Universität Konstanz
Geisteswissenschaftliche Sektion
Fachbereich Philosophie
Hauptfach Philosophie

Bachelorarbeit

Belief Revision in Dynamic Epistemic Logic and Ranking Theory

1. Gutachter: Dr. Franz Huber
2. Gutachter: Prof. Dr. Wolfgang Spohn

Vorgelegt von Peter Fritz
Konstanz, im Januar 2009

Abstract

I want to look at recent developments of representing AGM-style belief revision in dynamic epistemic logics and the options for doing something similar for ranking theory. Formally, my aim will be modest: I will define a version of basic dynamic doxastic logic using ranking functions as the semantics. I will show why formalizing ranking theory this way is useful *for the ranking theorist* first by showing how it enables one to compare ranking theory more easily with other approaches to belief revision. I will then use the logic to state an argument for defining ranking functions on larger sets of ordinals than is customary. Secondly, I will argue that the only way to extend the account of belief revision given by ranking theory to higher-order beliefs and revisions is by continuing the approach taken by me and defining ranking theoretical equivalents of dynamic epistemic logics. *For proponents of dynamic epistemic logic*, such logics will naturally be of interest provided they are convinced of the revision operator defined by ranking theory.

> *As we know, there are known knowns; there are things we know we know. We also know, there are known unknowns, that is to say, we know there are some things we do not know. But there are also unknown unknowns, the ones we do not know we do not know.*
>
> Donald Rumsfeld, stressing the importance of higher-order epistemic states in a Department of Defense news briefing, 12 February 2002

1 Introduction

The philosophical project of belief revision theory is a part of formal epistemology, and its aim is to characterise how rational agents revise their belief states when presented with new information. Many theories of belief revision have been proposed in the last 20 years, of which a large number can be put into one of two groups: those using systems of spheres as representations of epistemic states, and those using ranking functions. The former are significantly more popular and numerous, and the project of putting statements expressing belief and belief revision into the form of a logic has therefore only been carried out for them. Here, I want to look at the prospects of carrying over these technical advances to ranking theory, and the potential benefits of such a development.

The structure of the text is as follows. In chapter 2, I introduce AGM theory and ranking theory, chapter 3 defines the logics used, in chapter 4, I provide some results on comparing theories of belief revision using the logics defined in the previous chapter, chapter 5 states an argument on codomains of ranking functions, and chapter 6 considers the role of logics in ranking theory of higher-order beliefs and revisions. A mathematical appendix provides definitions of mathematical tools some readers may not be familiar with. All proofs in this text are straightforward, but given nonetheless.

2 Theories of Belief Revision

Theories of belief revision usually employ formal languages to characterise belief states of agents as well as information that is learned by them. A belief state is normally required to at least fix the set of sentences of the language the agent believes, although it may be richer than that. This set of believed sentences is called the *belief set* of the agent, and it is required to be closed under logical consequence. That is, if a sentence logically follows from a set of sentences in the belief set, it is in the belief set as well. New information in turn is represented by a sentence of the language.

Note that I wrote of agents believing *sentences*, not propositions. In so writing I mean to express that the agent believes what the sentence expresses. As is customary in belief revision theory, I pass over this distinction quickly, acknowledging that much of great philosophical subtlety could be said about it.

There is a broad consensus that when the agent learns a sentence such that she doesn't disbelieve it (an agent is said to *disbelieve* a sentence iff she believes its negation), the new belief set should be the set of sentences that are logical consequences of the union of the old belief set and the set containing only the new sentence. This process is called *expansion*.

More problematic are the cases where an agent retracts her belief in a sentence, called *contraction*, and the case where an agent comes to believe a sentence she formerly disbelieved, called *revision*. It is generally accepted that these three processes are linked by the so-called Levi identity:

> Revision by φ is equivalent to contraction with $\neg\varphi$ and subsequent expansion with φ.

The problem with contraction (and through the Levi identity, with revision) is this: The most natural formalization of contracting a belief state with a sentence φ would be to take the largest (with respect to the subset relation) closed subset of the prior belief set that fails to imply φ. However, this is not a unique characterisation, as in general there are several such sets. Using the intersection of all such sets is not an option either, as this would only leave those sentences in the belief set that are consequences of the negation of the sentence with which was contracted, which is clearly too small.

2.1 AGM Belief Revision

This problem was stated much more rigorously in Alchourrón et al. (1985), and a solution was proposed, which laid the foundation for all subsequent work on belief revision. It is often called the "AGM theory" after the first letters of the names of the authors of this paper, viz. Charlos Alchourrón, Peter Gärdenfors, and David Makinson. I want to present this theory in the form of Grove (1988), as this will make it easier to compare it to other theories discussed in the text. I will only give an informal account, for a rigorous definition, see Grove (1988, pp 157–163). Note that Grove is primarily concerned with revision, not contraction. As stated, this does not make a big difference, as any theory of revision is also a theory of contraction, on account of the uncontroversial Levi identity (and *vice versa*).

Grove's account starts like the usual Kripke-semantics for modal logics: There is a set of worlds W, each propositional variable of the formal language is assigned a set of worlds in which it holds, and the Boolean connectives $\neg, \wedge$ are interpreted as complementation and intersection. Grove then uses a variant of the notion of *systems of spheres* introduced in Lewis (1973): For any set of worlds X, a system of spheres, centered on X, is a set of sets of worlds that is a wellorder (with respect to $\subseteq$) with X as its minimal element and W as its maximal element. It follows that the individual elements of a system of spheres are sets of worlds which are nested like matryoshka dolls or the layers of an onion, and are called "spheres".

A belief state of an agent is represented by a system of spheres. The intended philosophical interpretation is that the worlds of the innermost sphere are the most plausible, and the worlds of each sphere are less plausible then the ones of the spheres it contains. A system of spheres is therefore equivalent to a plausibility order that is a wellorder. What does it mean for a world to be more plausible than another world? As worlds represent *ways the world might be*, it means that the agents holds it more plausible that the world is *that way* than the *other way*. The agent believes a sentence iff none of the most plausible worlds contradict it. Let a world be a φ-world iff φ is true in it. Then this means that an agent believes φ iff the set of minimal worlds (with respect to the plausibility relation) is a subset of the set of φ-worlds.

This enables Grove to define belief revision: Given the system of spheres representing the belief state of the agent, after revising with φ, the agent is in a belief state which is represented by a system of spheres whose smallest sphere consists of the minimal φ-worlds. After revising with φ, she therefore believes exactly those propositions which hold in all worlds of the intersection of the worlds in which φ holds with the smallest sphere (in the original system of spheres) that contains a world in which φ holds.

There is a problem with this theory: We starts with the belief state as given by a system of spheres, but after revision we only know which propositions the agents believes. Therefore, AGM theory cannot give an account of iterated revision. This criticism was voiced in Spohn (1988), and an alternative account of belief revision was developed there. It is now know under the name of *ranking theory*, which I present in the next section.

2.2 Belief Revision in Ranking Theory

This criticism against AGM theory is open to an obvious reply: AGM theory attempts to explain single revisions,and it succeeds in that. Obviously, an account of iterated revision would have to specify which system of spheres results from revising a belief state as given by a certain system of spheres. Nothing said so far shows that this cannot be done.

Spohn tries to do exactly this in Spohn (1988, pp 112–115). He argues as follows: A *prima facie* plausible suggestion how to define revision is this: After revision with φ, the φ-worlds are more plausible than the non-φ-worlds. In terms of system of spheres: After revising with φ, there is a new sphere for the set of φ-worlds of each existing sphere.

One problem Spohn raises for this solution is that this makes any world in which $\neg\varphi$ holds less plausible than any world in which φ holds. One could therefore look for another account of iterated revision in which the changes made to the belief state are minimal. According to AGM theory, the minimal sphere after revision will be the intersection of the worlds in which φ holds with the smallest sphere that contains a world in which φ holds. So why not create a new sphere with these worlds, make it the innermost sphere and add all of these worlds to all the other spheres (otherwise it wouldn't be a sphere system)?

(A note on terminology: Hans van Ditmarsch has called the former of these proposals "maximal-Spohn revision" and the latter "minimal-Spohn revision", e.g. in van Ditmarsch (forthcoming). As Spohn clearly does not support either of these accounts of belief revision, I'm unsure whether one should use these terms. An alternative terminology would be the one used in Rott (2008). He calls the former *lexicographic* or *moderate* revision and the latter *natural* or *conservative* revision.)

Spohn raises the following problems for these two accounts: If revision is defined in one of these two ways, a revision cannot be reversed. That is, in general, there is no other revision such that performing the two successively does not change the belief state. Also, revising with φ_1 and then with φ_2 does not result in the same system of spheres as revising with φ_2 and then with φ_1. Particularly the first of these problems points to what Spohn perceives as the underlying problem of the approach using systems of spheres: As spheres are just sets of worlds, if a revision results in two spheres having the same elements, they thereby are identical, which makes it impossible to define a revision rule that allows for reversibility.

Spohn's conclusion is that systems of spheres do not provide enough information about an agent's belief state. Just as AGM theory uses system of spheres instead of belief sets to model single belief revisions, as the latter do not provide enough information, so Spohn wants to use ranking functions instead of systems of spheres to model *iterated* belief revision.

What are ranking functions? In Spohn (1988), ranking functions are defined as functions from the set of possible worlds to the class of ordinals such that at least one world is mapped to 0. Subsequent publications vary the codomain of ranking functions, and this question will be discussed in detail in section 5. The following may help picturing ranking functions: As a ranking function assigns an ordinal number to each world, it is equivalent to a system of spheres in which the spheres are assigned ordinals. Any sphere contained in another sphere will be assigned a lower ordinal, but the numbering is allowed to contain gaps. As we identify propositions with sets of worlds, we can define a derivative ranking function from propositions to ordinals as follows: Any proposition is mapped to the minimum of the set of ordinals the worlds it contains are mapped to. The set of worlds a ranking functions assigns 0 to is called the *core* of that ranking function. The core of a ranking function takes the place of the smallest sphere in systems of spheres: in ranking theory, the agent believes φ iff the core of the ranking function representing her belief state is a subset of the φ-worlds.

What is the ranking-theoretic account of revision? This is described in Spohn (1988, p 117). In his account of revision, Spohn wants to represent the credibility of new information, which he measures with ordinals. His account of revision therefore has the following form: Given a ranking function representing the belief state of the agent, a set of worlds A representing the proposition with which the belief state is to be revised, and an ordinal α representing the credibility of this information, the revised ranking function is as follows. The new rank of any A-world is the rank of A subtracted from the old rank of the world, and the rank of any $\neg A$-world is α plus the rank of $\neg A$ subtracted from the old rank of the world. This is called A, α-*conditionalization*, and I will give a mathematically more precise definition of it after having defined the logics. Other kinds of revision can be defined in ranking theory representing different kinds of learning processes. One such alternate definition of revision is Shenoy conditionalization, which will be discussed later as well.

The ranking theoretical account of belief revision obviously allows for iterated revisions. Furthermore, it can be proven that it has neither of the problems posed to accounts using systems of spheres, so it solves the problems posed to AGM theory by Spohn.

3 Logics of Belief Revision

AGM theory of belief revision has been put into the form of a logic called dynamic doxastic logic (DDL). This family of logics uses systems of spheres for its semantics, and it can accommodate a range of belief revision theories operating with systems of spheres. (Systems of spheres are sometimes called *hypertheories*, e.g. in Segerberg (1995).) DDL was mainly developed by Krister Segerberg, and I will present a variant of the version presented in Leitgeb and Segerberg (2007).

As I also define a similar logic for ranking theory, I make significant changes to the logic as defined in their paper, so theorems proved for their logic may not hold in mine. I call the logics "basic dynamic doxastic logic" (BDDL), and "rankified basic dynamic doxastic logic" (RBDDL). They are *basic*, as there are certain restrictions already present on the syntactic level. For definitions of *full* or *unlimited* DDL and further classifications properly between DDL and BDDL, see Segerberg (2001, pp 62–63) and Lindström and Rabinowicz (1997). RBDDL is *rankified*, as it is obtained from BDDL by replacing systems of spheres with ranking functions (plus some minor adjustments).

Ranking theory uses ordinal parameters for its revisions. To completely represent this theory, a logic of ranking theory must syntactically represent this parameter. If the full class of ordinals would be used for this, the syntax would specify a proper class of syntactically well-formed expressions. To avoid this, some countable ordinal number λ is chosen as the limit of the ordinals used. Note that $\lambda + 1 = \{\alpha \in Ord | \alpha \leq \lambda\}$, and that, as the choice of λ affects semantics as well as *syntax*, one should properly speak of the *family* of BDDL or RBDDL logics, with e.g. ω^2-RBDDL being a member of RBDDL. For convenience, I will speak just of BDDL and RBDDL in most of the following, assuming λ to be some chosen ordinal limit.

I start by defining the syntax for BDDL as well as RBDDL.

3.1 Syntax

First, I define the set of pure Boolean formulas pB for a set of propositional letters P:

If $p \in P$, then $p \in pB$.

If $\varphi_1, \varphi_2 \in pB$, then $\neg\varphi_1 \in pB$ and $(\varphi_1 \wedge \varphi_2) \in pB$.

With this, the set of formulas F_B is defined:

If $\varphi \in pB$, then $\varphi \in F_B$.

If $\psi_1, \psi_2 \in F_B$, then $\neg\psi_1 \in F_B$ and $(\psi_1 \wedge \psi_2) \in F_B$.

If $\varphi \in pB, \psi \in F_B, \alpha \leq \lambda$, then $B\varphi \in F_B$ and $[\varphi, \alpha]\psi \in F_B$.

Other Boolean operators will be understood to be syntactic variations of the ones presented, e.g. $\psi_1 \rightarrow \psi_2$ is to be treated as $\neg(\psi_1 \wedge \neg\psi_2)$. Additionally, for any $n \in \mathbf{N}$, $[\varphi, \alpha]^n\psi$ is used to stand for $\overbrace{[\varphi, \alpha] \ldots [\varphi, \alpha]}^{n}\psi$.

3.2 Semantics of Pure Boolean Formulas

Before specifying the semantics of BDDL and RBDDL, here are those definitions common to both:

A model is an algebraic structure containing a non-empty set of worlds W and a valuation V on W. A *valuation V on W* is a function from the propositional letters to the propositions on W, while the *set of propositions on W* is the power set of W. With this, the semantics of pure Boolean formulas is defined by a function $[\![\cdot]\!]_V$ from the pure Boolean formulas to the propositions for any valuation V on W:

For any propositional letter p, $[\![p]\!]_V = V(p)$.

For any pure Boolean formula φ, $[\![\neg\varphi]\!]_V = W \backslash [\![\varphi]\!]_V$.

For any pure Boolean formulas φ_1 and φ_2, $[\![(\varphi_1 \wedge \varphi_2)]\!]_V = [\![\varphi_1]\!]_V \cap [\![\varphi_2]\!]_V$.

3.3 Semantics of BDDL

In BDDL, belief states are modelled by systems of spheres. Revisions are functions from propositions and ordinals up to λ to relations between systems of spheres and systems of spheres. The idea behind this is that a revision specifies for each proposition and ordinal which systems of spheres can rationally be entertained after revising a certain system of spheres by that proposition with that ordinal parameter. Now for the formal definitions:

A *system of spheres s on W* is a set of subsets of W that is a wellorder with respect to the subset relation such that $\emptyset \notin s$, $\bigcup s = W$, and s's order type is in $\lambda + 1$. For a set S of systems of spheres, a *revision on S* is a function from the set of propositions on W and $\lambda + 1$ to the binary relations on S. A *revision frame* is a structure $\langle W, S, \star \rangle$, where W is a non-empty set of worlds, S is a set of systems of spheres on W, and $\star$ is a revision on S. A *revision model* is a structure $\langle W, S, \star, V \rangle$ such that $\langle W, S, \star \rangle$ is a revision frame and V is a valuation on W.

Truth of a formula ψ in a revision model $\mathcal{M} = \langle W, S, \star, V \rangle$ is defined relative to a system of spheres s (representing the epistemic state of the agent) and a world w (representing the factual state) denoted by $\mathcal{M}, s, w \models \psi$. It is defined as follows:
For any pure Boolean formula φ:

$$\mathcal{M}, s, w \models \varphi \qquad \text{iff } w \in [\![\varphi]\!]_V$$

For any formulas ψ, ψ' and pure Boolean formula φ:

$$\mathcal{M}, s, w \models \neg\psi \qquad \text{iff not } \mathcal{M}, s, w \models \psi$$

$$\mathcal{M}, s, w \models (\psi \wedge \psi') \qquad \text{iff } \mathcal{M}, s, w \models \psi \text{ and } \mathcal{M}, s, w \models \psi'$$

$$\mathcal{M}, s, w \models B\varphi \qquad \text{iff } \bigcap s \subseteq [\![\varphi]\!]_V$$

$$\mathcal{M}, s, w \models [\varphi, \alpha]\psi \qquad \text{iff } \forall s' \in S(s \star ([\![\varphi]\!]_V, \alpha)s' \rightarrow \mathcal{M}, s', w \models \psi)$$

The truth conditions for pure Boolean formulas as well as for formulas of the form $\neg\psi$ and $(\psi \wedge \psi')$ are straightforward. An agent believes a proposition iff the proposition holds in all worlds she considers most plausible, so $B\varphi$ is true iff the innermost sphere of the system of spheres representing the belief state of the agent is a subset of the set of φ-worlds. If the agent revises her belief state by φ with parameter α, ψ holds iff ψ holds in all systems of spheres accessible from the current one via the relation obtained by applying the revision to φ and α.

One can now define a certain revision operator such that the true revisions are exactly those of AGM theory. For their version of BDDL, this has been done in Leitgeb and Segerberg (2007). In this paper, there is also an axiomatisation, but for the completeness proof, the reader is referred to an unpublished manuscript by Segerberg from 2005. I also want to stress that the proof might not be applicable to the variant of BDDL presented here.

Note that AGM theory of belief revision knows no revision parameters, so when representing it in BDDL, the revision will ignore the parameter. It may seem to be the case that no sensible revision could be defined for BDDL that does not ignore the parameter, as systems of spheres are purely relational. However, this is not the case – an example for such a revision can be derived from Nayak et al. (2007, definition 3, p 2572).

A few more notions are introduced which will be used later: For any countable ordinal α, S_W^α is the set of systems of spheres on W whose order type is in $\alpha+1$. For any functional revision $\star$, $s^\star_{A,\alpha}$ is the unique s' such that $s \star (A,\alpha)s'$.

As examples, I define lexicographic and natural revision $\star^{lex}$ and $\star^{nat}$, which are both functional. For any proposition A, $\alpha \leq \lambda$, $s \in S_W^\lambda$,

$$s^{\star^{lex}}_{A,\alpha} = ((\{x \cap A | x \in s\} \cup \{x \cup A | x \in s\}) \backslash \emptyset)|^\lambda$$

$$s^{\star^{nat}}_{A,\alpha} = (m \cup \{x \cup m | x \in s\})|^\lambda, \text{ where } m = \bigcap((\{x \cap A | x \in s\}) \backslash \emptyset)$$

These definitions are straightforward implementations of the revisions described earlier, except for the function $\cdot|^\lambda$ from sets of subsets of W that are wellorders with respect to $\subseteq$ to S_W^λ, which is needed to make sure the range of both relations is a subset of S_W^λ, and defined such that for any s,

$$s|^\lambda = \{s_\alpha | \alpha < \lambda\} \cup \{W\}, \text{ where } s_\alpha = \bigcap(s \backslash \{s_\beta | \beta < \alpha\})$$

3.4 Semantics of RBDDL

The semantics of RBDDL works strictly analogous to the semantics of BDDL, except that ranking functions are used instead of systems of spheres:

A *ranking function on* W is a function from W to $\lambda+1$ such that at least one world is mapped to 0. The *core* of a ranking function is the set of worlds the function maps to 0, written $core(\kappa)$. For a set K of ranking functions, a *revision on* K is a function from the set of propositions on W and $\lambda+1$ to the binary relations on K. A *revision frame* is a structure $\langle W, K, * \rangle$, where W is a set of possible worlds, K is a set of ranking functions on W, and $*$ is a revision on K. A *revision model* is a structure $\langle W, K, *, V \rangle$ such that $\langle W, K, * \rangle$ is a revision frame and V is a valuation on W.

Truth of a formula ψ in a revision model $\mathcal{M} = \langle W, S, *, V\rangle$ is defined relative to a ranking function κ (representing the epistemic state of the agent) and a world w (representing the factual state) denoted by $\mathcal{M}, \kappa, w \models \psi$. It is defined as follows:
For any pure Boolean formula φ:

$$\mathcal{M}, \kappa, w \models \varphi \qquad \text{iff } w \in [\![\varphi]\!]_V$$

For any formulas ψ, ψ' and pure Boolean formula φ:

$$\mathcal{M}, \kappa, w \models \neg\psi \qquad \text{iff not } \mathcal{M}, \kappa, w \models \psi$$

$$\mathcal{M}, \kappa, w \models (\psi \wedge \psi') \qquad \text{iff } \mathcal{M}, \kappa, w \models \psi \text{ and } \mathcal{M}, \kappa, w \models \psi'$$

$$\mathcal{M}, \kappa, w \models B\varphi \qquad \text{iff } core(\kappa) \subseteq [\![\varphi]\!]_V$$

$$\mathcal{M}, \kappa, w \models [\varphi, \alpha]\psi \qquad \text{iff } \forall \kappa' \in K(\kappa \star ([\![\varphi]\!]_V, \alpha)\kappa' \rightarrow \mathcal{M}, \kappa', w \models \psi)$$

I am not concerned with axiomatising the theorems of any class of frames of RBDDL here. The reason is that this thesis defends the (non-trivial) claim that a logic of ranking theory as defined here is *useful* and *interesting*. Only if this is established, axiomatisation is sensible. The following references provide starting points on axiomatisation: Aucher (2005), Hild and Spohn (2008), Spohn (2008, chap. 8).

Again, a few more notions are introduced: For any countable ordinal α, $K^\alpha_W = \{\kappa : W \rightarrow \alpha + 1 | \exists w \in W(\kappa(w) = 0)\}$. $[\cdot]^{\cdot} : Ord^2 \rightarrow Ord$ is defined by $[\alpha]^\beta = min\{\alpha, \beta\}$. For any functional revision $*$, $\kappa^*_{A,\alpha}$ is the unique κ' such that $\kappa * (A, \alpha)\kappa'$. For functional revisions, the truth conditions can be simplified as follows. If $*$ is a functional revision, then $\mathcal{M}, \kappa, w \models [\varphi, \alpha]\psi$ iff $\mathcal{M}, \kappa^*_{[\![\varphi]\!]_V, \alpha}, w \models \psi$.

I define two functional revision operators, Spohn revision (also called "A, α-conditionalization", defined in Spohn (1988)) $*^{\rightarrow}$ and Shenoy revision (also called "Shenoy conditionalization", defined in Shenoy (1991)) $*^{\uparrow}$ as follows. For any proposition A, $\alpha \leq \lambda$, $\kappa \in K^\lambda_W$, $w \in W$,

$$\text{Spohn: } \kappa^{*^{\rightarrow}}_{A,\alpha}(w) = \begin{cases} \kappa(w) - \kappa(A) & \text{if } w \in A \\ [\alpha + (\kappa(w) - \kappa(W \backslash A))]^\lambda & \text{else} \end{cases}$$

$$\text{Shenoy: } \kappa^{*^{\uparrow}}_{A,\alpha}(w) = \begin{cases} \kappa(w) - m & \text{if } w \in A \\ [(\kappa(w) + \alpha) - m]^\lambda & \text{else} \end{cases}$$

Where $m = min\{\kappa(A), \alpha\}$.

4 Relations Between BDDL and RBDDL

RBDDL provides a lot of possibilities of specifying revision operators. Some of them take full advantage of the richness of structure ranking functions provide, and some don't. Some are insensitive to the quantitative information ranking functions convey, and just depend on the order of worlds. Such revision behaviour can already be captured in BDDL. In this section, I want to make this difference precise. This in turn will enable a precise account of which revision operators require a ranking semantics, and which do not.

I start by defining a function $\xi : K^\lambda_W \rightarrow S^\lambda_W$ that assigns to each ranking function the system of spheres it represents. Earlier I compared ranking functions to system of spheres with labels – carrying on this simile, ξ is an *ordinal label eraser*. For any ranking function κ:

$$\xi(\kappa) = \{\{w \in W | \kappa(w) \leq \alpha\} | \alpha \in Ord\}$$

It is clear from the ordinal structure of ranking functions and systems of spheres that ξ is well-defined. As for each system of spheres, there are several ranking functions which represent it, I conversely define a function $\zeta : S^\lambda_W \rightarrow \wp(K^\lambda_W)$ that assigns to each system of spheres the set of ranking functions that represent it. For any system of spheres s:

$$\zeta(s) = \{\kappa \in K^\lambda_W | \xi(\kappa) = s\}$$

Note that ξ and ζ are defined relative to the choice of W and λ. Earlier, I mentioned that ranking functions are at least as informative as systems of spheres. In the formal setting of this section, this can be put in the following theorem:

Theorem 1. *For any system of spheres $s \in S^\lambda_W$ there is a ranking function $\kappa \in K^\lambda_W$ such that $\xi(\kappa) = s$.*

Proof. Let $s \in S^\lambda_W$. For each ordinal number α, let s_α be the αth sphere by transfinite recursion: $s_\alpha = \bigcap(s \backslash \{s_\beta | \beta < \alpha\})$. Now let κ be the ranking function that assigns the each world the lowest ordinal number such that the world is an element of the sphere with this number. That is, κ is defined as the ranking function such that for every world w, $\kappa(w) = min\{\alpha \in Ord | w \in s_a\}$. It is clear from this construction that $\kappa \in K^\lambda_W$ and $\xi(\kappa) = s$. □

Towards the goal of making revisions of ranking functions and systems of spheres easier to compare, formal notions of comparability and equivalence are defined:

A revision $*$ on a set of ranking functions K and a revision $\star$ on a set of systems of spheres S are *comparable* iff for every $\kappa \in K$, $\xi(\kappa) \in S$ and for every $s \in S$, there is a $\kappa \in K$ such that $\xi(\kappa) = s$.

Let $*$ be a revision on a set of ranking functions K and $\star$ be a revision on a set of systems of spheres S that are comparable. $*$ and $\star$ are *equivalent*, written $* \sim \star$, iff for any valuation V, $\kappa \in K$, $w \in W$, and formula ψ,

$$\langle W, K, *, V\rangle, \kappa, w \models \psi \text{ iff } \langle W, S, \star, V\rangle, \xi(\kappa), w \models \psi$$

From this definition, it follows that any model using a revision with an equivalent revision of the other type can be converted into a model of that type. To go from a BDDL to an RBDDL model, replace the revision with an equivalent ranking revision and replace the set of systems of spheres with the set of ranking functions the new revision is defined on. The belief state of the agent can now be given by any ranking function that represents the system of spheres (the existence of which is secured by the fact that the revisions are comparable). Going from an RBDDL model to a BDDL model is strictly analogous, and the new belief state can now be obtained by use of the ξ function.

Another application of the notion of equivalence between revisions is this: If two revisions are equivalent, then for every formula that is true in a certain model with this revision, there is a model of the other type with the equivalent revision that makes the formula true as well. This means that they share the same axiomatisations.

Although a useful notion of equivalence between revisions, the condition it was defined with is hard to use: How should one go about when one wants to find out whether two given revisions are equivalent? This is made more easy by the following theorem:

Theorem 2. *Let $*$ be a revision on a set of ranking functions K and $\star$ be a revision on a set of systems of spheres S that are comparable. If for any proposition A, $\alpha \leq \lambda$, and $\kappa_1, \kappa_2 \in K$, $\kappa_1 * (A, \alpha)\kappa_2$ iff $\xi(\kappa_1) \star (A, \alpha)\xi(\kappa_2)$ then $* \sim \star$.*

To prove this theorem, I first prove the following lemma:

Lemma 1. *For any ranking function κ, $core(\kappa) = \bigcap \xi(\kappa)$.*

Proof of lemma 1. $\bigcap \xi(\kappa) = \bigcap \{\{w \in W | \kappa(w) \leq \alpha\} | \alpha \in Ord\}$
$= \{w \in W | \kappa(w) = 0\} = core(\kappa)$ □

Proof of theorem 2. Let V be any valuation, κ any ranking function on K, ψ any formula and w any world. Let $\mathcal{M}^* = \langle W, K, *, V\rangle$ and $\mathcal{M}^\star = \langle W, S, \star, V\rangle$. Show $* \sim \star$ by induction on formulas:

($\psi = \varphi$, $\psi = \neg\psi'$, or $\psi = (\psi_1 \wedge \psi_2)$). If ψ is pure Boolean, a negation or a conjunction, the truth definitions are identical.

($\psi = B\varphi$). Show that $\mathcal{M}^*, \kappa, w \models B\varphi$ iff $\mathcal{M}^\star, \xi(\kappa), w \models B\varphi$. This is equivalent to $core(\kappa) \subseteq [\![\varphi]\!]_V$ iff $\bigcap \xi(\kappa) \subseteq [\![\varphi]\!]_V$, which holds, as $core(\kappa) = \bigcap \xi(\kappa)$, which was shown in Lemma 1.

($\psi = [\varphi, \alpha]\psi'$). Show that $\mathcal{M}^*, \kappa, w \models [\varphi, \alpha]\psi'$ iff $\mathcal{M}^\star, \xi(\kappa), w \models [\varphi, \alpha]\psi'$, which is equivalent to $\forall \kappa' \in K(\kappa * ([\![\varphi]\!]_V, \alpha)\kappa' \rightarrow \mathcal{M}^*, \kappa', w \models \psi')$ iff $\forall s \in S(\xi(\kappa) \star ([\![\varphi]\!]_V, \alpha)s \rightarrow \mathcal{M}^\star, s, w \models \psi')$. I show only direction $\Rightarrow$, the other follows analogously. Take any $s \in S$. If $\xi(\kappa) \star ([\![\varphi]\!]_V, \alpha)s$, then there is a κ' such that $\kappa * ([\![\varphi]\!]_V, \alpha)\kappa'$ and $\xi(\kappa') = s$ (the existence of κ' is guaranteed by comparability of the revisions). It follows that $\mathcal{M}^*, \kappa', w \models \psi'$, and, using the induction premise, that $\mathcal{M}^\star, \xi(\kappa'), w \models \psi'$. □

I will put this theorem to use right away. As ranking functions are more fine-grained than systems of spheres, any BDDL revision has an equivalent RBDDL revision. This is stated precisely in the next theorem:

Theorem 3. *Let $\star$ be a revision defined on a set of systems of spheres S. There is a revision that is equivalent to $\star$.*

Proof. Let $\star$ be any revision defined on a set of systems of spheres S. Let $*$ be a revision defined on $K = \bigcup\{\zeta(s) | s \in S\}$. For any ranking functions $\kappa_1, \kappa_2 \in K$, proposition A, and $\alpha \leq \lambda$, let $\kappa_1 * (A, \alpha)\kappa_2$ iff $\xi(\kappa_1) * (A, \alpha)\xi(\kappa_2)$. By theorem 2, the revisions are equivalent. □

Theorem 2 states that if a certain condition is met, two revisions are equivalent. One might wonder whether this can be strengthened to the proposition that the revisions are equivalent iff the condition is met. The following theorem shows that it can't, as the language is in a sense not expressive enough to make certain distinctions between ranking functions:

Theorem 4. *There is a revision $*$ on a set of ranking functions K and a revision $\star$ on a set of systems of spheres S that are comparable such that $* \sim \star$, but not for any $\kappa_1, \kappa_2 \in K$, proposition A, and $\alpha \leq \lambda$, $\kappa_1 * (A, \alpha)\kappa_2$ iff $\xi(\kappa_1) \star (A, \alpha)\xi(\kappa_2)$.*

Proof. Proof by construction of an example. Let κ be a ranking function and $\hat{\kappa}$ be the ranking function such that for any world w, $\hat{\kappa}(w) = \kappa(w) * 2$. It is clear that $\xi(\kappa) = \xi(\hat{\kappa})$. Define $K = \{\kappa, \hat{\kappa}\}$, $S = \{\xi(\kappa)\}$, and the following revisions on them. For any proposition A and $\alpha \leq \lambda$, $*(A, \alpha) = \{\langle\kappa, \kappa\rangle, \langle\hat{\kappa}, \hat{\kappa}\rangle\}$ and $\star(A, \alpha) = \{\langle\xi(\kappa), \xi(\kappa)\rangle\}$. Prove that these revisions fulfill the conditions described:

$(* \sim \star)$. Show that for any valuation V, $\kappa \in K$, formula ψ, and world w, $\langle W, K, *, V\rangle, \kappa, w \models \psi$ iff $\langle W, S, \star, V\rangle, \xi(\kappa), w \models \psi$. It is easy to see that this is true by induction: If ψ is pure Boolean, or of the form $\neg\psi'$, $(\psi' \wedge \psi'')$, or $B\psi'$, this is trivially true (see proof of theorem 2). If ψ is of the form $[\varphi, \alpha]\psi'$, the truth of the equivalence can be seen as follows: No revision changes the belief state, so $[\varphi, \alpha]\psi'$ is true iff ψ' is true, and the equivalence holds for ψ' on account of the induction.

(Not for all $\kappa_1, \kappa_2 \in K$, proposition A, and $\alpha \leq \lambda$, $\kappa_1 * (A, \alpha)\kappa_2$ iff $\xi(\kappa_1) \star (A, \alpha)\xi(\kappa_2)$). Example: for any proposition A and $\alpha \leq \lambda$, $\kappa * (A, \alpha)\hat{\kappa}$ is false, but $\xi(\kappa) \star (A, \alpha)\xi(\hat{\kappa})$ is true. □

A, α-conditionalization was introduced earlier as a revision operator of ranking theory that was supposed to overcome defects shared by all revision operators defined on systems of spheres. Therefore, one should expect $*^{\rightarrow}$ not to have an equivalent revision. This is indeed the case, provided W is not trivial, in the sense of having at least three elements:

Theorem 5. *If the set of worlds W has at least three elements, then there is no revision defined on systems of spheres that is equivalent to $*^{\rightarrow}$.*

Proof. Proof by construction of an example. Let a, b, and c be distinct elements of W. Let κ_1 and κ_2 be ranking functions such that $\kappa_1(a) = \kappa_2(a) = 0$, $\kappa_1(b) = \kappa_2(b) = 1$, $\kappa_1(c) = 2$, $\kappa_2(c) = 3$, and for any other world $w \in W \backslash \{a, b, c\}$, $\kappa_1(w) = \kappa_2(w) = 4$. Applying $*^{\rightarrow}$, it turns out that $({\kappa_1}^{*^{\rightarrow}}_{\{b,c\},1})^{*^{\rightarrow}}_{\{a,c\},1}(c) = 0$, and $({\kappa_2}^{*^{\rightarrow}}_{\{b,c\},1})^{*^{\rightarrow}}_{\{a,c\},1}(c) = 1$. Using bc and ac as propositional variables, then in a model $\mathcal{M}$ with $*^{\rightarrow}$ as its revision containing a valuation V such that $V(bc) = \{b, c\}$ and $V(ac) = \{a, c\}$, $\mathcal{M}, \kappa_1, w \models [bc, 1][ac, 1]Bc$ and $\mathcal{M}, \kappa_2, w \models \neg[bc, 1][ac, 1]Bc$. As $\xi(\kappa_1) = \xi(\kappa_2)$, there cannot be a revision defined on systems of spheres that is equivalent to $*^{\rightarrow}$. □

Theorem 3 says that any BDDL revision has an equivalent RBDDL revision. That the converse is not true follows immediately from theorem 5, as stated in the following corollary:

Corollary 1. *There is a revision defined on a set of ranking functions that has no equivalent revision.*

5 The Ordinal Limit

Originally, ranking functions were introduced as certain functions from possible worlds to ordinals in Spohn (1988). In other publications, the codomain of ranking functions has been specified differently. For example, in Spohn

(2008) and Huber (2007) the natural numbers plus infinity are used, in Spohn (1999) it is the natural numbers, and Aucher (2005) uses the natural numbers up to some fixed finite limit. Some publications even use the positive real numbers and infinity (e.g. Hild and Spohn (2008)), but I won't consider this choice here.

In most of these papers, the reason for choosing whichever codomain is chosen is a tradeoff between generality and technical convenience. However, as ranking theory is a theory of belief revision, choosing between different codomains of ranking functions is a philosophical choice, and as it effects philosophically relevant differences in the theory, it needs philosophical argumentation. As the theory should be as simple as possible, the codomain of ranking functions should be chosen as small as possible, provided this does not introduce any unwanted consequences. This line of reasoning is suggested in Spohn (1988, fn 16):

> It would be a natural idea to restrict the range of OCFs [ordinal conditional functions, another term for ranking functions, P.F.] to the set of natural numbers. In fact, much of the following could thereby be simplified since usual arithmetic is simpler than the arithmetic of ordinals. For the sake of formal generality I do not impose this restriction. But larger ranges may also be intuitively needed. For example, it is tempting to use OCFs with larger ranges to represent the stubbornness with which some beliefs are held in the face of seemingly arbitrarily augmentable counter-evidence.

In Spohn (2008), he comments on this footnote:

> I have never elaborated this remark and do not know whether it can be done. So, not being faced with good applications we better stay content with the simpler notion. [of using natural numbers plus infinity as the codomain for ranking functions, P.F.] For, there would be a price to pay for the generality. [of using ordinals as the codomain for ranking functions, P.F.] We shall soon start calculating with the ranks, and then we would have to engage into ordinal arithmetic, which is far less well-behaved than ordinary arithmetic; for instance, ordinal addition and multiplication are no longer commutative. This would be a severe handicap.

In the following, I want to start investigations in this area by formulating one argument using the logic defined earlier. This is intended to further illustrate the logic's usefulness. However, I want to stress that the argument could as well be formulated without it.

In the logic given, a restriction of the codomain of ranking functions is simply a choice of the limiting ordinal λ. In the argument to follow, I will formally argue against the use of the ordinals up to and including ω, which is the equivalent of using the natural numbers plus infinity. Therefore, I will consider the choice $\lambda = \omega$. The argument will use Shenoy revision, as the situation that will be described is one in which revisions may be unsuccessful. It works as follows: A certain epistemic situation is described such that a good theory of belief revision should intuitively be able to describe it. I then formally prove that this is not possible if $\lambda = \omega$. After that, I will look at options of extending the argument to higher ordinals.

The situation is this: You strongly disbelieve some proposition φ, e.g. that there are 10 spatial dimensions. If someone unbeknown to you comes along and tells you that φ, you will not believe it. Indeed, your conviction is so strong that no matter how many strangers come along and tell you φ, you do not change your mind. This is not because you generally completely mistrust strangers; should they tell you something you have no beliefs about, you will readily believe it. Although strong, your disbelief in φ is not dogmatic. If a well-known theoretical physicist comes along and tells you that recent advances in string theory have established φ beyond doubt, or you read this in a well-known journal of theoretical physics, you will successfully revise your beliefs, and believe φ afterwards. Although you then believe φ, you disbelieve some proposition φ' that implies φ. For example, you may disbelieve the conjunction that there are 10 spatial dimensions and that Hadrian succeeded Trajan as Emperor of Rome, as you may have false beliefs about the Emperors.

Intuitively, I think this is a perfectly reasonable story. Even if you disagree in this particular instance, I think that it is even more plausible to hold that an account of belief revision should not *prima facie* exclude all cases of this structure. I now describe how this situation is modelled in ranking theory and the prove that this cannot be done if $\lambda \leq \omega$.

In the situation described, not all revisions are successful, so Shenoy revision is used. The parameter of Shenoy revision indicates how credible the information is. Which parameter should be used for information given by a complete stranger? It may not be 0, as Shenoy revision with parameter 0 is only successful in the trivial case in which you already believe the proposition you revise with. I use 1, as using a higher number would not change matters in any important way, which will be apparent in the following. To represent the story told, a model is needed that valuates φ and φ' such that any φ'-world is a φ-world, and that makes false any statement of the form

$[\varphi, 1]^n B\varphi$ for all $n \in \mathbf{N}$ but makes true the statement $[\varphi, \gamma](B\varphi \wedge B\neg\varphi')$ for some parameter γ. The following formal investigations prove that there is no such model if $\lambda \leq \omega$.

Theorem 6. *If $\lambda \leq \omega$, then for any non-empty set W, $K \subseteq K_W^\lambda$, valuation V on W, $\kappa \in K$, $w \in W$, $\varphi, \varphi' \in pB$ such that $[\![\varphi']\!]_V \subseteq [\![\varphi]\!]_V$ and $[\![\varphi']\!]_V \neq \emptyset$, if for all $n \in \mathbf{N}$,*
$\langle W, K, *^\uparrow, V\rangle, \kappa, w \models \neg[\varphi, n]B\varphi$
then for all $\gamma \leq \lambda$:
$\langle W, K, *^\uparrow, V\rangle, \kappa, w \models \neg[\varphi, \gamma](B\varphi \wedge B\neg\varphi')$

Proof. Let $\lambda \leq \omega$. Take any non-empty set W, $K \subseteq K_W^\lambda$, valuation V on W, $\kappa \in K$, $w \in W$, $\varphi, \varphi' \in pB$ such that $[\![\varphi']\!]_V \subseteq [\![\varphi]\!]_V$ and $[\![\varphi']\!]_V \neq \emptyset$, and for all $n \in \mathbf{N}$, $\langle W, K, *^\uparrow, V\rangle, \kappa, w \models \neg[\varphi, n]B\varphi$.

First, prove that $\kappa([\![\varphi]\!]_V) = \omega$, which implies that $\lambda = \omega$. Assume for contradiction that $\kappa([\![\varphi]\!]_V) \neq \omega$, then there is a $n \in \mathbf{N}$ such that $\kappa([\![\varphi]\!]_V) = n$. Let s be a world such that $s \notin [\![\varphi]\!]_V$. Then $\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, n+1}(s) = [(\kappa(s) + (n+1)) - min\{n+1, \kappa([\![\varphi]\!]_V)\}]^\lambda = [\kappa(s) + 1]^\lambda \geq 1$. Therefore $core(\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, n+1}) \subseteq [\![\varphi]\!]_V$, so $\langle W, K, *^\uparrow, V\rangle, \kappa, w \models [\varphi, n+1]B\varphi$ in contradiction to the assumptions. ↯

For the next step, assume for contradiction that there is a $\gamma \leq \omega$ such that $\langle W, K, *^\uparrow, V\rangle, \kappa, w \models [\varphi, \gamma](B\varphi \wedge B\neg\varphi')$. $\gamma < \omega$ immediately contradicts the assumptions, so consider $\gamma = \omega$. $\langle W, K, *^\uparrow, V\rangle, \kappa, w \models [\varphi, \omega](B\varphi \wedge B\neg\varphi')$ is equivalent to $\langle W, K, *^\uparrow, V\rangle, \kappa^{*^\uparrow}_{[\![\varphi]\!]_V, \omega}, w \models (B\varphi \wedge B\neg\varphi')$, which implies $core(\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, \omega}) \subseteq W \backslash [\![\varphi']\!]_V$, which will be contradicted in the following. As $\lambda = \omega$, $[\![\varphi]\!]_V \leq \omega$, so from the definition of $*^\uparrow$, it follows that for all $w \in W$:

$$\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, \omega}(w) = \begin{cases} \kappa(w) - \kappa([\![\varphi]\!]_V) & \text{if } w \in [\![\varphi]\!]_V \\ [(\kappa(w) + \omega) - \kappa([\![\varphi]\!]_V)]^\omega & \text{else} \end{cases}$$

As proven above, $\kappa([\![\varphi]\!]_V) = \omega$. Also, $\lambda = \omega$, so for any $t \in [\![\varphi]\!]_V$, $\kappa(t) = \omega$, and therefore $\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, \omega}(t) = \kappa(t) - \kappa([\![\varphi]\!]_V) = \omega - \omega = 0$. Therefore, as $[\![\varphi']\!]_V \subseteq [\![\varphi]\!]_V$, for any $v \in [\![\varphi']\!]_V$, $\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, \omega}(v) = 0$. As $[\![\varphi']\!]_V \neq \emptyset$, it follows that $core(\kappa^{*^\uparrow}_{[\![\varphi]\!]_V, \omega}) \nsubseteq W \backslash [\![\varphi']\!]_V$.↯ □

Lemma 2. *For any non-empty set W, $K \subseteq K_W^\lambda$, valuation V on W, $\kappa \in K$, $w \in W$, $n \in \mathbf{N}$, $\varphi \in pB$, $\psi \in F_B$, $\langle W, K, *^\uparrow, V\rangle, \kappa, w \models [\varphi, n]\psi \leftrightarrow [\varphi, 1]^n\psi$.*

Proof. Take any non-empty set W, $K \subseteq K_W^\lambda$, valuation V on W, $\kappa \in K$, $w \in W$, $n \in \mathbf{N}$, $\varphi \in pB$, $\psi \in F_B$. Let $\mathcal{M} = \langle W, K, *^\uparrow, V\rangle$. Show lemma by induction on n.

($n = 0$). Show $\mathcal{M}, \kappa, w \models [\varphi, 0]\psi \leftrightarrow \psi$, that is, $\mathcal{M}, \kappa^{*\uparrow}_{[\![\varphi]\!]_V,0}, w \models \psi$ iff $\mathcal{M}, \kappa, w \models \psi$. This holds, as for any proposition A, $\kappa^{*\uparrow}_{A,0} = \kappa$.

($n = n + 1$). Show $\mathcal{M}, \kappa, w \models [\varphi, n + 1]\psi \leftrightarrow [\varphi, 1]^{n+1}\psi$. It follows from the induction premise that this is equivalent to $\mathcal{M}, \kappa^{*\uparrow}_{[\![\varphi]\!]_V,n+1}, w \models \psi$ iff $\mathcal{M}, \kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}, w \models [\varphi, 1]\psi$, which is equivalent to $\mathcal{M}, \kappa^{*\uparrow}_{[\![\varphi]\!]_V,n+1}, w \models \psi$ iff $\mathcal{M}, (\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n})^{*\uparrow}_{[\![\varphi]\!]_V,1}, w \models \psi$. Show this by showing that

$$(\dagger) \quad \kappa^{*\uparrow}_{[\![\varphi]\!]_V,n+1} = (\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n})^{*\uparrow}_{[\![\varphi]\!]_V,1}$$

For the following, observe that $\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}([\![\varphi]\!]_V) = min\{\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}(v)|v \in [\![\varphi]\!]_V\} = min\{\kappa(v) - min\{\kappa([\![\varphi]\!]_V), n\}|v \in [\![\varphi]\!]_V\}$. Notice also that as n is natural, if $\kappa([\![\varphi]\!]_V) \leq n$, so is $\kappa([\![\varphi]\!]_V)$. Prove ($\dagger$) by distinguishing two cases, which will each be split up in two cases as well.

(Case 1: $w \in [\![\varphi]\!]_V$). ($\dagger$) $\Leftrightarrow$
$\kappa(w) - min\{\kappa([\![\varphi]\!]_V), n + 1\} = \kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}(w) - min\{\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}([\![\varphi]\!]_V), 1\} \Leftrightarrow$
$\kappa(w) - min\{\kappa([\![\varphi]\!]_V), n + 1\} = (\kappa(w) - min\{\kappa([\![\varphi]\!]_V), n\})$
$-min\{min\{\kappa(v) - min\{\kappa([\![\varphi]\!]_V), n\}|v \in [\![\varphi]\!]_V\}, 1\}$
(Case 1a: $\kappa([\![\varphi]\!]_V) \leq n$). $\ldots \Leftrightarrow$
$\kappa(w) - \kappa([\![\varphi]\!]_V) = (\kappa(w) - \kappa([\![\varphi]\!]_V)) - min\{min\{\kappa(v) - \kappa([\![\varphi]\!]_V)|v \in [\![\varphi]\!]_V\}, 1\}$
$\Leftrightarrow \kappa(w) - \kappa([\![\varphi]\!]_V) = (\kappa(w) - \kappa([\![\varphi]\!]_V)) - min\{0, 1\}$✓
(Case 1b: $\kappa([\![\varphi]\!]_V) > n$). $\ldots \Leftrightarrow$
$\kappa(w) - (n + 1) = (\kappa(w) - n) - min\{min\{\kappa(v) - n|v \in [\![\varphi]\!]_V\}, 1\} \Leftrightarrow$
$\kappa(w) - (n + 1) = (\kappa(w) - n) - 1$✓

(Case 2: $w \notin [\![\varphi]\!]_V$). ($\dagger$) $\Leftrightarrow$
$[(\kappa(w) + (n + 1)) - min\{\kappa([\![\varphi]\!]_V), (n + 1)\}]^\lambda$
$= [(\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}(w) + 1) - min\{\kappa^{*\uparrow}_{[\![\varphi]\!]_V,n}([\![\varphi]\!]_V), 1\}]^\lambda \Leftrightarrow$
$[(\kappa(w) + (n + 1)) - min\{\kappa([\![\varphi]\!]_V), (n + 1)\}]^\lambda = [([(\kappa(w) + n) -$
$min\{\kappa([\![\varphi]\!]_V), n\}]^\lambda + 1) - min\{min\{\kappa(v) - min\{\kappa([\![\varphi]\!]_V), n\}|v \in [\![\varphi]\!]_V\}, 1\}]^\lambda$
(Case 2a: $\kappa([\![\varphi]\!]_V) \leq n$). $\ldots \Leftrightarrow$
$[(\kappa(w) + (n + 1)) - \kappa([\![\varphi]\!]_V)]^\lambda = [([(\kappa(w) + n) - \kappa([\![\varphi]\!]_V)]^\lambda + 1)$
$-min\{min\{\kappa(v) - \kappa([\![\varphi]\!]_V)|v \in [\![\varphi]\!]_V\}, 1\}]^\lambda \Leftrightarrow$
$[(\kappa(w) + (n + 1)) - \kappa([\![\varphi]\!]_V)]^\lambda = [([(\kappa(w) + n) - \kappa([\![\varphi]\!]_V)]^\lambda + 1)]^\lambda$✓
(Case 2b: $\kappa([\![\varphi]\!]_V) > n$). $\ldots \Leftrightarrow$
$[(\kappa(w) + (n + 1)) - (n + 1)]^\lambda = [([(\kappa(w) + n) - n]^\lambda + 1) - 1]^\lambda$✓ □

From the preceding theorem and lemma, it immediately follows that:

Corollary 2. *If $\lambda \leq \omega$, then for any non-empty set W, $K \subseteq K^\lambda_W$, valuation V on W, $\kappa \in K$, $w \in W$, $\varphi, \varphi' \in pB$ such that $[\![\varphi']\!]_V \subseteq [\![\varphi]\!]_V$ and $[\![\varphi']\!]_V \neq \emptyset$, if for all $n \in \mathbf{N}$,*

$\langle W, K, *^{\uparrow}, V\rangle, \kappa, w \models \neg[\varphi, 1]^n B\varphi$
then for all $\gamma \leq \lambda$*:*
$\langle W, K, *^{\uparrow}, V\rangle, \kappa, w \models \neg[\varphi, \gamma](B\varphi \wedge B\neg\varphi')$

With this, it is proven that for the situation given, $\lambda \leq \omega$ is not sufficient. In the next paragraphs, I will informally investigate how far this argument can reasonably be extended.

It is not hard to see that the problem posed is already repaired by using $\lambda = \omega + 1$. This move is patently *ad hoc*, and one can argue that it is unsatisfying as follows: There must not only be one φ' that implies φ, but a series of pure Boolean formulas of order type ω such that each implies the next and revising with one does not entail not disbelieving all that imply it. Such a demand is clearly not excessive as any conjunction implies its conjuncts. It is straightforward that this means that λ may not be smaller than $\omega * 2$. This is a very mild extension; it seems to me that whoever accepts the argument at all must accept it as well.

One can also be bolder and ask not for one revision parameter representing an authority, but a series of higher and higher authorities of order type ω. Imagine a company in which there are some number of levels of management. For each level, there is a proposition such that if any number of managers below the level tell it to you, you will not believe it. If someone of that level or higher tells it to you, you do believe it. The first authority would get parameter ω, the second $\omega * 2$, the third $\omega * 3$ and so on, so λ may not be smaller than ω^2.

This extension may already seem to put some strain on the intuitions used so far, but the following goes further. One can ask: why stop at one series of authorities? Couldn't there be a series *plus* an authority that is more credible than any of the authorities considered so far? This *hyper*-authority would get credibility ω^2, and as before, there could not only be one, but a whole series of order type ω. This would then push the limit ordinal from ω^2 to ω^3. The game continues on another level: Why not have a series of levels of *hyper*-authorities of order type ω? This would mean $\lambda = \omega^\omega$.

By now, it becomes difficult to picture the structure of ordinals and apply it to epistemic situations. But one can push even further by arguing as follows: Belief revision theory is the theory of rational revision of belief in general. At least as a philosophical theory, it should not be subjected to constraints that are based merely on the limited ability of our human minds, nor the contingent facts of this world.

So one can carry on and conceive of *gods* (with respect to their epistemic status) having credibility ω^ω, series of higher and higher gods (with credibility $\omega^{\omega+n}$), series of *hyper*-gods (with credibility $\omega^{\omega*2+n}$), series of levels of *hyper*-gods (with credibility $\omega^{\omega * n}$), and even series of *meta-hyper*-gods (with credibility ω^{ω^n}), ending with $\lambda = \omega^{\omega^\omega}$.

The next (and here, last) jump to the next level brings the series of *hyper*-authorities, *meta-hyper*-gods, and yet unnamed further classes, leading to the series $\omega, \omega^\omega, \omega^{\omega^\omega}, \ldots$. Its limit is called ε_0, and it is with this ordinal that I end this line of extending the argument. It is not to be expected that the intuitions evoked in the story I told extend beyond this point.

Of course, this whole argument depends on the assumption of Shenoy revision, so one could argue that rather than accepting the fact that λ must be chosen higher than ω, one should look for better accounts of revision. This is a sensible suggestion, but there are two things to bear in mind: First, the burden of showing the existence of such better accounts rests on the objector against this argument. Second, although the proofs essentially use properties of Shenoy revision, I expect them to be generalizable to classes of revisions satisfying certain properties. Should these properties turn out to be ones we intuitively require from non-successful revision operators, the objection would be rebutted. Of course, here, the burden of proof lies on the defender of the argument.

6 Ranking Theory of Higher-Order Beliefs and Revisions

The approaches using systems of spheres as well as the ones using ranking functions discussed, whether they are put in a logic or not, are only concerned with beliefs about factual propositions and revisions with factual propositions. By *factual propositions*, I mean propositions that do not talk about the beliefs of agents or about what would happen if the agent would revise her belief state with a certain proposition. Also, the discussion so far was only concerned with the epistemic status of a single agent. Stating that Anne believes that Paul believes that φ is not possible. One might object against this view of the theories discussed, and hold that although these frameworks lack the formal flexibility to explicitly express such things, we can the interpret propositional letters and possible worlds in any way we please, e.g. the set of worlds $V(p)$ could be taken to be the worlds in which Paul believes that φ. Letting Anne be the agent, a BDDL representation of Anne's belief would just be Bp.

Of course, one *can* interpret propositional letters to represent such meta-information, but I hold that one shouldn't, because if one does, AGM as well as ranking theory of belief revision will turn out to be utterly implausible.

Here's a very simple example: Assume you neither believe that John is bald nor that John isn't bald. Suppose you learn that John is bald. How does AGM revision deal with this example?

I use the propositional letter p for John's being bald, and q for the agent's believing John to be bald. The following diagram visualizes the AGM modelling:

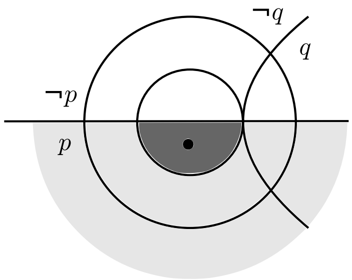

The concentric circles represent the initial system of spheres, the black dot the actual world. The two strokes each divide the space of worlds in two sections, the worlds where p or $\neg p$ holds, and the worlds where q or $\neg q$ holds. The actual world is in the quadrant in which p and $\neg q$ hold, as John is bald, but the agent does not believe it. Also, the agent does not believe John not to be bald, but she does believe not to believe that John is bald. So far so good. p is true in the area shaded light gray, so after revision with this proposition, the dark gray area becomes most plausible. As expected, after the revision, the agent believes p, so she believes that John is bald. The problem is now: $\neg q$ is still true in the actual world, and $\neg q$ means that the agent does *not* believe that John is bald! The problem is that the system has two representations of the agent's belief in John's being bald, and in this case, they contradict each other. It is an simple exercise to go through the steps of this argument if ranking theory is used instead of AGM theory – they share the same problems in this area.

One approach would be to hold that any proposition expressing higher-order belief or revision can be reduced to one expressing first-order states by reduction axioms. For iterated belief, one could try to defend the controversial axioms of positive and negative introspection ($B\varphi \leftrightarrow BB\varphi$ and $\neg B\varphi \leftrightarrow B\neg B\varphi$). However, this approach is little promising for revision with propositions containing belief ascriptions, as argued in Baltag and Smets (2008, pp 38–40). Even worse is the case of multi-agent belief. To represent beliefs of multiple agents, one usually employs multi-modal languages containing a belief operator B_a for every agent a. Obviously, the truth-values of propositions like $B_aB_c\varphi$ and $B_bB_c\varphi$ are completely independent, and no reduction axiom can be given for such cases in general.

A more promising approach that tackles such problems starts with a logic of higher-order multi-agent belief, and adds the notion of belief revision to this. This is an area of current research, and it goes by the name of "belief revision in dynamic logic", where "epistemic" or "doxastic" may be inserted before "logic". A terminological note on the difference between dynamic doxastic logic (DDL) and dynamic epistemic logic (DEL): "epistemic logic" is the name usually applied to modal logics in which the modal operators are understood to represent knowledge or belief. "Doxastic logic" is sometimes used to refer to these logics if the modal operators are understood to express belief. The expressions "doxastic" and "epistemic" are derived from the greek words for belief and knowledge, viz. "doxa" and "episteme". "Dynamic modal logic" is a term used to refer to modal logics with dynamic operators. The names "DDL" and "DEL" somewhat continue these conventions, but not exactly. "DDL" is used for logics in the tradition of Segerberg, similar to BDDL. "DEL" in turn is used for epistemic logics enriched with any of a variety of dynamic operators, in which the modal operator can be interpreted as belief as well as knowledge. As logics of belief revision are never based on Kripke models, one often doesn't talk about belief revision in DEL (as does the title of the present thesis), but just about *dynamic logics*. Still, this research can properly be seen as a further generalization of DEL approaches. It is to be hoped that as the research program of formulating theories of belief revision in terms of dynamic logics matures, terminological unity will emerge. I end this digression with a suggestion in this direction from Leitgeb and Segerberg (2007, p 189):

> We predict that the two research programmes of DDL and DEL will merge in the long run into the single logical endeavour of DBC: *dynamic logics of belief change*.

6.1 Belief Revision in Dynamic Logic

In the past few years, a number of proposals have appeared that try to model belief revision in logics derived from DEL, for example (Baltag and Smets, 2008, forthcoming), van Benthem (2007), and van Ditmarsch (2005). In the following, I briefly highlight some ideas from van Benthem (2007).

Van Benthem starts with a logic of belief based on plausibility orders. Here, a plausibility order $\leq$ is an order on the set of worlds, and the intended interpretation is that $w \leq w'$ iff w is considered to be at least as plausible as w'. Plausibility orders generalize systems of spheres, as every system of spheres can be mapped to a unique wellorder on the set of world. As van Benthem wants to model beliefs of several agents, the plausibility orders are indexed to a set of agents, and as he wants to model higher-order beliefs, they are also indexed to the set of worlds.

Syntactically, this first logic is just a multi-agent doxastic modal language, that is, propositional letters are formulas, and for any formulas φ, φ', and agent index i, $\neg\varphi$, $\varphi \wedge \varphi'$ and $B_i\varphi$ are formulas. For any set of agents I, a model is a structure $\mathcal{M} = \langle W, \{\leq_{i,s}\}_{i \in I, s \in W}, V\rangle$, where $\leq_{i,s}$ is the plausibility order of i in s. Truth is now defined relative to a model and a world – note that we don't have to specify plausibility orders representing epistemic states, as we had to in BDDL, as these are already specified in the model. Truth for factual propositions $(p, \neg\varphi, \varphi \wedge \varphi)$ are straightforward as in BDDL. Truth of belief statements also strictly similar and defined by

$$\mathcal{M}, s \models B_i\varphi \text{ iff } \forall t \in W(t \in min_{\leq_{i,s}} W \rightarrow \mathcal{M}, t \models \varphi)$$

We see that the reason why van Benthem can define higher-order beliefs is that in contrast to the semantics of BDDL, models here specify a plausibility order for each world. He now adds a dynamic operator to this logic representing revision. The expression $[\Uparrow \varphi]\psi$ is used to state that after revising with φ, ψ holds. Truth for this is defined by:

$$\mathcal{M}, s \models [\Uparrow \varphi]\psi \text{ iff } \mathcal{M} \Uparrow \varphi, s \models \psi$$

where $\mathcal{M} \Uparrow \varphi$ is obtained from $\mathcal{M}$ by replacing each plausibility relation $\leq$ with $\leq\Uparrow \varphi$, which can be defined as follows: For any $w, w' \in W$, $w \leq\Uparrow \varphi w'$ iff either φ is true in w, and not in w', or $w \leq w'$. $\Uparrow$ is here just used as one example for a revision operator, others can analogously be defined.

The case sketched is just the simplest form of a dynamic logic of belief revision, as belief revision is represented by a special dynamic operator. Generalizing this approach, a semantic representation of changes to the plausibility order can be devised, and a multi-purpose dynamic operator be introduced, which is parametrized by syntactic representations of these changes. Belief revision is then just one (or some, if a range of belief revision operators is needed) of many possible plausibility updates. For the definitions of such generalizations, the reader is referred to van Benthem (2007) and Baltag and Smets (forthcoming).

6.2 Rankifying

In this text, I have described something like $\Uparrow$ before: In section 2 *lexicographic revision* was introduced as a revision for systems of spheres. The context in which it was mentioned was Spohn's criticism of system of spheres models of belief revision. This criticism applies to van Benthem's system as well: The big difference between a BDDL were a suitable revision (modelling lexicographic revision) is used and the logic with $\Uparrow$ just considered is that the latter lets one talk about higher-order beliefs and revisions. If we restrict its syntax to BDDL expressions (syntactically converting between the revision operators), we see that on this restriction, the logics amount to the same. So if Spohn's arguments are correct, this approach is in trouble, too.

This is not always acknowledged. For example, in Baltag and Smets (2008, p 23), we read:

> Our models are the same as Board's "belief revision structures" [19], i.e. nothing but "Spohn models" as in [48], but with a purely relational description. Spohn models are usually described in terms of a map assigning ordinals to states. But giving such a map is equivalent to introducing a well pre-order $\leq$ on states, and it is easy to see that all the relevant information is captured by this order.

As we have seen, Spohn's arguments lead him to the opposite of what Baltag and Smets say: He explicitly holds that orders (no matter what kind) do *not* capture all the relevant information. Now, whether or not he's right with this, if there are such arguments, one should explain why one rejects them if one doesn't follow their conclusion. Of course, one could defend the quoted paragraph and say that with "relevant information", the relevant information *for the revision operator Baltag and Smets want to capture* is meant. This would of course be true. But this line of argument is not a good option, as they explicitly want to create a logic which is not limiting in this respect, as they hold that the different revision operators can be represented in it.

Assuming that Spohn's arguments are correct, what should a proponent of dynamic logic for belief revision do? My answer is simple: Rankify! Just as by replacing systems of spheres in BDDL with ranking functions, one obtains a logic of ranking theory with only minor adjustments, so can the plausibility orders of dynamic logics for belief revisions be replaced by ranking functions. This is the process I want to call *rankifying*.

On the other hand, a ranking theorist might see dynamic epistemic logics for belief revision and think to herself: "I want to be able to talk about higher-order beliefs and revisions too!". To her, my answer is as well: Rankify! There are a number of promising proposals for belief revision in dynamic logics, and there seems to be no reason why a ranking theorist shouldn't be able to express her account of revision in them – just rankify them.

7 Conclusion

I have shown that just like belief revision theories based on systems of spheres, theories based on ranking functions can be put into the form of a logic. The logic I presented has been shown to be a fruitful tool of formal epistemology by a few results comparing the two approaches to belief revision theory, and an argument for the use of extended codomains of ranking functions. Furthermore, the parallels between logics using systems of spheres and logics using ranking functions show how putting ranking theory in terms of a logic enables one to extend ranking theory to multi-agent, higher-order beliefs and revisions.

Acknowledgements

I thank Hans van Ditmarsch and Hans Rott for helpful discussions and my supervisors Franz Huber and Wolfgang Spohn for insightful comments on drafts of the present thesis.

Research on this thesis was supported by the German Research Foundation through its Emmy Noether Program.

Mathematical Appendix

A Relations

A binary relation R on a set S is a subset of the cartesian product of S with itself: $R \subseteq S \times S$. Therefore, such a relation is a set of tuples on S. As a notational variant, I write xRy for $\langle x, y \rangle \in R$. For a binary relation R on S, the following properties are defined:

R is *reflexive* iff for any $x \in S$, xRx.
R is *transitive* iff for any $x, y, z \in S$, if xRy and yRz, then xRz.
R is *antisymmetric* iff for any $x, y \in S$, if xRy and yRx then $x = y$.
R is *total* iff for any $x, y \in S$, xRy or yRx (or both).

If R is a reflexive and transitive binary relation on S, then for any $S' \subseteq S$:

the *minimum* of S': $min_R S' = \{x \in S' | \forall y \in S' : yRx \rightarrow xRy\}$

And for any reflexive and transitive relation R on S:

R is *wellordered* iff for any $S' \subseteq S$ such that $S' \neq \emptyset$, $min_R S' \neq \emptyset$.

Using these properties, the following diagram defines certain classes of relations. It is to be read such that e.g. a prewellorder is defined as a reflexive, transitive, wellordered, and total order.

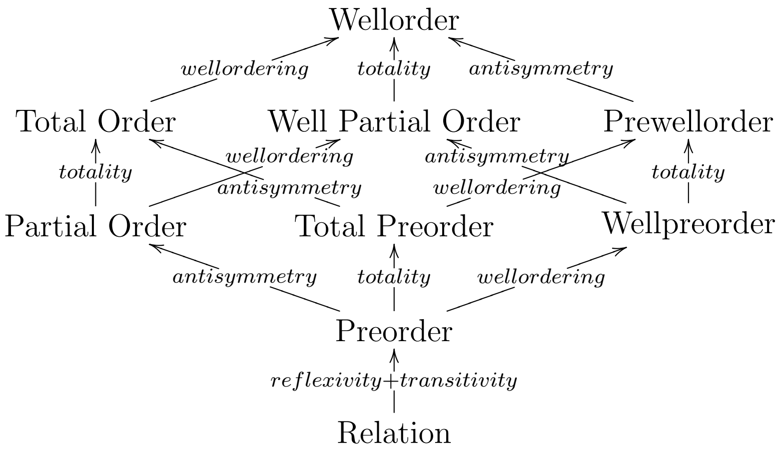

Additionally, a relation R is *functional* iff for any x, there is exactly one y such that xRy.

B Ordinal Numbers

I start with an informal introduction to ordinal numbers, and then show one way of defining them more formally.

B.1 Informal Introduction

Natural numbers serve two purposes: they can be used to represent the size of a finite set, and to represent the position of an element in a finite wellordered set. To represent the position of an element in an *infinite* wellordered set, *transfinite ordinal* numbers may be needed.

The ordinal numbers represent *positions* in a wellordered set. Imagine an infinite wellordered set. Take the minimal element and assign position 0 to it. Removing the minimal element of a wellordered set results in another wellordered set, so you can repeat the process on the set of yet unlabeled elements and assign 1 to the minimal unlabeled element. If you do this for all natural numbers and have not yet exhausted the set, you need *transfinite ordinal numbers*. The first of these is ω, so the minimal unlabeled elements after using up all natural numbers will be assigned position ω. The next will be $\omega + 1$, then $\omega + 2$ and so on. I have not yet introduced ordinal addition, so think of "$\omega + 1$" as just a name for the ordinal succeeding ω. Here's an example of a set transfinite ordinals are needed for: $\mathbf{N} \cup \{a, b, c\}$, ordered as follows: $1 < 2 < 3 < \cdots < a < b < c$.

B.2 Formal Definition

This is mostly taken from Levy (1979, pp 52ff) and Zuckerman (1974, pp 219ff). The alternative definitions of subtraction can be found in Klaua (1969, pp 173ff).

B.2.1 Ordinal Numbers

Let a set S be *transitive* iff $\bigcup S \subseteq S$. A set S is an *ordinal* iff it is transitive and wellordered by $\in$. The class of ordinals is called "*Ord*". The ordinals are ordered by $\leq$ and its corresponding strict (that is, non-reflexive) order $<$, which are defined as follows: for any ordinals x and y, $x < y$ iff $x \in y$ and $x \leq y$ iff $x < y$ or $x = y$. It can be proven that $\leq$ is a wellorder.

Further can be proven that for any ordinal α, $\alpha \cup \{\alpha\}$ is an ordinal, and there is no ordinal β such that $\alpha < \beta < \alpha \cup \{\alpha\}$. $\alpha \cup \{\alpha\}$ is called the *successor* of α.

Repeatedly applying this generation of successors on the empty set produces the following series of ordinals: $\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \ldots$ These are identified with the natural numbers, starting with $0 = \emptyset$. It can be shown that the set of natural numbers $\mathbf{N}$ is itself an ordinal. If $\mathbf{N}$ is used as an ordinal, "ω" is used for it.

An ordinal is a *successor ordinal* iff there is an ordinal β such that α is the successor of β. An ordinal is a *limit ordinal* iff $\alpha \neq 0$ and α is not a successor ordinal. It follows that ω is the smallest limit ordinal.

B.2.2 Ordinal Arithmetic

For any ordinals α, β, define

Addition $$\alpha + \beta = \begin{cases} \alpha & \text{if } \beta = 0 \\ \bigcup\{\alpha + \gamma | \gamma < \beta\} & \text{else} \end{cases}$$

Multiplication $$\alpha * \beta = \begin{cases} 0 & \text{if } \beta = 0 \\ (\alpha * \gamma) + \alpha & \text{if } \beta = \gamma + 1 \\ \bigcup\{\alpha * \gamma | \gamma < \beta\} & \text{else} \end{cases}$$

Exponentiation $\alpha^\beta = \begin{cases} 1 & \text{if } \beta = 0 \\ (\alpha^\gamma) * \alpha & \text{if } \beta = \gamma + 1 \\ \bigcup\{\alpha^\gamma | \gamma < \beta\} & \text{else} \end{cases}$

Note that neither addition nor multiplication is commutative. (An operation $\circ$ is commutative iff for any x, y, $x \circ y = y \circ x$.) It can be proven that for any ordinals α, β such that $\beta \leq \alpha$, there is exactly one ordinal γ such that $\beta + \gamma = \alpha$. Therefore, for any ordinals α, β such that $\beta \leq \alpha$:

Subtraction $\alpha - \beta =$ the $\gamma \in Ord$ such that $\beta + \gamma = \alpha$

Alternative Definitions of Subtraction In Spohn (1988), the subtraction definitions of Klaua (1969) are used. For comparison, they are given here. He defines two notions of subtraction for any ordinals α, β such that $\beta \leq \alpha$. Left-sided subtraction: $-\beta + \alpha =$ the $\gamma \in Ord$ such that $\beta + \gamma = \alpha$ and right-sided subtraction: if there is a $\gamma \in Ord$ such that $\gamma + \beta = \alpha$: $\alpha - \beta = min_{\leq}\{\gamma \in Ord | \gamma + \beta = \alpha\}$.

Zuckerman defines only one notion of subtraction, which is Klaua's left-sided subtraction. It seems natural to use this, as it is defined on all ordinals, which Klaua's right-sided subtraction is not. This is surely also the reason why Spohn uses it. I therefore only use subtraction in the sense of left-sided subtraction, but use the notation used by Zuckerman (which Klaua uses for right-sided subtraction).

B.2.3 Order Types

A function $f : A \to B$ is *bijective* iff for any $x \in B$, there is a $y \in A$ such that $f(y) = x$ and for any $x, y \in A$, if $x \neq y$, then $f(x) \neq f(y)$.

Two relations R and S on sets A and B are *order isomorphic* iff there is a bijective function $f : A \to B$ such that xRy iff $f(x)Sf(y)$ for any $x, y \in A$.

For any wellorder S on a set A, the *order type* of S is the unique ordinal number α such that $\in$ on α is order isomorphic to S. It can be proven that this is well-defined for any wellorder.

References

Charlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

Guillaume Aucher. A combined system for update logic and belief revision. In Michael W. Barley and Nik Kasabov, editors, *Pacific Rim International Workshop on Multi-Agents 2004 (PRIMA 2004)*, pages 1–18. Springer, 2005.

Alexandru Baltag and Sonja Smets. A qualitative theory of dynamic interactive belief revision. In Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, editors, *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, volume 3 of *Texts in Logic and Games*, pages 9–58. Amsterdam University Press, 2008.

Alexandru Baltag and Sonja Smets. The logic of conditional doxastic actions. In Robert van Rooij and Krzysztof R. Apt, editors, *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 9–31. Amsterdam University Press, forthcoming.

Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

Matthias Hild and Wolfgang Spohn. The measurement of ranks and the laws of iterated contraction. *Artificial Intelligence*, 172:1195–1218, 2008.

Franz Huber. The consistency argument for ranking functions. *Studia Logica*, 86:299–329, 2007.

Dieter Klaua. *Allgemeine Mengenlehre*, volume 2. Akademie-Verlag, 1969.

Hannes Leitgeb and Krister Segerberg. Dynamic doxastic logic: Why, how, and where to? *Synthese (Knowledge, Rationality & Action)*, 155:167–190, 2007.

Azriel Levy. *Basic Set Theory*. Springer, 1979.

David Lewis. *Counterfactuals*. Blackwell, 1973.

Sten Lindström and Wlodek Rabinowicz. Extending dynamic doxastic logic: Accommodating iterated beliefs and Ramsey conditionals within DDL. In Lars Lindahl, Paul Needham, and Rysiek Sliwinski, editors, *For Good Measure: Philosophical Essays Dedicated to Jan Odelstad on the Occasion of his Fiftieth Birthday*, volume 46 of *Uppsala Philosophy Studies*, pages 123–153. Uppsala University, 1997.

Abhaya C. Nayak, Randy Goebel, and Mehmet A. Orgun. Iterated belief contraction from first principles. In Manuela M. Veloso, editor, *Proceedings* *of the Twentieth International Joint Conference on Arificial Intelligence (IJCAI-07)*, pages 2568–2573, 2007.

Hans Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In David Makinson, Jacek Malinowski, and Heinrich Wansing, editors, *Towards Mathematical Philosophy. Papers from the Studia Logica conference Trends in Logic IV*, volume 28 of *Trends in Logic*, pages 1–28. Springer, 2008.

Krister Segerberg. The basic dynamic doxastic logic of AGM. In Mary-Anne Williams and Hans Rott, editors, *Frontiers in Belief Revision*, pages 57–84. Kluwer, 2001.

Krister Segerberg. Belief revision from the point of view of dynamic doxastic logic. *Bulletin of the IGPL*, 3:535–553, 1995.

Prakash P. Shenoy. On Spohn's rule for revision of beliefs. *International Journal of Approximate Reasoning*, 5:149–181, 1991.

Wolfgang Spohn. Ranking theory. Typescript, 2008.

Wolfgang Spohn. Ordinal conditional functions. A dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume 2, pages 105–134. Kluwer, 1988.

Wolfgang Spohn. Ranking functions, AGM style. In Sören Halldén, Bengt Hansson, Wlodek Rabinowicz, and Nils-Eric Sahlin, editors, *Internet Festschrift for Peter Gärdenfors*, pages 1–20. Internet Publication, `http://www.lucs.lu.se/spinning`, 1999.

Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17:129–155, 2007.

Hans van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese (Knowledge, Rationality & Action)*, 147:229–275, 2005.

Hans van Ditmarsch. Comments on 'The logic of conditional doxastic actions'. In Robert van Rooij and Krzysztof R. Apt, editors, *New Perspectives on Games and Interaction*, volume 4 of *Texts in Logic and Games*, pages 33–44. Amsterdam University Press, forthcoming.

Martin Zuckerman. *Sets and Transfinite Numbers*. Macmillan, 1974.