, and to reason about, rational agents, we define actions that model three well-known changes of belief, viz. *expansions*, *contractions*, and *revisions*. We treat these belief changes as fully fledged actions by defining both the *opportunity for* and the *result of* these actions, and the *ability* of agents to apply these belief-changing actions. In defining the result of the contraction action we introduce the concept of selection functions. These are special functions that select a set of states which is to be added to the set of doxastic alternatives of an agent, thereby contracting its set of beliefs. The action that models belief revisions is defined as the sequential composition of a contraction and an expansion. We show that these belief-changing actions are defined in an intuitively acceptable, reasonable way by proving that the AGM postulates for belief changes are validated. The ability of agents to apply belief-changing actions is defined in terms of their knowledge and belief. These definitions are such that actions that an agent is able to perform lead to desirable states of affairs. The resulting framework provides an intuitively acceptable and intelligible formalization of expansions, contractions and revisions as actions in an agent-oriented setting.

1.1 Introduction

The formalization of rational agents is a topic of continuing interest in Artificial Intelligence. Research on this subject has held the limelight ever since the pioneering work of Moore [Moore, 1980; Moore, 1985] in which knowledge and actions are considered. Over the years important contributions have been made on both *informational* attitudes like knowledge and belief [Halpern and Moses, 1992; Meyer and Hoek, 1995], and *motivational* attitudes[2] like commitments and obligations [Cohen and Levesque, 1990]. Recent developments include the work on agent-oriented programming [Shoham, 1993], the Belief-Desire-Intention architecture [Rao and Georgeff, 1991], and cognitive robotics [Lesperance *et al.*, 1995].

---

[1]Utrecht University, Department of Computer Science, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands. Email: bernd@cs.ruu.nl.

[2]The terms 'informational attitudes' and 'motivational attitudes' are due to Shoham & Cousins [Shoham and Cousins, 1994].

This paper is part of a series of papers in which we developed a *theorist* logic for rational agents, i.e., a logic that is used to *specify*, and to *reason about*, rational agents. We concentrate on informational attitudes and aspects of action, leaving motivational attitudes (for the moment) out of consideration. In the basic architecture the *knowledge*, *belief* and *abilities* of agents, as well as the *opportunities* for, and the *results* of, their actions are formalized [Hoek *et al.*, 1994a]. Subsequent papers extended this framework with the possibility to deal with nondeterministic actions [Hoek *et al.*, 1994b], epistemic tests [Linder *et al.*, 1994d], communication between agents [Linder *et al.*, 1994b], and reasoning by default [Linder *et al.*, 1994c]. The aim of this paper is a formalization of *belief-changing* actions. Belief-changing actions are intrinsically interesting from a philosophical point of view, but in addition they are very important for the formalization of rational agents that acquire information from multiple sources. For whenever some source provides reliable information that contradicts the information that an agent already has, the agent has to change its beliefs if it wants to incorporate this new information while keeping its set of beliefs consistent. In this paper we concentrate on three kinds of belief-changing actions; the formalization of information acquisition from multiple sources is dealt with elsewhere [Linder *et al.*, 1995]. The actions that we formalize are meant to model *rational changes of belief*. The probably best known and most prominent formal approach towards rational belief change is the so called AGM framework as proposed by Alchourrón, Gärdenfors and Makinson [Alchourrón *et al.*, 1985; Gärdenfors, 1988]. In this framework rationality postulates — we refer to these as the *AGM 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. For any of these changes of belief, a corresponding action is defined in our framework. As for any action, we define the states of affairs resulting from execution of belief-changing actions, conditions that need to be satisfied in order for agents to have the opportunity to perform these actions, and capacities that the agents must posses in order to be capable of performing these actions. The resulting definitions can be seen as providing *a dynamic logic of belief change*, in which belief changes are modelled as actions in such a way that all AGM postulates are validated. Several other modal approaches to belief change have been proposed, of which we briefly mention a few. The approaches of Fuhrmann [Fuhrmann, 1990] and Grove [Grove, 1988] are inspired by conditional rather than dynamic logic, and differ substantially from our approach, in particular with respect to the object language that is used. De Rijke proposes a highly expressive dynamic formalism, called dynamic modal logic, which can be used as a unifying framework to compare various dynamic approaches to belief change [Rijke, 1994]. The language and semantics of dynamic modal logic are completely different from ours. Furthermore, it is not straightforward to define the AGM belief changes in dynamic modal logic; the definitions proposed by De Rijke do not validate all of the AGM postulates. The formalism proposed by Segerberg [Segerberg, 1994], which is based on propositional dynamic logic, is in spirit very close to our approach. In effect though, it is different, the most notable difference being the fact that not all AGM postulates are validated in Segerberg's framework.

The rest of the paper is organized as follows. To sketch the context and the area of application of this research, we start in §1.2 with the (re)introduction of some of our ideas on knowledge, belief, abilities, opportunities and results in a context of multiple agents. In §1.3 we introduce actions that model expansions, revisions, and contractions. We define the states of affairs following these actions, conditions that need to be fulfilled for agents to have the opportunity to perform these actions, and (mental) capacities that agents should have to be capable of performing these actions. Furthermore we show that these actions satisfy (slightly adapted) versions of the AGM postulates. In §1.4 we round off. Proofs that we found too elaborate, trivial, or tedious to include here can be found elsewhere [Linder *et al.*, 1994a].

1.2 Informational attitudes and actions

At the informational level we consider both *knowledge and belief*. Formalizing these notions has been a subject of continuing research both in analytical philosophy and in AI [Halpern and Moses, 1992; Hintikka, 1962]. In representing knowledge and belief we follow, both from a syntactical and a semantic point of view, the approach common in epistemic and doxastic logic: the formula $\mathbf{K}_i\varphi$ denotes the fact that agent i knows φ, and $\mathbf{B}_i\varphi$ that agent i believes φ. For the semantics we use Kripke-style possible worlds models.

At the action level we consider *results*, *abilities* and *opportunities*. In defining the result of an action, we follow ideas of Von Wright [Wright, 1963], in which the state of affairs brought about by execution of the action is defined to be its result. An important aspect of any investigation of action is the relation that exists between ability and opportunity. In order to successfully complete an action, both the opportunity and the ability to perform the action are necessary. Although these notions are interconnected, they are surely not identical [Kenny, 1975]: the ability of agents comprises mental and physical powers, moral capacities, and human and physical possibility, whereas the opportunity to perform actions is best described by the notion of circumstantial possibility. A nice example that illustrates the difference between ability and opportunity is that of a lion in a zoo [Elgesem, 1993]: although the lion will (ideally) never have the opportunity to eat a zebra, it certainly has the ability to do so. We postulate that in order to make our formalization of rational agents, like for instance robots, as accurate and realistic as possible, abilities and opportunities need also be distinguished in AI environments. The abilities of agents are formalized via the $\mathbf{A}_i$ operator; the formula $\mathbf{A}_i\alpha$ denotes that agent i has the ability to do α. When using the definitions of opportunities and results as given above, the framework of (propositional) dynamic logic provides an excellent means to formalize these notions. Using events $\mathrm{do}_i(\alpha)$ to refer to the performance of the action α by the agent i, we consider the formulae $\langle \mathrm{do}_i(\alpha)\rangle\varphi$ and $[\mathrm{do}_i(\alpha)]\varphi$. In our deterministic framework, $\langle \mathrm{do}_i(\alpha)\rangle\varphi$ is the stronger of these formulae; it represents the fact that agent i has the opportunity to do α and that doing α results in φ being true. The formula $[\mathrm{do}_i(\alpha)]\varphi$ is noncommittal about the opportunity of the agent to do α, but states that, should the opportunity arise, only states of affairs satisfying φ would result. Besides the possibility to formalize both opportunities and results when using dynamic logic, another advantage lies in the compatibility of epistemic, doxastic and dynamic logic from a semantic point of view: a Kripke-style semantics can be used to provide meaning to epistemic, doxastic and dynamic notions.

Definition 1.2.1 Let a finite set $\mathcal{A} = \{1, \ldots, n\}$ of agents, and some denumerable sets Π of propositional symbols and At of atomic actions be given. The language $\mathcal{L}$ and the class of actions Ac are defined by mutual induction as follows.

1. $\mathcal{L}$ is the smallest superset of Π such that
   - if $\varphi, \psi \in \mathcal{L}$ then $\neg\varphi, \varphi \vee \psi \in \mathcal{L}$
   - if $i \in \mathcal{A}$, $\alpha \in Ac$ and $\varphi \in \mathcal{L}$ then $\mathbf{K}_i\varphi, \mathbf{B}_i\varphi, \langle \mathrm{do}_i(\alpha)\rangle\varphi, \mathbf{A}_i\alpha \in \mathcal{L}$
2. Ac is the smallest superset of At such that
   - if $\varphi \in \mathcal{L}$ then $\texttt{confirm}\ \varphi \in Ac$
   - if $\alpha_1 \in Ac$ and $\alpha_2 \in Ac$ then $\alpha_1; \alpha_2 \in Ac$
   - if $\varphi \in \mathcal{L}$ and $\alpha_1, \alpha_2 \in Ac$ then $\texttt{if}\ \varphi\ \texttt{then}\ \alpha_1\ \texttt{else}\ \alpha_2\ \texttt{fi} \in Ac$
   - if $\varphi \in \mathcal{L}$ and $\alpha_1 \in Ac$ then $\texttt{while}\ \varphi\ \texttt{do}\ \alpha_1\ \texttt{od} \in Ac$

The purely propositional fragment of $\mathcal{L}$ is denoted by $\mathcal{L}_0$. The constructs $\wedge$, $\rightarrow, \leftrightarrow, \mathbf{tt}, \mathbf{ff}$ and $[\mathrm{do}_i(\alpha)]\varphi$ are defined in the usual way. Other constructs are introduced by definitional abbreviation: `skip` is `confirm` $\mathbf{tt}$, α^0 is `skip`, and α^{n+1} is $\alpha; \alpha^n$.

The `confirm` action behaves essentially like the test actions in dynamic logic [Harel, 1984]. As such this action differs substantially from tests as they are looked upon by humans: these genuine tests are usually assumed to contribute to the information of the agent that performs the test [Linder *et al.*, 1994d], whereas by performing `confirm` φ it is just *confirmed* (verified, checked) that φ holds. The other actions in Ac denote respectively sequential composition, conditional composition, and repetitive composition; `skip` denotes the void action.

In the following definitions it is assumed that some set $\{\mathbf{0}, \mathbf{1}\}$ of truth values is given.

Definition 1.2.2 The class M of Kripke models contains tuples $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c}\rangle$ such that

- $\mathcal{S}$ is a set of possible worlds, or states.
- $\pi : \Pi \times \mathcal{S} \rightarrow \{\mathbf{0}, \mathbf{1}\}$ is a total function that assigns a truth value to propositional symbols in possible worlds.
- $\mathrm{R} : \mathcal{A} \rightarrow \wp(\mathcal{S} \times \mathcal{S})$ is a function that yields the epistemic accessibility relations for a given agent. This function is such that $\mathrm{R}(i)$ is an equivalence relation for all i. For reasons of practical convenience we define $[s]_{\mathrm{R}(i)}$ to be $\{s' \in \mathcal{S} \mid (s, s') \in \mathrm{R}(i)\}$.

- $\mathrm{B} : \mathcal{A} \times \mathcal{S} \to \wp(\mathcal{S})$ is a function that yields the set of doxastic alternatives for a given agent in a given state. To model the kind of belief that we like to model, it is demanded that for all agents i and for all possible worlds s and s' it holds that:
  - $\mathrm{B}(i, s) = \mathrm{B}(i, s')$ if $s' \in [s]_{\mathrm{R}(i)}$
  - $\mathrm{B}(i, s) \subseteq [s]_{\mathrm{R}(i)}$
- $\mathtt{r} : \mathcal{A} \times At \to \mathcal{S} \to \wp(\mathcal{S})$ is such that $\mathtt{r}(i, a)(s)$ yields the (possibly empty) state transition in s caused by the event $\mathrm{do}_i(a)$. This function is such that for all atomic actions a it holds that $|\mathtt{r}(i, a)(s)| \leq 1$ for all i and s, i.e., these events are *deterministic*.
- $\mathtt{c} : \mathcal{A} \times At \to \mathcal{S} \to \{\mathbf{0}, \mathbf{1}\}$ is the capability function such that $\mathtt{c}(i, a)(s)$ indicates whether the agent i is capable of performing the action a in s.

Using an equivalence relation to provide the semantics for the knowledge operator, results in the fairly standard and well accepted S5 modal system modelling knowledge. The common approach to defining the semantics of the belief operator consists of using a serial, transitive and euclidean relation to denote doxastic alternatives [Kraus and Lehmann, 1988], thereby ending up with a KD45 axiomatization for belief. The approach that we propose in Definition 1.2.2 differs from the more common one in that we use a (possibly empty) set instead of a (serial) accessibility relation to denote doxastic alternatives. The reason for using a set instead of a relation is a technical one: using sets, while essentially equivalent to using accessibility relations, allows for concise definitions of the semantics of belief-changing actions. The fact that we allow empty sets, whereas Kraus & Lehmann demand the accessibility relation to be serial, results in our notion of belief not validating the D-axiom $\neg(\mathbf{B}_i\varphi \wedge \mathbf{B}_i\neg\varphi)$. The reason for this is that the AGM approach towards belief changes, upon which we base our definitions of belief-changing actions, presupposes the existence of inconsistent belief sets: expansions and revisions may result in the agent having absurd beliefs. Apart from this difference, knowledge and belief are related to each other as in the system of Kraus & Lehmann.

Definition 1.2.3 Let $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ be some model from M. For φ a propositional symbol, a negation, or a disjunction, $\mathcal{M}, s \models \varphi$ is defined as usual. For the other clauses it is thus defined:

$$\begin{array}{ll}
\mathcal{M}, s \models \mathbf{K}_i\varphi & \Leftrightarrow \forall s' \in \mathcal{S}[(s,s') \in \mathrm{R}(i) \Rightarrow \mathcal{M}, s' \models \varphi] \\
\mathcal{M}, s \models \mathbf{B}_i\varphi & \Leftrightarrow \forall s' \in \mathcal{S}[s' \in \mathrm{B}(i,s) \Rightarrow \mathcal{M}, s' \models \varphi] \\
\mathcal{M}, s \models \langle \mathrm{do}_i(\alpha)\rangle\varphi & \Leftrightarrow \exists \mathcal{M}', s'[\mathcal{M}', s' \in \mathtt{r}(i,\alpha)(\mathcal{M},s) \,\&\, \mathcal{M}', s' \models \varphi] \\
\mathcal{M}, s \models \mathbf{A}_i\alpha & \Leftrightarrow \mathtt{c}(i,\alpha)(\mathcal{M},s) = \mathbf{1}
\end{array}$$

where r and c are extended as follows:

$$\begin{array}{ll}
\mathtt{r}(i,a)(\mathcal{M},s) & = \mathcal{M}, \mathtt{r}(i,a)(s) \\
\mathtt{r}(i,\mathtt{confirm}\ \varphi)(\mathcal{M},s) & = \{(\mathcal{M},s)\} \text{ if } \mathcal{M},s \models \varphi \text{ and } \emptyset \text{ otherwise} \\
\mathtt{r}(i,\alpha_1;\alpha_2)(\mathcal{M},s) & = \mathtt{r}(i,\alpha_2)(\mathtt{r}(i,\alpha_1)(\mathcal{M},s)) \\
\mathtt{r}(i,\mathtt{if}\ \varphi\ \mathtt{then}\ \alpha_1 & = \mathtt{r}(i,\alpha_1)(\mathcal{M},s) \text{ if } \mathcal{M},s \models \varphi \text{ and} \\
\quad \mathtt{else}\ \alpha_2\ \mathtt{fi})(\mathcal{M},s) & \quad \mathtt{r}(i,\alpha_2)(\mathcal{M},s) \text{ otherwise} \\
\mathtt{r}(i,\mathtt{while}\ \varphi & = \{(\mathcal{M}',s') \mid \exists k \in \mathbb{N} \exists \mathcal{M}_0, s_0 \ldots \exists \mathcal{M}_k, s_k \\
\quad \mathtt{do}\ \alpha_1\ \mathtt{od})(\mathcal{M},s) & \quad [\mathcal{M}_0, s_0 = \mathcal{M}, s \,\&\, \mathcal{M}_k, s_k = \mathcal{M}', s' \,\&\, \forall j < k \\
& \quad [\mathcal{M}_{j+1}, s_{j+1} = \mathtt{r}(i,\mathtt{confirm}\ \varphi;\alpha_1)(\mathcal{M}_j, s_j)] \\
& \quad \& \, \mathcal{M}', s' \models \neg\varphi]\}
\end{array}$$

where $\mathtt{r}(i,\alpha)(\emptyset) = \emptyset$

and

$$\begin{array}{ll}
\mathtt{c}(i,a)(\mathcal{M},s) & = \mathtt{c}(i,a)(s) \\
\mathtt{c}(i,\mathtt{confirm}\ \varphi)(\mathcal{M},s) & = \mathbf{1} \text{ if } \mathcal{M},s \models \varphi \text{ and } \mathbf{0} \text{ otherwise} \\
\mathtt{c}(i,\alpha_1;\alpha_2)(\mathcal{M},s) & = \mathtt{c}(i,\alpha_1)(\mathcal{M},s) \,\&\, \mathtt{c}(i,\alpha_2)(\mathtt{r}(i,\alpha_1)(\mathcal{M},s)) \\
\mathtt{c}(i,\mathtt{if}\ \varphi\ \mathtt{then}\ \alpha_1 & = \mathtt{c}(i,\mathtt{confirm}\ \varphi;\alpha_1)(\mathcal{M},s) \text{ or} \\
\quad \mathtt{else}\ \alpha_2\ \mathtt{fi})(\mathcal{M},s) & \quad \mathtt{c}(i,\mathtt{confirm}\ \neg\varphi;\alpha_2)(\mathcal{M},s) \\
\mathtt{c}(i,\mathtt{while}\ \varphi & = \mathbf{1} \text{ if } \exists k \in \mathbb{N}[\mathtt{c}(i,(\mathtt{confirm}\ \varphi;\alpha_1)^k; \\
\quad \mathtt{do}\ \alpha_1\ \mathtt{od})(\mathcal{M},s) & \quad \mathtt{confirm}\ \neg\varphi)(\mathcal{M},s) = \mathbf{1}] \\
& \quad \text{and } \mathbf{0} \text{ otherwise}
\end{array}$$

where $\mathtt{c}(i,\alpha)(\emptyset) = \mathbf{1}$

Validity on the class M of models is defined as usual.

The notion of actions as considered in Definition 1.2.3 generalizes that of state-transformers as it is typical for dynamic logic [Harel, 1984], and allows also for actions that transform pairs (Model, State). The reason for this generalization lies in the fact that we also account for non-standard actions like 'to observe' [Linder *et al.*, 1994d], and 'to inform' [Linder *et al.*, 1994b] in addition to the mundane actions of dynamic logic, and these non-standard actions transform models rather than states. Also belief-changing actions are naturally interpreted as transforming models rather than states; we elaborate on this in § 1.3. Note that as a consequence of this generalized notion of actions, the functions r and c should in fact be retyped for non-atomic actions: instead of taking elements from $\mathcal{S}$ as their third argument they take elements from $M \times \mathcal{S}$. To prevent our notation from becoming to baroque, we leave out this retyping.

With regard to the abilities of agents, the motivation for the choices made in Definition 1.2.3 is the following. The definition of $\mathtt{c}(i, \mathtt{confirm}\ \varphi)$ expresses that an agent is able to get confirmation for a formula φ iff φ holds. Note that the definitions of $\mathtt{r}(i, \mathtt{confirm}\ \varphi)$ and $\mathtt{c}(i, \mathtt{confirm}\ \varphi)$ imply that in circumstances such that φ holds, agents have both the opportunity and the ability to confirm φ. An agent is capable of performing a sequential composition $\alpha_1; \alpha_2$ iff it is capable of performing α_1 (now), and it is capable of executing α_2 after it has performed α_1. An agent is capable of performing a conditional composition, if either it is able to get confirmation for the condition and thereafter perform the then-part, or it is able to confirm the negation of the condition and perform the else-part afterwards. An agent is capable of performing a repetitive composition $\mathtt{while}\ \varphi\ \mathtt{do}\ \alpha_1\ \mathtt{od}$ iff it is able to perform the action $(\mathtt{confirm}\ \varphi; \alpha_1)^k; \mathtt{confirm}\ \neg\varphi$ for some natural number k.

The following example gives an idea of the expressive power of the framework defined above.

Example 1.2.4 (Egocentric and altruistic agents) The egocentric action is described as follows:

> 'If agent i knows that it will believe that it feels better after helping its neighbor and it knows that it is able to help, it will do so, otherwise it will do nothing'

It is possible to express in our system the fact that if agent i believes to be feeling well before performing the egocentric action, it will believe to feel either well or even better after performing it. Using some obvious abbreviations the situation sketched above amounts to the following validity:

$$\begin{aligned}\models \mathbf{B}_i w_i \to [\mathrm{do}_i(&\mathtt{if}\ \ \mathbf{K}_i[\mathrm{do}_i(h)](\mathbf{B}_i b_i) \wedge \mathbf{K}_i \mathbf{A}_i h \\ &\mathtt{then}\ h\ \mathtt{else}\ \mathtt{skip})](\mathbf{B}_i w_i \vee \mathbf{B}_i b_i)\end{aligned}$$

It is also possible to formalize an altruistic action: agent i will help if it knows that this will make (some other) agent j feel better. Now it is possible to express the fact that if agent j is feeling good right now, it will feel good or even better after i has performed the altruistic action:

$$\models w_j \rightarrow [\text{do}_i(\texttt{if } \mathbf{K}_i[\text{do}_i(h)]b_j \texttt{ then } h \texttt{ else skip})](w_j \vee b_j)$$

1.3 Actions that change one's mind

Our approach towards belief changes in the agent-oriented, semantics-based framework of §1.2 is based on the idea that belief changes are brought about by actions that the agents may perform, i.e., by performing belief-changing actions, agents expand, contract and revise their beliefs. We consider the knowledge of agents concerning propositional formulae to be immune for changes: it persists under the execution of belief-changing actions. This non-defeasibility of knowledge constitutes one of two — the other one being veridicality — major differences between the two informational attitudes dealt with in our framework. Due to its non-defeasibility, knowledge in our framework plays the part that is played by the theorems of the underlying logic in the AGM framework, viz. that of *unassailable beliefs*. As such, we will on occasion refer to knowledge as providing the *principles* of agents, the information that the agents will never part from.

From a syntactical point of view, the class of actions Ac is extended with three new, belief-changing, actions.

Definition 1.3.1 The class Ac of actions (and hence the language $\mathcal{L}$) as defined in 1.2.1 is extended as follows:

if $\varphi \in \mathcal{L}_0$ then $\texttt{expand } \varphi, \texttt{contract } \varphi, \texttt{revise } \varphi \in Ac$

The main reason underlying the restriction to propositional formulae in Definition 1.3.1 is the fact that changes of belief concerning doxastic or epistemic formulae are not well understood. It is for instance not at all clear what it means to revise the beliefs of some agent i with the formula $p \wedge \neg\mathbf{B}_i p$: does or doesn't the agent believe $p \wedge \neg\mathbf{B}_i p$ after its beliefs are revised with this formula?[3] Belief changes with propositional formulae do not suffer from problems like these. From a semantic point of view the property that makes belief changes with propositional formulae understandable and more or less predictable is the fact that the truth

[3] The natural language variant 'p, and i does not believe p' of the formula $p \wedge \neg\mathbf{B}_i p$ is considered by Thijsse ([Thijsse, 1992], pp. 131-132) to be a typical example of a non-contradictory sentence but a contradictory utterance. This implies that although the sentence in itself is consistent, it is not consistent to believe the sentence.

value of these formulae in a state of a model depends on the valuation for that state only.

Proposition 1.3.2 *Let* $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ *be some Kripke model with* $s \in \mathcal{S}$*, and let* $\mathcal{M}' = \langle \mathcal{S}', \pi', \mathrm{R}', \mathrm{B}', \mathtt{r}', \mathtt{c}' \rangle$ *be some Kripke model with* $s' \in \mathcal{S}'$*. Then it holds that:*

$$\forall p \in \Pi[\pi(p, s) = \pi'(p, s')] \Rightarrow \forall \psi \in \mathcal{L}_0[\mathcal{M}, s \models \psi \Leftrightarrow \mathcal{M}', s' \models \psi]$$

In fact it is not necessary to make a restriction to propositional formulae; it would suffice to restrict oneself to formulae whose truth in the state of a model is not dependent on the information fluents — represented by the epistemic accessibility function R and the doxastic accessibility function B — of the model. In a straightforward extension of the framework presented here one could allow agents to change their beliefs on their own — or other agents' — abilities and opportunities, or on the results of the agents' actions. However, in order to expose our ideas on belief-changing actions in their purest form, we have decided to follow common practice in the literature on belief changes and restrict ourselves to propositional formulae.

The semantics for belief-changing actions as presented in this paper is based on ideas that generalize those underlying our definition of informative actions like epistemic tests [Linder *et al.*, 1994d] and communication actions [Linder *et al.*, 1994b]. Basically, the idea is that changes in the beliefs of an agent correspond to including and dropping certain items of information: in the case of an expansion (new) information is included, contractions lead to the dropping of information, and revisions drop some items of information and include others. In our modal framework the information of an agent is formalized through its set of epistemic and doxastic alternatives. Hence a natural implementation of the dropping and inclusion of information is given by the inclusion and dropping of doxastic alternatives, where the dropping of information corresponds to the inclusion of new worlds into the set of doxastic alternatives of the agent and the inclusion of (new) information corresponds to a restriction of this set. Given this intuitive meaning of the semantics of belief-changing actions, it is obvious that these actions cause transitions between models, thus generalizing the usual actions from dynamic logic that cause inter-state transitions within a model.

Convention 1.3.3 In the rest of this paper we follow the convention that whenever some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ is clear from the context, $[\![\varphi]\!]$ denotes the set of states that satisfy φ, i.e., $[\![\varphi]\!] = \{s \in \mathcal{S} \mid \mathcal{M}, s \models \varphi\}$. The relation $\vdash_{cpl} \subseteq \wp(\mathcal{L}_0) \times \mathcal{L}_0$ is the derivability relation of classical propositional logic. The function $\mathrm{Th} : \wp(\mathcal{L}_0) \to \wp(\mathcal{L}_0)$ that yields for every set Φ of propositional formula the set $\{\varphi \in \mathcal{L}_0 \mid \Phi \vdash_{cpl} \varphi\}$ is the deductive closure operator associated with the derivability relation $\vdash_{cpl}$. Note that both $\vdash_{cpl}$ and Th work strictly on a propositional level; expressions like $\{\mathbf{B}_i p\} \vdash_{cpl} p \vee \neg p$ and $\mathrm{Th}(\mathbf{B}_i p \vee \neg \mathbf{B}_i p)$ are not well-defined. The indexes i and j, possibly marked, always refer to agents.

In the following, we successively define the semantics of the three belief-changing actions, and study their specific properties. We impose the following properties on *all* of them:

Definition 1.3.4 We distinguish the following properties of actions α, where $\chi \in \mathcal{L}$.

- $\models \langle \mathrm{do}_i(\alpha) \rangle \mathbf{tt}$ *realizability*
- $\models \langle \mathrm{do}_i(\alpha) \rangle \chi \to [\mathrm{do}_i(\alpha)] \chi$ *determinism*
- $\models \langle \mathrm{do}_i(\alpha; \alpha) \rangle \chi \leftrightarrow \langle \mathrm{do}_i(\alpha) \rangle \chi$ *idempotence*

Realizability of an action implies that agents have the opportunity to perform the action regardless of circumstances, determinism of an action means that performing the action results in a unique state of affairs, and idempotence of an action implies that performing the action an arbitrary number of times has the same effect as performing the action just once. We say that $\Delta \in \{\mathtt{expand}, \mathtt{contract}, \mathtt{revise}\}$ satisfies any of the properties of Definition 1.3.4 if the action $\Delta\varphi$ satisfies that property for all $\varphi \in \mathcal{L}_0$.

1.3.1 The `expand` action

Informally, a belief expansion is an action that leads to a state of affairs in which some formula is included in the set of beliefs of an agent. In our framework uncertainties of agents are formalized through the different doxastic alternatives that the agent has: if an agent believes neither φ nor $\neg\varphi$ then it considers both doxastic alternatives supporting φ and doxastic alternatives supporting $\neg\varphi$ possible. Expanding the beliefs of the agent with φ may then be implemented by declaring all alternatives supporting $\neg\varphi$ to be 'doxastically impossible', i.e., on the ground of its beliefs the agent no longer considers these alternatives to be possible. Hence the *expansion* of the belief set of an agent can be modelled through a *restriction* of its set of doxastic alternatives. The definition of the function $\mathtt{r}$ for expansions is a direct formalization of these intuitive ideas: if some agent i performs an expansion with some formula φ in a world s in the model $\mathcal{M}$, the result of this is that afterwards i has restricted its set of doxastic alternatives to those states that satisfy φ (even if there are no such states).

Definition 1.3.5 Let some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ with $s \in \mathcal{S}$, and $\varphi \in \mathcal{L}_0$ be given. We define:

$$\begin{array}{l} \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}, s) = \mathcal{M}', s \text{ where} \\ \quad \mathcal{M}' = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}', \mathtt{r}, \mathtt{c} \rangle \text{ with} \\ \quad\quad \mathrm{B}'(i', s') = \mathrm{B}(i', s') \text{ if } i' \neq i \text{ or } s' \notin [s]_{\mathrm{R}(i)} \\ \quad\quad \mathrm{B}'(i, s') = \mathrm{B}(i, s') \cap [\![\varphi]\!] \text{ if } s' \in [s]_{\mathrm{R}(i)} \end{array}$$

Definition 1.3.5 provides for an intuitively acceptable formalization of belief expansions as can be seen in the following proposition.

Proposition 1.3.6 *For all $\varphi, \psi \in \mathcal{L}_0$ we have:*

- $\models [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\varphi$
- $\models \mathbf{B}_i\psi \rightarrow [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi$
- $\models \mathbf{B}_i\varphi \rightarrow (\mathbf{B}_i\psi \leftrightarrow [\mathrm{do}_i(\mathtt{expand}\ \varphi)]\mathbf{B}_i\psi)$

The first clause of Proposition 1.3.6 states that an expansion with some formula results in the formula being believed. The second clause states that beliefs are persistent under expansions. In this clause the restriction to *propositional* formulae ψ is in general necessary. For consider a situation in which an agent does not believe φ and by negative introspection believes that it does not believe φ. After expanding its beliefs with φ, the agent believes φ and, assuming that the resulting belief set is not the absurd one, it no longer believes that it does not believe φ. Hence not all beliefs of the agent persist in situations like these. Note that the first two clauses combined indicate that our definition of belief, and in particular the fact that we allow absurd belief sets, is a good one when dealing with expansions. For an expansion with some formula φ in a situation in which $\neg\varphi$ is already believed, results in the agent believing both φ and $\neg\varphi$ and hence having inconsistent beliefs. The third clause states that in situations where some formula is already believed, nothing is changed as the result of an expansion with that formula. This latter property is suggested by the *criterion of informational economy* [Gärdenfors, 1988], which states that since information is in general not gratuitous, unnecessary losses of information are to be avoided. The validities expressed in Proposition 1.3.6 can be seen as the representation in $\mathcal{L}$ of the second, third and fourth of the AGM postulates for expansions. The other AGM postulates are not as neatly expressed in $\mathcal{L}$; nevertheless in §1.3.2 we show that all of the postulates for expansion are in fact validated.

Besides the properties given above, some other properties — particularly dealing with multiple agents — can be shown to hold in our formalization.

Proposition 1.3.7 *For all $\varphi, \psi \in \mathcal{L}_0$ and $\chi \in \mathcal{L}$ we have:*

- $\models \mathbf{K}_j \mathbf{B}_i \psi \to \mathbf{K}_j [\mathrm{do}_i(\texttt{expand}\ \varphi)] \mathbf{K}_j \mathbf{B}_i \psi$
- $\models \mathbf{B}_j \mathbf{B}_i \psi \to \mathbf{B}_j [\mathrm{do}_i(\texttt{expand}\ \varphi)] \mathbf{B}_j \mathbf{B}_i \psi$
- $\models \mathbf{B}_i \varphi \to (\chi \leftrightarrow [\mathrm{do}_i(\texttt{expand}\ \varphi)] \chi)$
- `expand` *satisfies realizability, determinism and idempotence*

The first clause of Proposition 1.3.7 states that the knowledge of another agent — the watching agent — on the beliefs of the agent performing an expansion is known (by the watching agent) to persist; the second clause states the same for the beliefs of the watching agent. The third clause formalizes the idea that expansions with formulae that are already believed, cause an universally minimal change: nothing in the universe changes as the result of such an expansion. The last clause has already been phrased in the introduction to this section.

It turns out that expansions as formalized in Definition 1.3.5 can be completely characterized as follows.

Proposition 1.3.8 *For all $\varphi, \psi \in \mathcal{L}_0$ we have:*

- $\models [\mathrm{do}_i(\texttt{expand}\ \varphi)] \mathbf{B}_i \psi \leftrightarrow \mathbf{B}_i(\varphi \to \psi)$

Proposition 1.3.8 states that some (propositional) formula ψ is believed after an expansion with φ if and only if the agent believes that φ implies ψ beforehand. As a special case of Proposition 1.3.8 we can prove that an expansion with some formula results in the agent having absurd beliefs if and only if the agent believes the negation of the formula beforehand.

Corollary 1.3.9 *For all $\varphi \in \mathcal{L}_0$ we have:*

- $\models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\mathbf{ff} \leftrightarrow \mathbf{B}_i\neg\varphi$

1.3.2 Expansions and the AGM postulates

As for contractions and revisions, the AGM postulates are by now the standard ones to describe rational belief expansions [Alchourrón *et al.*, 1985; Gärdenfors, 1988]. These postulates describe how changes in the belief set of an agent should work out. In the AGM framework, belief sets are defined as follows.

Definition 1.3.10 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 (unique) absurd belief set, consisting of all formulae from $\mathcal{L}_0$, is denoted by $K_\perp$.

In the following definition, K and H denote arbitrary AGM belief sets, φ denotes some propositional formula, and the expansion of K with φ is denoted by K^+_φ.

Definition 1.3.11 The AGM postulates for belief expansion:

(G^+1) K^+_φ is an AGM belief set.
(G^+2) $\varphi \in K^+_\varphi$.
(G^+3) $K \subseteq K^+_\varphi$.
(G^+4) If $\varphi \in K$ then $K^+_\varphi = K$.
(G^+5) If $K \subseteq H$ then $K^+_\varphi \subseteq H^+_\varphi$.
(G^+6) For all K, and all φ, K^+_φ is the smallest set that satisfies G^+1 – G^+5.

It turns out that our `expand` action can be seen as providing a belief expansion in the sense of the AGM postulates. To formulate the AGM postulates in our framework we introduce our own kind of belief sets. These belief sets are model-based and indexed with a particular agent. Furthermore the notion of knowledge sets, as providing the principles or prejudices that the agent will never part from, is defined below.

Definition 1.3.12 Let $\mathcal{M}$ be some Kripke model with $s \in \mathcal{M}$. The *belief set* of agent i in $\mathcal{M}, s$, notation $\mathcal{B}(i, \mathcal{M}, s)$, is defined by:

- $\mathcal{B}(i, \mathcal{M}, s) = \{\varphi \in \mathcal{L}_0 \mid \mathcal{M}, s \models \mathbf{B}_i\varphi\}$

The *knowledge set* of agent i in $\mathcal{M}, s$, notation $\mathcal{K}(i, \mathcal{M}, s)$, is defined by:

- $\mathcal{K}(i, \mathcal{M}, s) = \{\varphi \in \mathcal{L}_0 \mid \mathcal{M}, s \models \mathbf{K}_i\varphi\}$

The *expansion* of $\mathcal{B}(i, \mathcal{M}, s)$ with a formula $\varphi \in \mathcal{L}_0$, notation $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$, is defined by:

- $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \{\psi \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{expand}\ \varphi)]\mathbf{B}_i\psi\}$

The unique absurd belief set $\mathcal{B}_\perp$ is defined to be $\mathcal{L}_0$.

Theorem 1.3.13 *Let $\mathcal{M}$ and $\mathcal{M}'$ be Kripke models, with s in $\mathcal{M}$ and s' in $\mathcal{M}'$. The following is true for all $\varphi \in \mathcal{L}_0$.*

(B^+1) $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ *is an AGM belief set.*
(B^+2) $\varphi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
(B^+3) $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
(B^+4) *If* $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.
(B^+5) *If* $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}', s')$, *then* $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}', s')$.
(B^+6) $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ *is the smallest set that satisfies B^+1 – B^+5 as given above.*

Proof: Let $\mathcal{M}$ be a model with state s, and let $\varphi \in \mathcal{L}_0$ be arbitrary. Let $\mathcal{M}''$ be such that $\mathcal{M}'', s = \mathtt{r}(i, \texttt{expand}\ \varphi)(\mathcal{M}, s)$.

- (B^+1) This postulate follows straightforwardly from the definition of $\mathtt{r}$ for the `expand` action. If $\mathcal{M}''$ is such that $\mathrm{B}''(i, s) = \emptyset$ then $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathcal{B}_\perp$, and otherwise $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is consistent and deductively closed by definition of $\models$ for belief formulae.
- (B^+2) Since φ is in $\mathcal{L}_0$, it follows by definition of $\mathtt{r}$ for `expand` that $\mathcal{M}'', s' \models \varphi$ for all $s' \in \mathrm{B}''(i, s)$. Hence $\mathcal{M}'', s \models \mathbf{B}_i\varphi$ and thus $\varphi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
- (B^+3) Let $\psi \in \mathcal{B}(i, \mathcal{M}, s)$. Then $\mathcal{M}, s' \models \psi$ for all $\psi \in \mathrm{B}(i, s)$. Now $\mathrm{B}''(i, s) \subseteq \mathrm{B}(i, s)$ and since ψ is propositional, it follows that $\mathcal{M}'', s' \models \psi$ for all $s' \in \mathrm{B}''(i, s)$. Hence $\mathcal{M}'', s \models \mathbf{B}_i\psi$ and thus $\psi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
- (B^+4) Suppose $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$. Then $\mathcal{M}, s' \models \varphi$ for all $s' \in \mathrm{B}(i, s)$. Since φ is propositional, it follows that $\mathrm{B}(i, s) \cap [\![\varphi]\!] = \mathrm{B}(i, s)$ and hence $\mathrm{B}''(i, s) = \mathrm{B}(i, s)$. Then $\mathcal{M}'' = \mathcal{M}$, and $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}'', s) = \mathcal{B}(i, \mathcal{M}, s)$.
- (B^+5) The proof that this postulate is validated is most easily given as a direct consequence of Proposition 1.3.14. For if $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}', s')$ for some model $\mathcal{M}'$ with state s', then $\mathrm{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup$

$\{\varphi\}) \subseteq \text{Th}(\mathcal{B}(i, \mathcal{M}', s') \cup \{\varphi\})$ and $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}', s')$. Since the proof of Proposition 1.3.14 does not depend on B^+5, this postulate is validated.

- (B^+6) From B^+2, B^+3 and the fact that belief sets are deductively closed, it follows that $\text{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\}) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$. From Proposition 1.3.14, the proof of which does not depend on B^+6, it follows that $\mathcal{B}^+_\varphi(i, \mathcal{M}, s)$ is indeed the smallest set that satisfies B^+1 through B^+5. □

It is shown in the AGM framework that the postulates formulated in Definition 1.3.11 completely determine expansions, i.e., whereas the postulates for contraction and revision leave some degrees of freedom in defining contractions and revisions, G^+1 through G^+5 uniquely define expansions. Proposition 1.3.14 — a rephrasing of Proposition 1.3.8 — states that the same holds in our framework, and furthermore, the unique definition of expansions that we end up with is identical to the one given for the AGM framework.

Proposition 1.3.14 *For all models $\mathcal{M}$ with states s, and for all $\varphi \in \mathcal{L}_0$:*

- $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \text{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\})$

Proof: We prove that the two sets are equal by proving that each set is a subset of the other one. So let $\mathcal{M}$ be some model with state s and let $\varphi \in \mathcal{L}_0$ be arbitrary.

'$\supseteq$' This is shown by the argument given in the proof of B^+6: from B^+2, B^+3 and the fact that beliefs sets are deductively closed, it follows that $\text{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\}) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.

'$\subseteq$' Suppose that $\psi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$. If $\mathcal{M}', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}, s)$ then $\mathcal{M}', s' \models \psi$ for all $s' \in \text{B}'(i, s)$. Since ψ is in $\mathcal{L}_0$, and since $\text{B}'(i, s) = \text{B}(i, s) \cap [\![\varphi]\!]$, it follows that $\mathcal{M}, s' \models \psi$ for all $s' \in \text{B}(i, s)$ such that $\mathcal{M}, s' \models \varphi$. Hence $\mathcal{M}, s' \models \varphi \to \psi$ for all $s' \in \text{B}(i, s)$. Then $\varphi \to \psi \in \mathcal{B}(i, \mathcal{M}, s)$ and $\psi \in \text{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\})$. Since ψ is arbitrary, it follows that $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \text{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\})$. □

Note that it is of no use to try and consider the properties formalized in Proposition 1.3.7 in terms of the AGM framework. For the latter framework deals with the belief set of a single — implicit — agent only, and can therefore for instance not express the knowledge of (other) agents.

1.3.3 The `contract` action

A belief contraction is the change of belief through which in general some formula that is believed beforehand is no longer believed afterwards. As such, apparent beliefs that an agent has are turned into doubts as the result of a contraction. In terms of our framework, this change of belief may be implemented by *extending* the set of doxastic alternatives of an agent in order to encompass at least one state not satisfying the formula that is to be contracted. Consider for example the situation of an agent i that believes p, i.e., p holds in all its doxastic alternatives. When contracting p from the belief set of the agent, some $\neg p$-worlds are added to the set of doxastic alternatives of the agent. In order to end up with well-defined Kripke models, these worlds that are to be added, need to be in the set of epistemic alternatives of s. For in the Kripke models defined in 1.2.2, the set of doxastic alternatives for a given agent in a given state is contained in its set of epistemic alternatives in that state. Thus the worlds that are to be added to the set of doxastic alternatives of the agent are elements of the set of epistemic alternatives not supporting p.

The problem with defining contractions in this way, is that it is not straightforward as to decide which worlds need to be added. From the basic idea that knowledge — acting as the principles of agents — provides some sort of lower bound of the belief set of an agent, it is clear that in the case of a contraction with φ some states need to be added that are elements of the set of epistemic alternatives of the agent and do not support φ, but it is not clear exactly *which* elements of this set need to be chosen.

The approach that we propose to solve this problem is based on the use of so called *selection functions*. These are functions that (whenever possible) select a subset of the set of epistemic alternatives in such a way that the resulting `contract` action behaves rationally.

Definition 1.3.15 Let some model $\mathcal{M}$ be given. A function $\sigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ is a *selection function* for $\mathcal{M}$ if and only if it meets the following constraints for all $s, s' \in \mathcal{S}$ and $\varphi, \psi \in \mathcal{L}_0$.

Σ0. $\sigma(i, s, \varphi) = \sigma(i, s', \varphi)$ if $s' \in [s]_{\mathrm{R}(i)}$

Σ1. $\sigma(i, s, \varphi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$

Σ2. $\sigma(i, s, \varphi) \subseteq \mathrm{B}(i, s)$ if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$

Σ3. $\sigma(i, s, \varphi) = \emptyset$ iff $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] = \emptyset$

Σ4. if $[s]_{\mathrm{R}(i)} \cap [\![\varphi]\!] = [s]_{\mathrm{R}(i)} \cap [\![\psi]\!]$ then $\sigma(i, s, \varphi) = \sigma(i, s, \psi)$
Σ5. $\sigma(i, s, \varphi \wedge \psi) \subseteq \sigma(i, s, \varphi) \cup \sigma(i, s, \psi)$
Σ6. if $\sigma(i, s, \varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$ then $\sigma(i, s, \varphi) \subseteq \sigma(i, s, \varphi \wedge \psi)$

The first two of the demands given in Definition 1.3.15 ensure that applications of the `contract` action result in well-defined Kripke models: in all the states of its epistemic equivalence class the agent holds the same beliefs (Σ1), and the agent's set of beliefs still encompasses its knowledge after a contraction (Σ2). Demand Σ2 furthermore ensures that no worlds are added that are superfluous in the sense that they do not invalidate the formula that is to be contracted. Demands Σ2 to Σ6 enforce our notion of contraction to validate the AGM postulates (in Proposition 1.3.22 we elaborate on the relation between these demands and the postulates for contraction). Note that the expressions on the right-hand side of Σ3 and on the left-hand side of Σ4 are equivalently phrased as $\mathcal{M}, s \models \mathbf{K}_i\varphi$ and $\mathcal{M}, s \models \mathbf{K}_i(\varphi \leftrightarrow \psi)$, respectively.

The definition of `r` for the `contract` action is based on the use of selection functions: a contraction is performed by adding to the set of doxastic alternatives of the agent exactly those worlds that are picked out by the selection function.

Definition 1.3.16 Let some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ with $s \in \mathcal{S}$ and $\varphi \in \mathcal{L}_0$ be given. Furthermore, let σ be an arbitrary but fixed selection function for $\mathcal{M}$. We define:

$\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s) = \mathcal{M}', s$ where
$\mathcal{M}' = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}', \mathtt{r}, \mathtt{c} \rangle$ with
$\mathrm{B}'(i', s') = \mathrm{B}(i', s')$ if $i' \neq i$ or $s' \notin [s]_{\mathrm{R}(i)}$
$\mathrm{B}'(i, s') = \mathrm{B}(i, s') \cup \sigma(i, s, \varphi)$ for all $s' \in [s]_{\mathrm{R}(i)}$

Using selection functions to define the semantics for the `contract` action indeed results in an acceptable formalization of belief contraction, as can be seen in the following proposition.

Proposition 1.3.17 *For all* $\varphi, \psi, \vartheta \in \mathcal{L}_0$ *we have:*

- $\models [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\psi \to \mathbf{B}_i\psi$
- $\models \neg\mathbf{B}_i\varphi \to ([\mathrm{do}_i(\mathtt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i\psi)$
- $\models \neg\mathbf{K}_i\varphi \to [\mathrm{do}_i(\mathtt{contract}\ \varphi)]\neg\mathbf{B}_i\varphi$
- $\models \mathbf{B}_i\varphi \to (\mathbf{B}_i\psi \to [\mathrm{do}_i(\mathtt{contract}\ \varphi; \mathtt{expand}\ \varphi)]\mathbf{B}_i\psi)$

- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\texttt{contract}\ \psi)]\mathbf{B}_i\vartheta)$
- $\models ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\vartheta \wedge [\mathrm{do}_i(\texttt{contract}\ \psi)]\mathbf{B}_i\vartheta) \rightarrow$
  $[\mathrm{do}_i(\texttt{contract}\ \varphi \wedge \psi)]\mathbf{B}_i\vartheta$
- $\models [\mathrm{do}_i(\texttt{contract}\ \varphi \wedge \psi)]\neg\mathbf{B}_i\varphi \rightarrow$
  $([\mathrm{do}_i(\texttt{contract}\ \varphi \wedge \psi)]\mathbf{B}_i\vartheta \rightarrow [\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\vartheta)$

The first clause of Proposition 1.3.17 states that after a contraction an agent believes at most the formulae that it believed before the contraction. The second clause states that in situations in which an agent does not believe φ, nothing changes as the result of contracting φ. Again this property reflects the criterion of informational economy. The third clause states that a contraction with a contractable formula, this is a formula not belonging to the agent's principles, results in the agent not believing the contracted formula. The fourth clause states that whenever an agent believes a formula, all beliefs in its original belief set are recovered after a contraction with that formula followed by an expansion with the same formula. The fifth clause states that contractions with formulae that are known to be equivalent, result in identical belief sets. The sixth clause formalizes the idea that all formulae that are believed both after a contraction with φ and after a contraction with ψ are believed after a contraction with $\varphi \wedge \psi$. Clause 7 states that if a contraction with $\varphi \wedge \psi$ results in φ not being believed, then in order to contract φ no more formulae need to be removed than those that were removed in order to contract $\varphi \wedge \psi$. This last clause is related to the property of *minimal change* for contractions. The validities given in Proposition 1.3.17 represent all of the AGM postulates for contraction with the exception of the first one; in §1.3.4 we prove that also the first AGM postulate is validated in our framework.

Again, some other properties — dealing with multiple agents and universally minimal change — can be shown to hold.

Proposition 1.3.18 *For all $\varphi, \psi \in \mathcal{L}_0$ and for all $\chi \in \mathcal{L}$ we have:*

- $\models [\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{K}_j\mathbf{B}_i\psi \rightarrow \mathbf{K}_j\mathbf{B}_i\psi$
- $\models [\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_j\mathbf{B}_i\psi \rightarrow \mathbf{B}_j\mathbf{B}_i\psi$
- $\models \neg\mathbf{B}_i\varphi \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\chi \leftrightarrow \chi)$
- $\models \mathbf{B}_i\varphi \rightarrow (\chi \leftrightarrow [\mathrm{do}_i(\texttt{contract}\ \varphi; \texttt{expand}\ \varphi)]\chi)$
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \rightarrow ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\chi \leftrightarrow [\mathrm{do}_i(\texttt{contract}\ \psi)]\chi)$
- `contract` *satisfies realizability, determinism and idempotence*

The first two clauses of Proposition 1.3.18 state that the *a posteriori* knowledge and belief of watching agents on the beliefs of an agent performing a contraction are contained in their *a priori* knowledge and belief. The third clause states that contractions with disbelieved formulae cause no change at all. The fourth clause states that whenever an agent believes a formula, a contraction with that formula followed by an expansion with the same formula reduces to the void action and therefore causes no change. The fifth clause states that for formulae that an agent knows to be equivalent, a contraction with one formula causes exactly the same universal change as a contraction with the other formula. By the last clause, contractions obey the properties given in Definition 1.3.4.

1.3.4 Contractions and the AGM postulates

The AGM postulates for belief contraction are given below. In these postulates K, φ, ψ and K^+_φ are assumed to have their usual connotation, and $K^\perp_\varphi$ denotes the contraction of K with the formula φ.

Definition 1.3.19 The AGM postulates for belief contraction:

($G^\perp 1$) $K^\perp_\varphi$ is an AGM belief set.
($G^\perp 2$) $K^\perp_\varphi \subseteq K$.
($G^\perp 3$) If $\varphi \notin K$ then $K^\perp_\varphi = K$.
($G^\perp 4$) If $\not\vdash_{cpl} \varphi$ then $\varphi \notin K^\perp_\varphi$.
($G^\perp 5$) If $\varphi \in K$ then $K \subseteq (K^\perp_\varphi)^+_\varphi$.
($G^\perp 6$) If $\vdash_{cpl} \varphi \leftrightarrow \psi$ then $K^\perp_\varphi = K^\perp_\psi$.
($G^\perp 7$) $K^\perp_\varphi \cap K^\perp_\psi \subseteq K^\perp_{\varphi\wedge\psi}$.
($G^\perp 8$) If $\varphi \notin K^\perp_{\varphi\wedge\psi}$ then $K^\perp_{\varphi\wedge\psi} \subseteq K^\perp_\varphi$.

Using the definition of the `contract` action as given in 1.3.16, it is indeed the case that this action models contractions in the sense of the AGM postulates. As was the case for belief expansions, we have to modify the postulates for belief contraction somewhat to account for the agent-oriented, semantics-based character of our framework.

Definition 1.3.20 The *contraction* of $\mathcal{B}(i, \mathcal{M}, s)$ with $\varphi \in \mathcal{L}_0$, notation $\mathcal{B}^\perp_\varphi(i, \mathcal{M}, s)$ is defined by:

- $\mathcal{B}^\perp_\varphi(i, \mathcal{M}, s) = \{\psi \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi\}$

The sequence of a contraction with φ followed by an expansion with ψ of $\mathcal{B}(i, \mathcal{M}, s)$, notation $\mathcal{B}^{\perp+}_{\varphi\psi}(i, \mathcal{M}, s)$, is defined by:

- $\mathcal{B}^{\perp+}_{\varphi\psi}(i, \mathcal{M}, s) = \{\vartheta \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{contract}\ \varphi; \texttt{expand}\ \psi)]\mathbf{B}_i\vartheta\}$

Theorem 1.3.21 *Let $\mathcal{M}$ be some Kripke model. For all $s \in \mathcal{M}$ and for all $\varphi, \psi \in \mathcal{L}_0$ the following are true.*

($B^{\perp}1$) $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$ *is an AGM belief set.*
($B^{\perp}2$) $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) \subseteq \mathcal{B}(i, \mathcal{M}, s)$.
($B^{\perp}3$) *If* $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.
($B^{\perp}4$) *If* $\varphi \notin \mathcal{K}(i, \mathcal{M}, s)$ *then* $\varphi \notin \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$.
($B^{\perp}5$) *If* $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ *then* $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp+}_{\varphi\varphi}(i, \mathcal{M}, s)$.
($B^{\perp}6$) *If* $\varphi \leftrightarrow \psi \in \mathcal{K}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}^{\perp}_{\psi}(i, \mathcal{M}, s)$.
($B^{\perp}7$) $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) \cap \mathcal{B}^{\perp}_{\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$.
($B^{\perp}8$) *If* $\varphi \notin \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$.

Proof: Let $\mathcal{M}$ be some Kripke model with state s, and let σ be an arbitrary selection function for $\mathcal{M}$. Let $\varphi \in \mathcal{L}_0$ be arbitrary, and let $\mathcal{M}', s = \mathtt{r}(i, \texttt{contract}\ \varphi)(\mathcal{M}, s)$. We show that contractions based on σ satisfy the AGM postulates.

- ($B^{\perp}1$) This postulate is easily seen to be satisfied by the same argument as given for $B^{+}1$.
- ($B^{\perp}2$) By demand $\Sigma 1$ it follows that $\sigma(i, s, \varphi)$ yields a set of states from $\mathcal{M}$. It is easily seen that for $\varphi \in \mathcal{L}_0$ it holds that if $\mathcal{M}, s' \models \varphi$ for all $s' \in \mathcal{S}'$ then for all $\mathcal{S}'' \subseteq \mathcal{S}'$, $\mathcal{M}, s'' \models \varphi$ for all $s'' \in \mathcal{S}''$. Now if $\psi \in \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$, then $\mathcal{M}, s' \models \psi$ for all $s' \in \mathrm{B}(i, s) \cup \sigma(i, s, \varphi)$. Hence $\mathcal{M}, s' \models \psi$ for all $s' \in \mathrm{B}(i, s)$, and thus $\psi \in \mathcal{B}(i, \mathcal{M}, s)$.
- ($B^{\perp}3$) If $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$, then $\mathcal{M}, s' \models \neg\varphi$ for some $s' \in \mathrm{B}(i, s)$. Then $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$, and by $\Sigma 2$ it follows that $\sigma(i, s, \varphi) \subseteq \mathrm{B}(i, s)$. Thus $\mathrm{B}'(i, s) = \mathrm{B}(i, s)$ and hence $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.
- ($B^{\perp}4$) If $\varphi \notin \mathcal{K}(i, \mathcal{M}, s)$, then $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset$. Hence by $\Sigma 3$, $\sigma(i, s, \varphi) \neq \emptyset$, and thus, by $\Sigma 1$, $\mathrm{B}'(i, s)$ contains some s' such that $\mathcal{M}, s' \models \neg\varphi$. Since φ is propositional, then also $\mathcal{M}', s' \models \neg\varphi$, and hence $\mathcal{M}', s \not\models \mathbf{B}_i\varphi$. Thus $\varphi \notin \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$.
- ($B^{\perp}5$) Suppose $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$. We distinguish two cases:
  - $\varphi \in \mathcal{K}(i, \mathcal{M}, s)$. Then $\sigma(i, s, \varphi) = \emptyset$ by $\Sigma 3$. Hence $\mathrm{B}'(i, s) = \mathrm{B}(i, s)$, and $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$. Then $\mathcal{B}^{\perp+}_{\varphi\varphi}(i, \mathcal{M}, s) = \mathcal{B}^{+}_{\varphi}(i, \mathcal{M}, s)$. Now since $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ it follows by $B^{+}4$ that

$\mathcal{B}^{\perp+}_{\varphi\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.

- $\varphi \notin \mathcal{K}(i, \mathcal{M}, s)$. Then $\mathrm{B}'(i, s) = \mathrm{B}(i, s) \cup \mathcal{S}'$ with $\mathcal{S}' = \sigma(i, s, \varphi)$. By $\Sigma 1$, $\mathcal{S}' \subseteq [\![\neg\varphi]\!]$. Let $\mathcal{M}'', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}', s)$. Then by Definition 1.3.5 it follows that $\mathrm{B}''(i, s) = \mathrm{B}'(i, s) \cap [\![\varphi]\!]$. Now since $\mathrm{B}'(i, s) = \mathrm{B}(i, s) \cup \mathcal{S}'$ and since $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$, and thus $\mathrm{B}(i, s) \cap [\![\varphi]\!] = \mathrm{B}(i, s)$, we have that $\mathrm{B}''(i, s) = \mathrm{B}(i, s)$. Hence $\mathcal{B}^{\perp+}_{\varphi\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$.

Since in both cases $\mathcal{B}(i, \mathcal{M}, s) = \mathcal{B}^{\perp+}_{\varphi\varphi}(i, \mathcal{M}, s)$, also $\mathcal{B}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp+}_{\varphi\varphi}(i, \mathcal{M}, s)$, and hence we conclude that postulate $B^{\perp}5$ is validated.

- ($B^{\perp}6$) Suppose $\varphi \leftrightarrow \psi \in \mathcal{K}(i, \mathcal{M}, s)$. Then from $\Sigma 4$ it follows that $\sigma(i, s, \varphi) = \sigma(i, s, \psi)$. Thus $\mathtt{r}(i, \mathtt{contract}\ \psi)(\mathcal{M}, s) = \mathcal{M}', s$, and $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}^{\perp}_{\psi}(i, \mathcal{M}, s)$.
- ($B^{\perp}7$) Let $\rho \in \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) \cap \mathcal{B}^{\perp}_{\psi}(i, \mathcal{M}, s)$. This implies that $\mathcal{M}, s' \models \rho$ for all $s' \in \mathrm{B}(i, s) \cup \sigma(i, s, \varphi) \cup \sigma(i, s, \psi)$. Since by $\Sigma 5$, $\sigma(i, s, \varphi \wedge \psi) \subseteq \sigma(i, s, \varphi) \cup \sigma(i, s, \psi)$, it follows that $\mathcal{M}, s' \models \rho$ for all $s' \in \mathrm{B}(i, s) \cup \sigma(i, s, \varphi \wedge \psi)$. But then $\rho \in \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$. Since ρ was chosen arbitrarily it follows that $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) \cap \mathcal{B}^{\perp}_{\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$.
- ($B^{\perp}8$) Suppose $\varphi \notin \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$. Let $\mathcal{M}'', s = \mathtt{r}(i, \mathtt{contract}\ (\varphi \wedge \psi))(\mathcal{M}, s)$. We distinguish two cases:
  - $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$. Then $\varphi \wedge \psi \notin \mathcal{B}(i, \mathcal{M}, s)$, since $\mathcal{B}(i, \mathcal{M}, s)$ is an AGM belief set. From $B^{\perp}3$ it follows that $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$ and $\mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$ are both equal to $\mathcal{B}(i, \mathcal{M}, s)$.
  - $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$. Since $\varphi \notin \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$, it follows that some $s' \in \sigma(i, s, \varphi \wedge \psi)$ exists such that $\mathcal{M}'', s' \models \neg\varphi$, and since φ is propositional, also $\mathcal{M}, s' \models \neg\varphi$. But then $\sigma(i, s, \varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$, and by $\Sigma 6$ it follows that $\sigma(i, s, \varphi) \subseteq \sigma(i, s, \varphi \wedge \psi)$. Next, for all formulae $\rho \in \mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s)$, $\mathcal{M}, s' \models \rho$ for all $s' \in \mathrm{B}(i, s) \cup \sigma(i, s, \varphi \wedge \psi)$. Since $\sigma(i, s, \varphi) \subseteq \sigma(i, s, \varphi \wedge \psi)$ it follows that $\mathcal{M}, s' \models \rho$ for all $s' \in \mathrm{B}(i, s) \cup \sigma(i, s, \varphi)$, and hence $\mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$.

Since in both cases $\mathcal{B}^{\perp}_{\varphi\wedge\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$, we conclude that postulate $B^{\perp}8$ is validated. □

1.3.5 Selection functions revisited

In this section we elaborate on the concept of selection functions. In particular, we show that selection functions can be seen as providing the semantic counterpart of partial meet contraction functions as defined in the AGM framework. This link with partial meet contraction functions leads to a concrete, and fairly simple, instantiation of selection functions. Furthermore it is shown that our kind of selection functions provides a strengthening of the selection functions as proposed by Stalnaker in the context of a conditional logic [Stalnaker, 1968]. We start by relating the demands imposed on selection functions to the AGM postulates for belief contraction. This relation is presented in the following proposition, leading to a refinement of the results obtained in Theorem 1.3.21.

Proposition 1.3.22 *Let $\mathcal{M}$ be some model with state s, and let $\varphi \in \mathcal{L}_0$. Let $\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ be defined as in 1.3.16 but with σ replaced by an arbitrary function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$. Then it holds that:*

- *If ς satisfies $\Sigma 1$ then $B^{\perp}5$ is validated.*
- *If ς satisfies $\Sigma 2$ then $B^{\perp}3$ is validated.*
- *Given that ς satisfies $\Sigma 1$ it is the case that if ς satisfies $\Sigma 3$ then $B^{\perp}4$ is validated.*
- *If ς satisfies $\Sigma 4$ then $B^{\perp}6$ is validated.*
- *If ς satisfies $\Sigma 5$ then $B^{\perp}7$ is validated.*
- *If ς satisfies $\Sigma 6$ then $B^{\perp}8$ is validated.*

The implications given in Proposition 1.3.22 cannot be generalized to equivalences. That is, the conditions imposed on selection functions are *sufficient* to bring about validation of the postulates for belief contraction, but are not *necessary* to do so. The following example sheds some more light on this issue.

Example 1.3.23 Consider the single-agent language $\mathcal{L}$, based on $\Pi = \{p, q\}$ and $At = \{a\}$. Consider the model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ where

- $\mathcal{S} = \mathcal{T} \cup \mathcal{U}, \mathcal{T} = \{t_0, t_1, t_2\}, \mathcal{U} = \{u_0, u_1, u_2\}$
- $\pi(p, t_j) = \mathbf{1}$ iff $\pi(p, u_j) = \mathbf{1}$ iff $j = 0$ or $j = 1$
  $\pi(q, t_j) = \mathbf{1}$ iff $\pi(q, u_j) = \mathbf{1}$ iff $j = 0$
- $\mathrm{R}(1) = \mathcal{S}^2$
- $\mathrm{B}(1, s) = \{t_0\}$ for all $s \in \mathcal{S}$
- $\mathtt{r}$ is arbitrary

- $\mathsf{c}(1,a)(t) = \mathbf{1}$ for all $t \in \mathcal{T}$, $\mathsf{c}(1,a)(u) = \mathbf{0}$ for all $t \in \mathcal{U}$

Note that although the elements of $\mathcal{T}$ are copies of the elements of $\mathcal{U}$ on the propositional level, they do not satisfy the same set of formula. For $\mathcal{M}, t \models \mathbf{A}_1 a$ for each $t \in \mathcal{T}$, whereas $\mathcal{M}, u \models \neg\mathbf{A}_1 a$ for all $u \in \mathcal{U}$. Define the function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ for all $s \in \mathcal{S}$ as follows:

$$\begin{array}{ll}
\varsigma(1,s,\varphi) = \{t_0, u_0\} & \text{if } \mathcal{M}, t_0 \not\models \varphi \\
\varsigma(1,s,\varphi) = \{t_0\} & \text{if } \mathcal{M} \models \varphi \\
\varsigma(1,s,\varphi) = \{t \in \mathcal{T} \mid \mathcal{M}, t \not\models \varphi\} \cup \{t_0\} & \text{if } \mathcal{M}, t_0 \models \varphi \,\&\, \\
 & \quad \mathcal{M} \not\models \varphi \,\&\, \varphi \vdash_{cpl} p \\
\varsigma(1,s,\varphi) = \{u \in \mathcal{U} \mid \mathcal{M}, u \not\models \varphi\} \cup \{t_0\} & \text{if } \mathcal{M}, t_0 \models \varphi \,\&\, \\
 & \quad \mathcal{M} \not\models \varphi \,\&\, \varphi \not\vdash_{cpl} p
\end{array}$$

The function ς is not a selection function for $\mathcal{M}$. In particular, ς does meet only one of the demands given in Definition 1.3.15. To see this, take some arbitrary $s \in \mathcal{S}$.

- Since ς yields identical results for all $s \in \mathcal{S}$, demand $\Sigma 0$ is met.
- Since $t_0 \in \varsigma(1,s,p)$ and $t_0 \notin [\![\neg p]\!]$, demand $\Sigma 1$ is not met.
- Since $u_0 \in \varsigma(1,s,\neg p)$ and $u_0 \notin \mathrm{B}(1,s)$, demand $\Sigma 2$ is not met.
- Although $\mathcal{S} \cap [\![\neg\mathbf{tt}]\!] = \emptyset$, $\varsigma(1,s,\mathbf{tt}) = \{t_0\}$, and hence demand $\Sigma 3$ is not met.
- Although $\mathcal{M}, t_0 \models \mathbf{K}_1(\neg q \leftrightarrow (p \wedge \neg q))$, $\varsigma(1,s,\neg q) = \{t_0, u_0, u_2\}$ whereas $\varsigma(1,s,p \wedge \neg q) = \{t_0, t_2\}$, and hence demand Σ_4 is not met.
- Since it is the case that $\varsigma(1,s,p \wedge q) = \{t_0, t_1, t_2\}$, $\varsigma(1,s,p) = \{t_0, t_2\}$ and $\varsigma(1,s,q) = \{t_0, u_1, u_2\}$, $\varsigma(1,s,p \wedge q) \not\subseteq \varsigma(1,s,p) \cup \varsigma(1,s,q)$, and hence demand $\Sigma 5$ is not met.
- Although $\varsigma(1,s,q \wedge p) \cap [\![\neg q]\!] = \{t_1, t_2\} \neq \emptyset$, $\varsigma(1,s,q) \not\subseteq \varsigma(1,s,q \wedge p)$ and hence demand $\Sigma 6$ is not met.

Thus ς is by no means a selection function for $\mathcal{M}$. It is however easily seen that when defining $\mathtt{r}(1, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ based on the non-selection function ς, all the AGM postulates as phrased in Theorem 1.3.21 are validated[4].

Despite the negative results of Example 1.3.23, we can prove that when defining $\mathtt{r}$ for the $\mathtt{contract}$ action based on some function ς that adds

[4]One way to see this is by remarking that ς is a variant of the selection function σ_a as presented in Definition 1.3.26 in which one uses the fact that $\mathcal{T}$- and $\mathcal{U}$-worlds satisfy exactly the same sets of propositional formulae.

doxastic alternatives, validation of the AGM postulates imposes some weak variants of the demands for selection functions on ς.

Proposition 1.3.24 *Let $\mathcal{M}$ be some Kripke model with state s. Assume that* $\mathtt{r}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s)$ *is defined as in 1.3.16 with σ replaced by an arbitrary function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$. Then if* `contract` *is to meet the demands presented in Theorem 1.3.21 it follows that:*

- $\varsigma(i, s, \varphi) \subseteq [s]_{\mathrm{R}(i)}$
- $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] = \emptyset \;\&\; [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset \Rightarrow \varsigma(i, s, \varphi) \cap [\![\neg\varphi]\!] \neq \emptyset$

The fact that our approach using selection functions defines a contraction function which satisfies the AGM postulates is not as surprising as it might seem at first sight. Fact of the matter is that selection functions can be seen as *model-based, knowledge-restricted* variants of the *partial meet contraction functions* defined in the AGM framework. The idea behind partial meet contraction functions is that, given an underlying logic, contractions validating the AGM postulates can be implemented as follows. For a belief set K and a formula φ that is to be contracted, a partial meet contraction function yields the intersection of a set of maximal subsets of K that do not entail φ. More formal, $K_\varphi^\perp = \bigcap S(K \perp \varphi)$, where $K \perp \varphi$ is the set of belief sets K' that fail to imply φ and are maximal subsets of K, and S is a function that selects some of the elements of $K \perp \varphi$. Selection functions can be seen as the *semantic* counterpart of partial meet contraction functions. Whereas in partial meet contraction functions some of the maximal subsets not implying the contracted formula are selected, our selection function picks out some of the maximal subsets of the belief set not implying the formula, *given the model under consideration*. This can be seen as follows. Assume some Kripke model $\mathcal{M}$ with state s to be given. Assume furthermore that $\mathcal{M}, s \models \mathbf{B}_i\varphi \wedge \neg\mathbf{K}_i\varphi$ for some $\varphi \in \mathcal{L}_0$. Given this model, it is obvious that adding any of the worlds from $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ to the set of doxastic alternatives of i results in a model in which the agent no longer believes φ. Furthermore, it is even so obvious that it is sufficient to add exactly one of the worlds from $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ in order to result in a model in which the agent no longer believes φ. From this point of view, a selection function selects some of the maximal subsets of a belief set, and the resulting belief set is the intersection of these maximal subsets, all with respect to the model $\mathcal{M}$. The following proposition formalizes these informal ideas.

Proposition 1.3.25 *Let some Kripke model $\mathcal{M}$ with state s be given. Let σ be some selection function for $\mathcal{M}$. Define for $\varphi \in \mathcal{L}_0$:*

- $\mathcal{B}(i, \mathcal{M}, s) \perp \varphi = \{\{\mathcal{B}(i, \mathcal{M}, s)\}\}$ *if* $\varphi \in \mathcal{K}(i, \mathcal{M}, s)$
- $\mathcal{B}(i, \mathcal{M}, s) \perp \varphi = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i, s) \cup \{s''\}[\mathcal{M}, s' \models \psi]\} \mid s'' \in [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]\}$ *if* $\varphi \notin \mathcal{K}(i, \mathcal{M}, s)$
- $S(\mathcal{B}(i, \mathcal{M}, s) \perp \varphi) = \{\{\psi \in \mathcal{L}_0 \mid \forall s' \in \mathrm{B}(i, s) \cup \{s''\} [\mathcal{M}, s' \models \psi]\} \mid s'' \in \sigma(i, s, \varphi)\}$

Then $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \bigcap S(\mathcal{B}(i, \mathcal{M}, s) \perp \varphi)$.

The relation with the partial meet contraction functions as formalized in Proposition 1.3.25 suggest a concrete implementation of selection functions, the so called *All-is-Good*, or AiG, function. The idea underlying AiG functions is that, whenever necessary, all the states from the epistemic equivalence class that do not support the formula that is to be contracted, are added to the set of doxastic alternatives. The AiG function can be seen as the semantic counterpart of the *full meet contraction* function in the AGM framework. When performing full meet contraction of a belief set K with a formula φ, the intersection of all maximal subsets of K not implying φ results.

Definition 1.3.26 (The All-is-Good function) Let $\mathcal{M}$ be a Kripke model. The AiG function σ_a is for all $s \in \mathcal{M}$ and φ in $\mathcal{L}_0$ defined by:

- $\sigma_a(i, s, \varphi) = \mathrm{B}(i, s) \cap [\![\neg\varphi]\!]$ if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$
- $\sigma_a(i, s, \varphi) = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ otherwise

Proposition 1.3.27 *The AiG function σ_a as given in Definition 1.3.26 is a selection function.*

Proof: We successively show that the AiG function satisfies the demands for selection functions. So assume that σ_a is the AiG function for some model $\mathcal{M}$ with states s, s', and let $\varphi, \psi \in \mathcal{L}_0$ be arbitrary.

$\Sigma 0$. Suppose $s' \in [s]_{\mathrm{R}(i)}$. Then $\mathrm{B}(i, s) = \mathrm{B}(i, s')$ and $[s]_{\mathrm{R}(i)} = [s']_{\mathrm{R}(i)}$ and hence demand $\Sigma 0$ is met.

$\Sigma 1$. Since $\sigma_a(i, s, \varphi) = \mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$, and $\sigma_a(i, s, \varphi) = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$ otherwise, demand $\Sigma 1$ is indeed met.

$\Sigma 2$. Since $\sigma_a(i, s, \varphi) = \mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \subseteq \mathrm{B}(i, s)$ if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$, demand $\Sigma 2$ is obviously validated.

$\Sigma 3$. Demand $\Sigma 3$ follows directly from the definition of the AiG function.

Σ4. If $[s]_{\mathrm{R}(i)} \cap [\![\varphi]\!] = [s]_{\mathrm{R}(i)} \cap [\![\psi]\!]$, then both $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \mathrm{B}(i,s) \cap [\![\neg\psi]\!]$ and $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] = [s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!]$, which suffices to conclude Σ4.

Σ5. We distinguish four cases:

- $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$. In this case also $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \emptyset$. Hence $\sigma_a(i,s,\varphi \wedge \psi) = [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi \vee \neg\psi]\!] = [s]_{\mathrm{R}(i)} \cap ([\![\neg\varphi]\!] \cup [\![\neg\psi]\!]) = ([s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]) \cup ([s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!]) = \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$.
- $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$. In this case $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \mathrm{B}(i,s) \cap [\![\neg\psi]\!]$. Since $\mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$, it follows that $\sigma_a(i,s,\varphi \wedge \psi) = \sigma_a(i,s,\psi)$.
- $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$. This case is completely analogous to the previous one, resulting in $\sigma_a(i,s,\varphi \wedge \psi) = \sigma_a(i,s,\varphi)$.
- $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset, \mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$. Then also $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] \neq \emptyset$. In this case $\sigma_a(i,s,\varphi \wedge \psi) = \mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \mathrm{B}(i,s) \cap ([\![\neg\varphi]\!] \cup [\![\neg\psi]\!]) = (\mathrm{B}(i,s) \cap [\![\neg\varphi]\!]) \cup (\mathrm{B}(i,s) \cap [\![\neg\psi]\!]) = \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$.

Since in all four cases $\sigma_a(i,s,\varphi \wedge \psi) \subseteq \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$, we conclude that Σ5 is validated.

Σ6. We distinguish two cases:

- $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] = \emptyset$. In this case both $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] = \emptyset$ and $\mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$, and by an identical argument as given in the first case of the proof of Σ5, we conclude that $\sigma_a(i,s,\varphi \wedge \psi) = \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$.
- $\mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!] \neq \emptyset$. In this case $\sigma_a(i,s,\varphi \wedge \psi) = \mathrm{B}(i,s) \cap [\![\neg\varphi \vee \neg\psi]\!]$, and since $\sigma_a(i,s,\varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$, also $\mathrm{B}(i,s) \cap [\![\neg\varphi]\!] \neq \emptyset$. Hence $\sigma_a(i,s,\varphi) = \mathrm{B}(i,s) \cap [\![\neg\varphi]\!]$. Now if $\mathrm{B}(i,s) \cap [\![\neg\psi]\!] = \emptyset$ it follows by an identical argument as given in the third clause of the proof of Σ5 that $\sigma_a(i,s,\varphi \wedge \psi) = \sigma_a(i,s,\varphi)$. If $\mathrm{B}(i,s) \cap [\![\neg\psi]\!] \neq \emptyset$ it follows by an identical argument as given in the fourth clause of the proof of Σ5 that $\sigma_a(i,s,\varphi \wedge \psi) = \sigma_a(i,s,\varphi) \cup \sigma_a(i,s,\psi)$.

Since in both cases $\sigma_a(i,s,\varphi) \subseteq \sigma_a(i,s,\varphi \wedge \psi)$ we conclude that demand Σ6 is validated. □

Proposition 1.3.27 states that AiG functions are indeed selection functions, which implies that defining contractions in terms of AiG functions results in a validation of the AGM postulates. In light of the general feeling that full meet contraction — and therefore also AiG contraction — is not completely acceptable since it results in belief sets that are too small, it is important to recall that the AiG function is but a special instantiation of the general concept of selection functions. It is in particular not the case that the demands imposed on selection functions force AiG contraction.

The belief states resulting from an application of the `contract` action can be completely characterized in terms of *a priori* information, i.e., knowledge and belief, of the agent. In one of the clauses given below it is presupposed that `r` for the `contract` action is based on the AiG function, the other clause holds for general selection functions.

Proposition 1.3.28 *For all $\varphi, \psi \in \mathcal{L}_0$ we have:*

- $\models \neg \mathbf{B}_i\varphi \to ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i\psi)$
- $\models \mathbf{B}_i\varphi \to ([\mathrm{do}_i(\texttt{contract}\ \varphi)]\mathbf{B}_i\psi \leftrightarrow (\mathbf{B}_i\psi \wedge \mathbf{K}_i(\neg\varphi \to \psi)))$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for all models.*

Again the result of Proposition 1.3.28 can be rephrased to make it more in line with a characterization of full meet contraction in the AGM framework.

Proposition 1.3.29 *For all Kripke models $\mathcal{M}$ with state s, and for all $\varphi, \psi \in \mathcal{L}_0$ we have:*

- $\varphi \notin \mathcal{B}(i, \mathcal{M}, s) \Rightarrow (\psi \in \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) \Leftrightarrow \psi \in \mathcal{B}(i, \mathcal{M}, s))$
- $\varphi \in \mathcal{B}(i, \mathcal{M}, s) \Rightarrow (\psi \in \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) \Leftrightarrow \psi \in \mathcal{B}(i, \mathcal{M}, s)\,\&\,(\neg\varphi \to \psi) \in \mathcal{K}(i, \mathcal{M}, s))$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for $\mathcal{M}$.*

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

- Suppose $\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$. By $B^{\perp}3$ we have that $\mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}, s)$, and hence $\psi \in \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$ iff $\psi \in \mathcal{B}(i, \mathcal{M}, s)$.
- Suppose $\varphi \in \mathcal{B}(i, \mathcal{M}, s)$. Let $\mathcal{M}', s = \texttt{r}(i, \texttt{contract}\ \varphi)(\mathcal{M}, s)$. Then:

$$\psi \in \mathcal{B}^{\perp}_{\varphi}(i, \mathcal{M}, s)$$
$$\Leftrightarrow \psi \in \mathcal{B}(i, \mathcal{M}', s)$$

$\Leftrightarrow \mathcal{M}', s' \models \psi$ for all $s' \in \mathrm{B}'(i, s)$
$\Leftrightarrow \mathcal{M}', s' \models \psi$ for all $s' \in \mathrm{B}(i, s) \cup \sigma_a(i, s, \varphi)$
$\Leftrightarrow \mathcal{M}', s' \models \psi$ for all $s' \in \mathrm{B}(i, s)$ and
$\mathcal{M}', s' \models \psi$ for all $s' \in \sigma_a(i, s, \varphi)$
$\Leftrightarrow \mathcal{M}, s' \models \psi$ for all $s' \in \mathrm{B}(i, s)$ and
$\mathcal{M}, s' \models \psi$ for all $s' \in [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$
$\Leftrightarrow \mathcal{M}, s \models \mathbf{B}_i\psi$ and $\mathcal{M}, s \models \mathbf{K}_i(\neg\varphi \rightarrow \psi)$
$\Leftrightarrow \psi \in \mathcal{B}(i, \mathcal{M}, s)$ and $\neg\varphi \rightarrow \psi \in \mathcal{K}(i, \mathcal{M}, s)$ □

Besides the relation between selection functions and partial meet contraction functions, another interesting relation exists between our selection functions and those defined by Stalnaker [Stalnaker, 1968]. Stalnaker uses selection functions (to avoid confusion we use the term *conditional selection functions* to refer to selection functions in the sense of Stalnaker) in the context of a Kripke style semantics for conditional logic. Given a Kripke model $\mathcal{M}$ and a state s in $\mathcal{M}$, a conditional selection function f when applied to a pair (φ, s) yields the *most preferred* or *most reasonable* world given φ and s. Stalnaker gives four demands that a reasonable conditional selection function should meet:

1. φ is true at $f(\varphi, s)$.
2. $f(\varphi, s)$ is undefined only if s' is inaccessible from s for all worlds s' in which φ holds.
3. if φ is true at s then $f(\varphi, s) = s$.
4. if φ is true at $f(\psi, s)$ and ψ is true at $f(\varphi, s)$ then $f(\varphi, s) = f(\psi, s)$.

Intuitively, there seems to be at least some resemblance between the ideas underlying Stalnaker's conditional selection function and those underlying our selection functions. For although conditional selection functions aim at yielding a *single world* that *satisfies* a given formula and selection functions aim at yielding a *set of worlds* that *falsify* a given formula, both aim at yielding '*reasonable*' results. For conditional selection functions this reasonableness is enforced through the demands given above, whereas for selection functions this is enforced through the demands $\Sigma 1$ through $\Sigma 6$ as given in Definition 1.3.15. One could ask whether imposing demands similar to those proposed by Stalnaker would result in selection functions, and *vice versa*. To investigate this relation in our framework we introduce the notion of *s-selection functions*. Basically these s-selection functions can be seen as conditional selection functions that are at some points adapted to make them more in line with our framework.

Definition 1.3.30 Let some model $\mathcal{M}$ be given. A function $s : \mathcal{A} \times \mathcal{S} \times \wp(\mathcal{S})$ is an *s-selection function* for $\mathcal{M}$ if and only if it meets the following constraints for all $s \in \mathcal{S}$ and for all $\varphi, \psi \in \mathcal{L}_0$.

$S1$. $\neg\varphi$ is true at all states from $s(i, s, \varphi)$

$S2$. $s(i, s, \varphi)$ is empty only if s' is (epistemically) inaccessible from s for all worlds s' in which $\neg\varphi$ holds.

$S3$. if $\mathrm{B}(i, s) \cap [\![\neg\varphi]\!] \neq \emptyset$ then $s(i, s, \varphi) \subseteq \mathrm{B}(i, s)$

$S4$. if $\neg\varphi$ is true at all states from $s(i, s, \psi)$ and $\neg\psi$ is true at all states from $s(i, s, \varphi)$, then $s(i, s, \varphi) = s(i, s, \psi)$

It turns out that all selection functions for a given model, are also s-selection functions for the model. The converse does however not hold.

Proposition 1.3.31 *For all models $\mathcal{M}$, and for all functions $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ it holds that if ς is a selection function for $\mathcal{M}$ then ς is an s-selection function for $\mathcal{M}$.*

Proof: Let some Kripke model $\mathcal{M}$ and function ς be given such that ς is a selection function for $\mathcal{M}$. We show that ς satisfies the demands for an s-selection function. The properties $S1$ through $S3$ follow directly from demands $\Sigma 1$ through $\Sigma 3$, leaving only $S4$ to be proved. Hence assume that $\neg\varphi$ is true at all states from $\varsigma(i, s, \psi)$ and $\neg\psi$ is true at all states from $\varsigma(i, s, \varphi)$. This implies that $\varsigma(i, s, \psi) \subseteq [\![\neg\varphi]\!]$ and $\varsigma(i, s, \varphi) \subseteq [\![\neg\psi]\!]$. From $\Sigma 1$ it follows that $\varsigma(i, s, \varphi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!]$, and hence $\varsigma(i, s, \varphi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \cap [\![\neg\psi]\!]$. Analogously it follows that $\varsigma(i, s, \psi) \subseteq [s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \cap [\![\neg\psi]\!]$. If either $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] = \emptyset$ or $[s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!] = \emptyset$, then both $\varsigma(i, s, \varphi) = \emptyset$ and $\varsigma(i, s, \psi) = \emptyset$ and hence $S4$ would be met. Hence assume that $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \neq \emptyset$ and $[s]_{\mathrm{R}(i)} \cap [\![\neg\psi]\!] \neq \emptyset$. Then also $[s]_{\mathrm{R}(i)} \cap [\![\neg\varphi \vee \neg\psi]\!] \neq \emptyset$. From $\Sigma 3$ it follows that none of $\varsigma(i, s, \varphi)$, $\varsigma(i, s, \psi)$ and $\varsigma(i, s, \varphi \wedge \psi)$ is empty. By $\Sigma 5$ we have that $\varsigma(i, s, \varphi \wedge \psi) \subseteq \varsigma(i, s, \varphi) \cup \varsigma(i, s, \psi)$. Hence $\varsigma(i, s, \varphi \wedge \psi) \subseteq ([s]_{\mathrm{R}(i)} \cap [\![\neg\varphi]\!] \cap [\![\neg\psi]\!])$ (†). Then $\varsigma(i, s, \varphi \wedge \psi) \cap [\![\neg\varphi]\!] \neq \emptyset$, and hence by $\Sigma 6$ we have that $\varsigma(i, s, \varphi) \subseteq \varsigma(i, s, \varphi \wedge \psi)$ (‡). Analogously we have that $\varsigma(i, s, \psi) \subseteq \varsigma(i, s, \varphi \wedge \psi)$. Hence $\varsigma(i, s, \varphi \wedge \psi) = \varsigma(i, s, \varphi) \cup \varsigma(i, s, \psi)$. Now $\models \mathbf{K}_i(((\neg\varphi \vee \psi) \wedge \varphi) \leftrightarrow (\varphi \wedge \psi))$. Hence by $\Sigma 4$, $\varsigma(i, s, \varphi \wedge \psi) = \varsigma(i, s, (\neg\varphi \vee \psi) \wedge \varphi)$. From $\Sigma 5$ it follows that $\varsigma(i, s, \varphi \wedge \psi) \subseteq \varsigma(i, s, \neg\varphi \vee \psi) \cup \varsigma(i, s, \varphi)$. Since by $\Sigma 1$, $\varsigma(i, s, \neg\varphi \vee \psi) \subseteq$

$[\![\varphi \wedge \neg\psi]\!]$, it follows from (†) that $\varsigma(i, s, \varphi \wedge \psi) \cap \varsigma(i, s, \neg\varphi \vee \psi) = \emptyset$. Hence $\varsigma(i, s, \varphi \wedge \psi) \subseteq \varsigma(i, s, \varphi)$. Combining this with (‡) yields that $\varsigma(i, s, \varphi \wedge \psi) = \varsigma(i, s, \varphi)$. By an analogous argument we conclude from $\models \mathbf{K}_i(((\neg\psi \vee \varphi) \wedge \psi) \leftrightarrow (\varphi \wedge \psi))$ that $\varsigma(i, s, \varphi \wedge \psi) = \varsigma(i, s, \psi)$. Thus $\varsigma(i, s, \varphi) = \varsigma(i, s, \psi)$, which suffices to conclude that ς validates $S4$. Thus ς is an s-selection function. □

Proposition 1.3.32 *Some Kripke model $\mathcal{M}$ and function $\varsigma : \mathcal{A} \times \mathcal{S} \times \mathcal{L}_0 \to \wp(\mathcal{S})$ exist, such that ς is an s-selection function for $\mathcal{M}$, but not a selection function. Furthermore, when defining contractions based on the function ς, not all AGM postulates for belief contraction are validated.*

Proposition 1.3.32 states that it is not sufficient to use s-selection functions in defining contractions, for then not all AGM postulates are validated. As such, it is necessary to strengthen the demands proposed by Stalnaker in order to define AGM contraction in a modal context. Although Example 1.3.23 shows that the strengthening that we propose in our definition of selection functions is not *the necessary and sufficient* one, Theorem 1.3.21 shows that it is at least *a sufficient* one.

1.3.6 The `revise` action

Having defined actions that model expansions and contractions, we now turn to defining actions that model revisions. A revision is a change of belief through which some formula is added to the beliefs of an agent, while preserving consistency. Our definition of actions that model revisions is based on the *Levi identity* [Levi, 1977]. Levi suggested that revisions can be defined in terms of contractions and expansions: a revision with φ can be defined as a contraction with $\neg\varphi$ followed by an expansion with φ. Given the definitions of contractions and expansions of the previous sections and the fact that the class of actions Ac that we consider is closed under sequential composition, the Levi identity provides for a means to define revisions as the sequential composition of a contraction and an expansion action.

Definition 1.3.33 Let some model $\mathcal{M} = \langle \mathcal{S}, \pi, \mathrm{R}, \mathrm{B}, \mathtt{r}, \mathtt{c} \rangle$ with $s \in \mathcal{S}$ and $\varphi \in \mathcal{L}_0$ be given. We define:

- $\mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s) = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi; \mathtt{expand}\ \varphi)(\mathcal{M}, s)$

Definition 1.3.33 indeed provides for an intuitively acceptable formalization of belief revision.

Proposition 1.3.34 *For all $\varphi, \psi, \vartheta \in \mathcal{L}_0$ we have:*

- $\models [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\varphi$
- $\models [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\vartheta \to [\mathrm{do}_i(\texttt{expand } \varphi)]\mathbf{B}_i\vartheta$
- $\models \neg\mathbf{B}_i\neg\varphi \to ([\mathrm{do}_i(\texttt{expand } \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\vartheta)$
- $\models \mathbf{K}_i\neg\varphi \leftrightarrow [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\mathbf{ff}$
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \to ([\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\vartheta \leftrightarrow [\mathrm{do}_i(\texttt{revise } \psi)]\mathbf{B}_i\vartheta)$
- $\models [\mathrm{do}_i(\texttt{revise } \varphi \wedge \psi)]\mathbf{B}_i\vartheta \to [\mathrm{do}_i(\texttt{revise } \varphi; \texttt{expand } \psi)]\mathbf{B}_i\vartheta$
- $\models \neg[\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\neg\psi \to$
  $([\mathrm{do}_i(\texttt{revise } \varphi; \texttt{expand } \psi)]\mathbf{B}_i\vartheta \to [\mathrm{do}_i(\texttt{revise } \varphi \wedge \psi)]\mathbf{B}_i\vartheta)$

The first clause of Proposition 1.3.34 states that agents believe φ as the result of revising their beliefs with φ. The second clause states that a revision with φ results in the agent believing at most the formulae that it would believe after expanding its beliefs with φ, i.e., changing the belief set to incorporate φ consistently (if possible) — this is a revision with φ — results in a subset of the set of beliefs that results from straightforward inserting φ in the belief set — an expansion with φ. The third clause formalizes the idea that expansion is a special kind of revision: in cases where $\neg\varphi$ is not believed, expanding with φ and revising with φ amount to the same action. The left-to-right implication of the fourth clause states that if $\neg\varphi$ is known, i.e., $\neg\varphi$ is among the agent's principles, then the revision with φ results in the agent believing **ff**, i.e., the revision results in the absurd belief set. The right-to-left implication of the fourth clause states that the agent will believe **ff** only after it performs a revision with one of its principles. The fifth clause states that revisions with formulae that are known to be equivalent have identical results. The sixth clause formalizes the idea that the revision with the conjunction $\varphi \wedge \psi$ results in the agent believing at most the formulae that it would believe after a revision with φ followed by an expansion with ψ. The seventh clause states that if a revision with φ does not result in $\neg\psi$ being believed, then after revising with $\varphi \wedge \psi$ the agent believes at least the formulae that it would believe as the result of performing a revision with φ followed by an expansion with ψ. As Gärdenfors remarks, clauses 6 and 7 provide for some sort of *minimal change* condition on revisions. With the exception of the first of the AGM postulates for revision, all occur in Proposition 1.3.39

as validities. In §1.3.7 we come back to the AGM postulates for revision.

Also for revisions we can prove some properties dealing with multiple agents and universally minimal change.

Proposition 1.3.35 *For all* $\varphi, \psi \in \mathcal{L}_0$ *and for all* $\chi \in \mathcal{L}$ *we have:*

- $\models [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{K}_j\mathbf{B}_i\psi \to [\mathrm{do}_i(\texttt{expand } \varphi)]\mathbf{K}_j\mathbf{B}_i\psi$
- $\models [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_j\mathbf{B}_i\psi \to [\mathrm{do}_i(\texttt{expand } \varphi)]\mathbf{B}_j\mathbf{B}_i\psi$
- $\models \neg\mathbf{B}_i\varphi \to ([\mathrm{do}_i(\texttt{revise } \varphi)]\chi \leftrightarrow [\mathrm{do}_i(\texttt{expand } \varphi)]\chi)$
- $\models \mathbf{K}_i(\varphi \leftrightarrow \psi) \to ([\mathrm{do}_i(\texttt{revise } \varphi)]\chi \leftrightarrow [\mathrm{do}_i(\texttt{revise } \psi)]\chi)$
- `revise` *satisfies realizability, determinism and idempotence*

The first two clauses of Proposition 1.3.35 state that the knowledge and belief of watching agents on the beliefs of the executing agent following a revision are contained in the knowledge and belief following an expansion. The third clause states that whenever an agent does not believe some formula, a revision with the formula and an expansion with the formula constitute identical actions. The fourth clause states the property of identical change for known equivalences. The last clause states that revisions also validate the properties of Definition 1.3.4.

The belief sets resulting from application of the `revise` action are characterized as follows.

Proposition 1.3.36 *For all* $\varphi, \psi \in \mathcal{L}_0$ *we have:*

- $\models \mathbf{K}_i\neg\varphi \leftrightarrow [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\mathbf{ff}$
- $\models \neg\mathbf{B}_i\neg\varphi \to ([\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{B}_i(\varphi \to \psi))$
- $\models \neg\mathbf{K}_i\neg\varphi \wedge \mathbf{B}_i\neg\varphi \to ([\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\psi \leftrightarrow \mathbf{K}_i(\varphi \to \psi))$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for all models.*

The first clause of Proposition 1.3.36 states that in cases where an agent *knows* the negation of some formula to be true, a revision with this formula results in absurd beliefs. The second clause states that in situations where the negation of some formula is not *believed*, revising beliefs with the formula amounts to an expansion with the formula. The last clause states that in situations that are not of the kinds described in the first two clauses, formulae are believed after a revision with φ if it is *known* beforehand that φ implies the formula, i.e., the belief set of the agent after a revision with φ consists of all those formulae that are known to be implied by φ.

1.3.7 Revisions and the AGM postulates

The AGM postulates for belief revision are given below. In these postulates K, φ and K^+_φ are assumed to have their usual connotation, and $K^\star_\varphi$ denotes the revision of K with the formula φ.

Definition 1.3.37 The AGM postulates for belief revision:

$(G^\star 1)$ $K^\star_\varphi$ is an AGM belief set.
$(G^\star 2)$ $\varphi \in K^\star_\varphi$.
$(G^\star 3)$ $K^\star_\varphi \subseteq K^+_\varphi$.
$(G^\star 4)$ If $\neg\varphi \notin K$ then $K^+_\varphi \subseteq K^\star_\varphi$.
$(G^\star 5)$ $K^\star_\varphi = K_\perp$ if and only if $\vdash_{cpl} \neg\varphi$.
$(G^\star 6)$ If $\vdash_{cpl} \varphi \leftrightarrow \psi$ then $K^\star_\varphi = K^\star_\psi$.
$(G^\star 7)$ $K^\star_{\varphi\wedge\psi} \subseteq (K^\star_\varphi)^+_\psi$.
$(G^\star 8)$ If $\neg\psi \notin K^\star_\varphi$, then $(K^\star_\varphi)^+_\psi \subseteq K^\star_{\varphi\wedge\psi}$.

When defining revision through the Levi identity — starting from expansions and contractions that satisfy the appropriate postulates — the AGM postulates for belief revision are met [Gärdenfors, 1988; Levi, 1977]. The same holds in our framework.

Definition 1.3.38 The *revision* of $\mathcal{B}(i, \mathcal{M}, s)$ with $\varphi \in \mathcal{L}_0$, notation $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$, is defined by:

- $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \{\psi \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{revise } \varphi)]\mathbf{B}_i\psi\}$

The sequence of a revision with φ followed by an expansion with ψ of $\mathcal{B}(i, \mathcal{M}, s)$, notation $\mathcal{B}^{\star +}_{\varphi\psi}(i, \mathcal{M}, s)$, is defined by:

- $\mathcal{B}^{\star +}_{\varphi\psi}(i, \mathcal{M}, s) = \{\vartheta \in \mathcal{L}_0 \mid \mathcal{M}, s \models [\mathrm{do}_i(\texttt{revise } \varphi; \texttt{expand } \psi)]\mathbf{B}_i\vartheta\}$

Theorem 1.3.39 *Let $\mathcal{M}$ be some Kripke model. For all $s \in \mathcal{M}$ and for all $\varphi, \psi \in \mathcal{L}_0$ the following are true.*

$(B^\star 1)$ $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$ *is a belief set.*
$(B^\star 2)$ $\varphi \in \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.
$(B^\star 3)$ $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
$(B^\star 4)$ *If* $\neg\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) \subseteq \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.
$(B^\star 5)$ $\mathcal{B}(i, \mathcal{M}, s) = \mathcal{B}_\perp$ *if and only if* $\neg\varphi \in \mathcal{K}(i, \mathcal{M}, s)$.
$(B^\star 6)$ *If* $\varphi \leftrightarrow \psi \in \mathcal{K}(i, \mathcal{M}, s)$ *then* $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}^\star_\psi(i, \mathcal{M}, s)$.
$(B^\star 7)$ $\mathcal{B}^\star_{\varphi\wedge\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^{\star +}_{\varphi\psi}(i, \mathcal{M}, s)$.
$(B^\star 8)$ *If* $\neg\psi \notin \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$, *then* $\mathcal{B}^{\star +}_{\varphi\psi}(i, \mathcal{M}, s) \subseteq \mathcal{B}^\star_{\varphi\wedge\psi}(i, \mathcal{M}, s)$.

Proof: Let $\mathcal{M}$ be some Kripke model with state s. Let σ be an arbitrary selection function for $\mathcal{M}$. Let $\varphi \in \mathcal{L}_0$ be arbitrary, and let

- $\mathcal{M}', s = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi)(\mathcal{M}, s)$
- $\mathcal{M}'', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}', s) = \mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s)$
- $\mathcal{M}''', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}, s)$

We successively prove all clauses of Theorem 1.3.39.

- $(B^\star 1)$ This postulate is shown in the same way as the corresponding postulates for belief expansion and contraction.
- $(B^\star 2)$ Note that $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}'', s) = \mathcal{B}^+_\varphi(i, \mathcal{M}', s)$. From B^+2 it follows that $\varphi \in \mathcal{B}^+_\varphi(i, \mathcal{M}', s)$, which suffices to conclude that the postulate is validated.
- $(B^\star 3)$ By definition of $\mathtt{r}$ for `contract` and `expand` it follows that $\mathrm{B}''(i,s) = (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!]$. Now if some formula $\psi \in \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$, this means that $\mathcal{M}'', s'' \models \psi$ for all $s'' \in (\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!]$. Since ψ is propositional this implies that $\mathcal{M}, s'' \models \psi$ for all $s'' \in (\mathrm{B}(i,s) \cap [\![\varphi]\!])$. By definition of $\mathtt{r}(i, \mathtt{expand}\ \varphi)$ we have that $\mathrm{B}'''(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$. But then $\mathcal{M}''', s''' \models \psi$ for all $s''' \in \mathrm{B}'''(i,s)$, and hence $\mathcal{M}''', s \models \mathbf{B}_i\psi$, which implies that $\psi \in \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$.
- $(B^\star 4)$ If $\neg\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$, then $\mathrm{B}(i,s) \cap [\![\varphi]\!] \neq \emptyset$. By demand $\Sigma 2$ for selection functions, it follows that $\mathrm{B}'(i,s) = \mathrm{B}(i,s)$. Hence $\mathrm{B}''(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$. Also $\mathrm{B}'''(i,s) = \mathrm{B}(i,s) \cap [\![\varphi]\!]$, and $\mathcal{M}'', s \models \mathbf{B}_i\psi$ if and only if $\mathcal{M}''', s \models \mathbf{B}_i\psi$ for all $\psi \in \mathcal{L}_0$. Thus $\mathcal{B}^+_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}''', s) = \mathcal{B}(i, \mathcal{M}'', s) = \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$.
- $(B^\star 5)$ We prove two implications.
  - '$\Rightarrow$' Suppose $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}_\perp$. This implies that $\mathrm{B}''(i,s) = \emptyset$. Hence by definition of $\mathtt{r}$ for `contract` and `expand` this implies that $(\mathrm{B}(i,s) \cup \sigma(i,s,\neg\varphi)) \cap [\![\varphi]\!] = \emptyset$. In particular this implies that $\sigma(i,s,\neg\varphi) \cap [\![\varphi]\!] = \emptyset$, and since by $\Sigma 1$, $\sigma(i,s,\neg\varphi) \subseteq [\![\varphi]\!]$, we conclude that $\sigma(i,s,\neg\varphi) = \emptyset$. It follows by demand $\Sigma 3$ that $[s]_{\mathrm{R}(i)} \cap [\![\varphi]\!] = \emptyset$. This implies that $\mathcal{M}, s' \models \neg\varphi$ for all $s' \in [s]_{\mathrm{R}(i)}$ and thus $\mathcal{M}, s \models \mathbf{K}_i\neg\varphi$, and $\neg\varphi \in \mathcal{K}(i, \mathcal{M}, s)$.
  - '$\Leftarrow$' Suppose $\neg\varphi \in \mathcal{K}(i, \mathcal{M}, s)$. Then by demand $\Sigma 3$, $\sigma(i,s,\neg\varphi) = \emptyset$. Hence $\mathrm{B}(i,s) \subseteq [\![\neg\varphi]\!]$ and $\mathrm{B}'(i,s) \subseteq [\![\neg\varphi]\!]$. Then $\mathrm{B}''(i,s) = \mathrm{B}'(i,s) \cap [\![\varphi]\!] = \emptyset$, and thus $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}'', s) = \mathcal{B}_\perp$.

- $(B^\star 6)$ Suppose $\varphi \leftrightarrow \psi \in \mathcal{K}(i, \mathcal{M}, s)$. Then also $\neg\varphi \leftrightarrow \neg\psi \in \mathcal{K}(i, \mathcal{M}, s)$, and by $\Sigma 4$ it follows that $\sigma(i, s, \neg\varphi) = \sigma(i, s, \neg\psi)$. Thus $\texttt{r}(i, \texttt{contract}\ \varphi)(\mathcal{M}, s) = \texttt{r}(i, \texttt{contract}\ \psi)(\mathcal{M}, s) = \mathcal{M}', s$. Also $\mathrm{B}'(i, s) \cap [\![\varphi]\!] = \mathrm{B}'(i, s) \cap [\![\psi]\!]$, and hence $\texttt{r}(i, \texttt{expand}\ \varphi)(\mathcal{M}', s) = \texttt{r}(i, \texttt{expand}\ \psi)(\mathcal{M}', s)$. Then it follows that $\texttt{r}(i, \texttt{revise}\ \varphi)(\mathcal{M}, s) = \texttt{r}(i, \texttt{revise}\ \psi)(\mathcal{M}, s)$, and therefore $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}^\star_\psi(i, \mathcal{M}, s)$.
- $(B^\star 7)$ Assume that $\mathcal{M}_1, s = \texttt{r}(i, \texttt{revise}\ (\varphi \wedge \psi))(\mathcal{M}, s)$, and assume furthermore that $\mathcal{M}_2, s = \texttt{r}(i, \texttt{revise}\ \varphi; \texttt{expand}\ \psi)(\mathcal{M}, s)$. From the definitions of $\texttt{r}$ for revisions, contractions and expansions, it follows that $\mathrm{B}_1(i, s) = (\mathrm{B}(i, s) \cup \sigma(i, s, \neg\varphi \vee \neg\psi)) \cap [\![\varphi \wedge \psi]\!]$ and $\mathrm{B}_2(i, s) = (\mathrm{B}(i, s) \cup \sigma(i, s, \neg\varphi)) \cap [\![\varphi]\!] \cap [\![\psi]\!]$. Hence $\mathrm{B}_1(i, s) = (\mathrm{B}(i, s) \cap [\![\varphi \wedge \psi]\!]) \cup \sigma(i, s, \neg\varphi \vee \neg\psi)$ and $\mathrm{B}_2(i, s) = (\mathrm{B}(i, s) \cap [\![\varphi \wedge \psi]\!]) \cup (\sigma(i, s, \neg\varphi) \cap [\![\psi]\!])$. Hence, should $\sigma(i, s, \neg\varphi) \cap [\![\psi]\!] \subseteq \sigma(i, s, \neg\varphi \vee \neg\psi)$, then $\mathrm{B}_2(i, s) \subseteq \mathrm{B}_1(i, s)$, and therefore $\mathcal{B}(i, \mathcal{M}_1, s) \subseteq \mathcal{B}(i, \mathcal{M}_2, s)$. So to prove that $\sigma(i, s, \neg\varphi) \cap [\![\psi]\!] \subseteq \sigma(i, s, \neg\varphi \vee \neg\psi)$. Since $\models \mathbf{K}_i(((\neg\varphi \vee \neg\psi) \wedge (\neg\varphi \vee \psi)) \leftrightarrow \neg\varphi)$, we have by $\Sigma 4$ that $\sigma(i, s, \neg\varphi) = \sigma(i, s, (\neg\varphi \vee \neg\psi) \wedge (\neg\varphi \vee \psi))$. From $\Sigma 5$ it follows that $\sigma(i, s, (\neg\varphi \vee \neg\psi) \wedge (\neg\varphi \vee \psi)) \subseteq \sigma(i, s, \neg\varphi \vee \neg\psi) \cup \sigma(i, s, \neg\varphi \vee \psi)$. From $\Sigma 1$ we conclude that $\sigma(i, s, \neg\varphi \vee \neg\psi) \subseteq [\![\varphi \wedge \psi]\!]$ and $\sigma(i, s, \neg\varphi \vee \psi) \subseteq [\![\varphi \wedge \neg\psi]\!]$. Since $\sigma(i, s, \neg\varphi) \subseteq [\![\varphi]\!]$ we have that $\sigma(i, s, \neg\varphi) \cap [\![\psi]\!] \subseteq [\![\varphi \wedge \psi]\!]$. Hence $\sigma(i, s, \neg\varphi) \cap [\![\psi]\!] \subseteq \sigma(i, s, \neg\varphi \vee \neg\psi)$, which was to be proved.
- $(B^\star 8)$ Suppose $\neg\psi \notin \mathcal{B}^\star_\varphi(i, \mathcal{M}, s)$. To keep our proof understandable we introduce the following definitions:

  $$\begin{aligned} \mathcal{M}_{11}, s &= \texttt{r}(i, \texttt{revise}\ \varphi)(\mathcal{M}, s) \\ \mathcal{M}_1, s &= \texttt{r}(i, \texttt{expand}\ \psi)(\mathcal{M}_{11}, s) \\ \mathcal{M}_{21}, s &= \texttt{r}(i, \texttt{contract}\ \neg\varphi \vee \neg\psi)(\mathcal{M}, s) \\ \mathcal{M}_2, s &= \texttt{r}(i, \texttt{expand}\ \varphi \wedge \psi)(\mathcal{M}_{21}, s) = \texttt{r}(i, \texttt{revise}\ \varphi \wedge \psi)(\mathcal{M}, s) \end{aligned}$$

  Using similar arguments as in the proof of $B^\star 7$ we find that:

  $$\begin{aligned} \mathrm{B}_{11}(i, s) &= (\mathrm{B}(i, s) \cup \sigma(i, s, \neg\varphi)) \cap [\![\varphi]\!] \\ &= (\mathrm{B}(i, s) \cap [\![\varphi]\!]) \cup \sigma(i, s, \neg\varphi) \\ \mathrm{B}_1(i, s) &= \mathrm{B}_{11}(i, s) \cap [\![\psi]\!] \\ \mathrm{B}_{21}(i, s) &= \mathrm{B}(i, s) \cup \sigma(i, s, \neg\varphi \vee \neg\psi) \\ \mathrm{B}_2(i, s) &= (\mathrm{B}(i, s) \cup \sigma(i, s, \neg\varphi \vee \neg\psi)) \cap [\![\varphi \wedge \psi]\!] \\ &= (\mathrm{B}(i, s) \cap [\![\varphi \wedge \psi]\!]) \cup \sigma(i, s, \neg\varphi \vee \neg\psi) \end{aligned}$$

  Now to prove that $\mathrm{B}_2(i, s) \subseteq \mathrm{B}_1(i, s)$. For then it follows that $\mathcal{B}(i, \mathcal{M}_1, s) \subseteq \mathcal{B}(i, \mathcal{M}_2, s)$, which on its turn implies $\mathcal{B}^{\star+}_{\varphi\psi}(i, \mathcal{M}, s) \subseteq$

$\mathcal{B}^{\star}_{\varphi\wedge\psi}(i,\mathcal{M},s)$. We distinguish two cases:

- $\neg\varphi \in \mathcal{B}(i,\mathcal{M},s)$. Then $\mathrm{B}(i,s) \subseteq [\![\neg\varphi]\!]$, and thus $\mathrm{B}(i,s)\cap[\![\varphi]\!] = \emptyset$. Hence $\mathrm{B}_{11}(i,s) = \sigma(i,s,\neg\varphi)$ and $\mathrm{B}_1(i,s) = \sigma(i,s,\neg\varphi)\cap[\![\psi]\!]$. In this case also $\neg\varphi\vee\neg\psi \in \mathcal{B}(i,\mathcal{M},s)$. Hence $\mathrm{B}(i,s)\cap[\![\varphi\wedge\psi]\!] = \emptyset$, and thus $\mathrm{B}_2(i,s) = \sigma(i,s,\neg\varphi\vee\neg\psi)$. Since $\neg\psi \notin \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s)$, it follows that $\mathrm{B}_{11}(i,s)\cap[\![\psi]\!] \neq \emptyset$, and thus $\sigma(i,s,\neg\varphi)\cap[\![\psi]\!] \neq \emptyset$. From $\Sigma 1$ it follows that $\sigma(i,s,\neg\varphi)\cap[\![\varphi\wedge\psi]\!] = \sigma(i,s,\neg\varphi)\cap[\![\varphi]\!]\cap[\![\psi]\!] = \sigma(i,s,\neg\varphi)\cap[\![\psi]\!] \neq \emptyset$. Then since $\models \mathbf{K}_i(((\neg\varphi\vee\neg\psi)\wedge\neg\varphi)\leftrightarrow\neg\varphi)$, we have by $\Sigma 4$ that $\sigma(i,s,(\neg\varphi\vee\neg\psi)\wedge\neg\varphi)\cap[\![\varphi\wedge\psi]\!] \neq \emptyset$, and by $\Sigma 6$ that $\sigma(i,s,\neg\varphi\vee\neg\psi) \subseteq \sigma(i,s,(\neg\varphi\vee\neg\psi)\wedge\neg\varphi) = \sigma(i,s,\neg\varphi)$. Since by $\Sigma 1$, $\sigma(i,s,\neg\varphi\vee\neg\psi) \subseteq [\![\varphi\wedge\psi]\!]$, and given that $[\![\varphi\wedge\psi]\!] \subseteq [\![\psi]\!]$, it follows that $\sigma(i,s,\neg\varphi\vee\neg\psi) \subseteq \sigma(i,s,\neg\varphi)\cap[\![\psi]\!]$. Hence $\mathrm{B}_2(i,s) \subseteq \mathrm{B}_1(i,s)$.
- $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$. Note that in this case $\neg\varphi\vee\neg\psi \notin \mathcal{B}(i,\mathcal{M},s)$: for if $\neg\varphi\vee\neg\psi \in \mathcal{B}(i,\mathcal{M},s)$ and $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s)$, then from $B^{\star}4$, $B^{+}2$ and $B^{+}3$ it follows that $\{\neg\varphi\vee\neg\psi,\varphi\} \subseteq \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s)$. Since $\mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s)$ is deductively closed by $B^{\star}1$ it follows that $\neg\psi \in \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s)$ which contradicts the assumption that $\psi \notin \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s)$. Hence $\neg\varphi\vee\neg\psi \notin \mathcal{B}(i,\mathcal{M},s)$. This implies that both $\mathrm{B}(i,s)\cap[\![\varphi]\!] \neq \emptyset$ and $\mathrm{B}(i,s)\cap[\![\varphi\wedge\psi]\!] \neq \emptyset$. Then it follows by $\Sigma 2$ that both $\sigma(i,s,\neg\varphi) \subseteq \mathrm{B}(i,s)$ and $\sigma(i,s,\neg\varphi\vee\neg\psi) \subseteq \mathrm{B}(i,s)$. Thus $\mathrm{B}_{11}(i,s) = \mathrm{B}(i,s)\cap[\![\varphi]\!]$, $\mathrm{B}_1(i,s) = (\mathrm{B}(i,s)\cap[\![\varphi]\!])\cap[\![\psi]\!]$, $\mathrm{B}_{21}(i,s) = \mathrm{B}(i,s)$, and $\mathrm{B}_2(i,s) = \mathrm{B}(i,s)\cap[\![\varphi\wedge\psi]\!]$. Since for all $\mathcal{S}' \subseteq \mathcal{S}$ it holds that $(\mathcal{S}'\cap[\![\varphi]\!])\cap[\![\psi]\!] = \mathcal{S}'\cap[\![\varphi\wedge\psi]\!]$, for all φ and ψ in $\mathcal{L}_0$, it follows that $\mathrm{B}_1(i,s) = \mathrm{B}_2(i,s)$.

In both cases $\mathrm{B}_2(i,s) \subseteq \mathrm{B}_1(i,s)$, hence $\mathcal{B}(i,\mathcal{M}_1,s) \subseteq \mathcal{B}(i,\mathcal{M}_2,s)$, and thus $\mathcal{B}^{\star+}_{\varphi\psi}(i,\mathcal{M},s) \subseteq \mathcal{B}^{\star}_{\varphi\wedge\psi}(i,\mathcal{M},s)$. □

Also the result of Proposition 1.3.36 can be rephrased in terms that make it more in line with the AGM framework.

Proposition 1.3.40 *For all models $\mathcal{M}$ with state s, and for all $\varphi \in \mathcal{L}_0$:*

- $\neg\varphi \in \mathcal{K}(i,\mathcal{M},s) \Rightarrow \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathcal{B}_{\perp}$.
- $\neg\varphi \notin \mathcal{B}(i,\mathcal{M},s) \Rightarrow \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{B}(i,\mathcal{M},s)\cup\{\varphi\})$.
- $\neg\varphi \in \mathcal{B}(i,\mathcal{M},s)\setminus\mathcal{K}(i,\mathcal{M},s) \Rightarrow \mathcal{B}^{\star}_{\varphi}(i,\mathcal{M},s) = \mathrm{Th}(\mathcal{K}(i,\mathcal{M},s)\cup\{\varphi\})$ *if the definition of* `r` *for the* `contract` *action is based on the AiG function for $\mathcal{M}$.*

Proof: Let $\mathcal{M}$ be some model with state s, and let φ be some arbitrary propositional formula. Let $\mathcal{M}', s = \mathtt{r}(i, \mathtt{contract}\ \neg\varphi)(\mathcal{M}, s)$, and let $\mathcal{M}'', s = \mathtt{r}(i, \mathtt{expand}\ \varphi)(\mathcal{M}', s) = \mathtt{r}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s)$. We successively prove the three cases.

- Suppose $\neg\varphi \in \mathcal{K}(i, \mathcal{M}, s)$. Then by definition it follows that $\mathcal{M}', s = \mathcal{M}, s$, and hence $\mathcal{M}', s \models \mathbf{B}_i\neg\varphi$. Then the expansion with φ of the beliefs of agent i in $\mathcal{M}', s$ leads to a model $\mathcal{M}''$ such that $\mathrm{B}''(i, s) = \emptyset$, and hence $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}(i, \mathcal{M}'', s) = \mathcal{B}_\perp$.
- Suppose $\neg\varphi \notin \mathcal{B}(i, \mathcal{M}, s)$. Then it follows from $B^\star 3$ and $B^\star 4$ that $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathcal{B}^+_\varphi(i, \mathcal{M}, s)$, and by Proposition 1.3.14, $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathrm{Th}(\mathcal{B}(i, \mathcal{M}, s) \cup \{\varphi\})$.
- Suppose $\neg\varphi \in \mathcal{B}(i, \mathcal{M}, s) \backslash \mathcal{K}(i, \mathcal{M}, s)$. Then by definition of the AiG function it follows that $\mathrm{B}'(i, s) = \mathrm{B}(i, s) \cup ([s]_{\mathrm{R}(i)} \cap [\![\varphi]\!])$. By definition of $\mathtt{r}(i, \mathtt{expand}\ \varphi)$ it follows that $\mathrm{B}''(i, s) = \mathrm{B}'(i, s) \cap [\![\varphi]\!]$, hence $\mathrm{B}''(i, s) = (\mathrm{B}(i, s) \cup ([s]_{\mathrm{R}(i)} \cap [\![\varphi]\!])) \cap [\![\varphi]\!]$, and since $\neg\varphi \in \mathcal{B}(i, \mathcal{M}, s)$ it follows that $\mathrm{B}''(i, s) = [s]_{\mathrm{R}(i)} \cap [\![\varphi]\!]$. By an argument similar to that given in the proof of Proposition 1.3.14 it is shown that $\mathcal{B}^\star_\varphi(i, \mathcal{M}, s) = \mathrm{Th}(\mathcal{K}(i, \mathcal{M}, s) \cup \{\varphi\})$. □

Again note that, due to its lack of expressive power as compared to our framework, the reasonable and desirable properties of Proposition 1.3.35 cannot be formalized within the AGM framework.

1.3.8 The ability to change one's mind

In the previous (sub)sections, we dealt with the formalization of the opportunity for, and the result of, the actions that model the belief changes of agents. Here we look at the *ability* of agents to change their beliefs.

For 'mental' actions, like testing (observing) and communicating, the abilities of agents are closely related to their (lack of) information. This observation seems to hold *a fortiori* for the abstract actions that cause agents to change their beliefs. For when testing and communicating, at least some interaction takes place, either with the real world in case of testing, or with other agents when communicating, whereas the changing of beliefs is a strictly mental, agent-internal, activity. Therefore, it seems natural to let the ability of an agent to change its beliefs be determined by its informational state only.

The intuitive idea behind the definitions as we present them, is that the ability to change beliefs can be used to guide the changes that the beliefs of an agent undergo. In particular, if an agent is able to change its beliefs in a certain way, then this change of belief should work out as desired, i.e., it should neither result in an absurd belief set nor cause no change at all. Another point of attention is given by the observation that the Levi identity should also be respected for abilities, i.e., an agent is capable of revising its beliefs with a formula φ if and only if it is able to contract its beliefs with $\neg\varphi$ and thereafter perform an expansion with φ.

Definition 1.3.41 Let $\mathcal{M}$ be some Kripke model with state s, and let $\varphi \in \mathcal{L}_0$ be arbitrary. We define the capability function `c` for the `expand`, `contract` and `revise` actions in the following manner:

$$\begin{array}{lll} \mathtt{c}(i, \mathtt{expand}\ \varphi)(\mathcal{M}, s) = \mathbf{1} & \Leftrightarrow & \mathcal{M}, s \models \neg \mathbf{B}_i \neg\varphi \\ \mathtt{c}(i, \mathtt{contract}\ \varphi)(\mathcal{M}, s) = \mathbf{1} & \Leftrightarrow & \mathcal{M}, s \models \neg \mathbf{K}_i \varphi \\ \mathtt{c}(i, \mathtt{revise}\ \varphi)(\mathcal{M}, s) & = & \mathtt{c}(i, \mathtt{contract}\ \neg\varphi; \mathtt{expand}\ \varphi)(\mathcal{M}, s) \end{array}$$

The first clause of Definition 1.3.41 states that an agent is able to expand its set of beliefs with a formula if and only if it does not already believe the negation of the formula. The second clause formalizes the idea that an agent is able to remove some formula from its set of beliefs if and only if it does not consider the formula to be one of its principles. The ability for the `revise` action is defined through the Levi identity.

Proposition 1.3.42 *For all $\varphi \in \mathcal{L}_0$ we have:*

- $\models \mathbf{A}_i\mathtt{expand}\ \varphi \leftrightarrow \mathbf{K}_i\mathbf{A}_i\mathtt{expand}\ \varphi$
- $\models \mathbf{A}_i\mathtt{expand}\ \varphi \rightarrow \langle \mathrm{do}_i(\mathtt{expand}\ \varphi)\rangle \neg \mathbf{B}_i \mathbf{ff}$
- $\models \mathbf{A}_i\mathtt{contract}\ \varphi \leftrightarrow \mathbf{K}_i\mathbf{A}_i\mathtt{contract}\ \varphi$
- $\models \mathbf{A}_i\mathtt{contract}\ \varphi \rightarrow \langle \mathrm{do}_i(\mathtt{contract}\ \varphi)\rangle \neg \mathbf{B}_i \varphi$
- $\models \mathbf{A}_i\mathtt{revise}\ \varphi \leftrightarrow \mathbf{A}_i\mathtt{contract}\ \neg\varphi$
- $\models \mathbf{A}_i\mathtt{revise}\ \varphi \rightarrow \langle \mathrm{do}_i(\mathtt{revise}\ \varphi)\rangle (\mathbf{B}_i\varphi \wedge \neg \mathbf{B}_i \mathbf{ff})$

The first and third clause of Proposition 1.3.42 state that agents know of their ability to expand and contract their beliefs; a consequence of the fifth clause is that agents also know of their ability to revise their beliefs. The second, fourth and sixth clause formalize the idea that belief changes of which the agent is capable, behave as desired, i.e., an expansion does not result in absurd belief sets, a contraction leads to disbelief in the contracted formula, and a revision results in a combination of these.

1.4 Discussion

In this paper we defined actions that model three well-known rational changes of belief, viz. *expansions*, *contractions*, and *revisions*. We characterized the states of affairs that result from execution of these actions, the conditions that decide whether agents have the opportunity to perform these actions, and the capacities that agents should posses in order to be able to perform these actions. The action that models belief contractions is defined using selection functions. These are functions that select a subset of the set of epistemic alternatives of an agent that is to be added to its set of doxastic alternatives, in order to contract its set of beliefs. We proved that our kind of selection functions provides a strengthening of the selection functions proposed by Stalnaker [Stalnaker, 1968], and can furthermore be seen as the modal counterpart of the partial meet contraction functions as defined in the AGM framework. The action that models belief revision is defined in terms of a contraction and an expansion in a way suggested by the Levi identity [Levi, 1977]. Agents are capable of performing a belief-changing action only if execution of the action works out as desired. We showed that our belief-changing actions satisfy the AGM postulates for belief expansions, belief contractions, and belief revisions, thereby supporting our claim that the formalization that we present is both an intuitively and philosophically acceptable one.

Acknowledgements

Thanks are due to Maarten de Rijke and to an anonymous referee for their comments and criticism which helped improve this paper. This research is partially supported by Esprit III BRWG project No.8319 'ModelAge', Esprit III BRA project No.6156 'Drums II', and the Vrije Universiteit Amsterdam; the third author is furthermore partially supported by the Katholieke Universiteit Nijmegen.

1.5 References

[Alchourrón *et al.*, 1985] 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.

[Cohen and Levesque, 1990] P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[Elgesem, 1993] D. Elgesem. *Action Theory and Modal Logic*. PhD thesis, Institute for Philosophy, University of Oslo, Oslo, Norway, 1993.

[Fuhrmann, 1990] A. Fuhrmann. On the modal logic of theory change. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change*, volume 465 of *Lecture Notes in Artificial Intelligence*, pages 259–281. Springer-Verlag, 1990.

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

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

[Halpern and Moses, 1992] J.Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[Harel, 1984] 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.

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

[Hoek *et al.*, 1994a] W. van der Hoek, B. van Linder, and J.-J.Ch. Meyer. A logic of capabilities. In A. Nerode and Yu. V. Matiyasevich, editors, *Proceedings of the Third International Symposium on the Logical Foundations of Computer Science (LFCS'94)*, volume 813 of *Lecture Notes in Computer Science*, pages 366–378. Springer-Verlag, 1994.

[Hoek *et al.*, 1994b] W. van der Hoek, B. van Linder, and J.-J.Ch. Meyer. Unravelling nondeterminism: On having the ability to choose. In P. Jorrand and V. Sgurev, editors, *Proceedings of the Sixth International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA'94)*, pages 163–172. World Scientific, 1994.

[Kenny, 1975] A. Kenny. *Will, Freedom and Power.* Basil Blackwell, Oxford, 1975.

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

[Lesperance *et al.*, 1995] Y. Lesperance, H. Levesque, F. Lin, D. Marcu, R. Reiter, and R. Scherl. Foundations of a logical approach to agent programming. To appear in the proceedings of the IJCAI-95 workshop on Agent Theories, Architectures, and Languages.

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

[Linder *et al.*, 1994a] B. van Linder, W. van der Hoek, and J.-J.Ch. Meyer. Actions that make you change your mind: Belief revision in an agent-oriented setting. Technical Report UU-CS-1994-53, Utrecht University, December 1994.

[Linder *et al.*, 1994b] B. van Linder, W. van der Hoek, and J.-J.Ch. Meyer. Communicating rational agents. In B. Nebel and L. Dreschler-Fischer, editors, *KI-94: Advances in Artificial Intelligence*, volume 861 of *Lecture Notes in Computer Science (subseries LNAI)*, pages 202–213. Springer-Verlag, 1994.

[Linder *et al.*, 1994c] B. van Linder, W. van der Hoek, and J.-J.Ch. Meyer. The dynamics of default reasoning. To appear in the proceedings of ECSQARU'95.

[Linder *et al.*, 1994d] B. van Linder, W. van der Hoek, and J.-J.Ch. Meyer. Tests as epistemic updates. In A.G. Cohn, editor, *Proceedings of the 11th European Conference on Artificial Intelligence (ECAI'94)*, pages 331–335. John Wiley & Sons, 1994.

[Linder *et al.*, 1995] B. van Linder, W. van der Hoek, and J.-J.Ch. Meyer. Seeing is believing – and so are hearing and jumping. Technical Report UU-CS–1995-08, Utrecht University, March 1995.

[Meyer and Hoek, 1995] J.-J.Ch. Meyer and W. van der Hoek. *Epistemic Logic for AI and Computer Science.* Cambridge University Press, 1995. To appear.

[Moore, 1980] R.C. Moore. Reasoning about knowledge and action. Technical Report 191, SRI International, 1980.

[Moore, 1985] R.C. Moore. A formal theory of knowledge and action. In J.R. Hobbs and R.C. Moore, editors, *Formal Theories of the Commonsense World*, pages 319–358. Ablex Publ. Comp., 1985.

[Rao and Georgeff, 1991] A.S. Rao and M.P. Georgeff. Modeling rational agents within a BDI-architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484, San Mateo CA, 1991. Morgan Kaufmann.

[Rijke, 1994] M. de Rijke. Meeting some neighbours. In J. van Eijck and A. Visser, editors, *Logic and Information Flow*, pages 170–195. MIT Press, 1994.

[Segerberg, 1994] 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.

[Shoham, 1993] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60:51–92, 1993.

[Shoham and Cousins, 1994] Y. Shoham and S.B. Cousins. Logics of mental attitudes in AI. In G. Lakemeyer and B. Nebel, editors, *Foundations of Knowledge Representation and Reasoning*, volume 810 of *Lecture Notes in Computer Science (subseries LNAI)*, pages 296–309. Springer-Verlag, 1994.

[Stalnaker, 1968] R. Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in Logical Theory*, number 2 in American Philosophical Quarterly Monograph Series, pages 98–122. Blackwell, Oxford, 1968.

[Thijsse, 1992] E. Thijsse. *Partial Logic and Knowledge Representation.* PhD thesis, Katholieke Universiteit Brabant, 1992.

[Wright, 1963] G.H. von Wright. *Norm and Action.* Routledge & Kegan Paul, London, 1963.