



UvA-DARE (Digital Academic Repository)

DDL as an "Internalization" of Dynamic Belief Revision

Baltag, A.; Fiutek, V.; Smets, S.

DOI

[10.1007/978-94-007-7046-1_12](https://doi.org/10.1007/978-94-007-7046-1_12)

Publication date

2014

Document Version

Submitted manuscript

Published in

Krister Segerberg on logic of actions

[Link to publication](#)

Citation for published version (APA):

Baltag, A., Fiutek, V., & Smets, S. (2014). DDL as an "Internalization" of Dynamic Belief Revision. In R. Trypuz (Ed.), *Krister Segerberg on logic of actions* (pp. 253-280). (Outstanding contributions to logic; Vol. 1). Springer. https://doi.org/10.1007/978-94-007-7046-1_12

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Chapter 1

DDL as an “Internalization” of Dynamic Belief Revision

Alexandru Baltag, Virginie Fiutek and Sonja Smets

Abstract In this paper we re-evaluate Segerberg’s “full DDL” (Dynamic Doxastic Logic) from the perspective of Dynamic Epistemic Logic (DEL), in its belief-revision-friendly incarnation. We argue that a correct version of full DDL must give up the Success Postulate for dynamic revision. Next, we present (an appropriately generalized and simplified version of) full DDL, showing that it is a generalization of the so-called Topo-logic of Moss and Parikh. We construct AGM-friendly versions of full DDL, corresponding to various revising/contracting operations considered in the Belief Revision literature. We show that DDL can internalize inside one model the “external” doxastic dynamics of DEL, as well as the evidential dynamics investigated by van Benthem and Pacuit. In our Conclusions section, we compare three styles of modeling doxastic dynamics: DDL, DEL and PDL/ETL (the Propositional Dynamic Logic approach, with its Epistemic Temporal Logic variant).

1.1 Introduction

Following the seminal work of Hintikka [11], the field of epistemic/doxastic logic generated a series of interesting logical systems which have sparked the interest of several groups of researchers: the philosophers interested in using logical formalism to address the questions raised in the traditional study of epistemology, the researchers in AI studying agency, attitudes, non-monotonic reasoning and knowledge representation, and the computer scientists investigating distributed systems.

Alexandru Baltag
University of Amsterdam, ILLC Amsterdam, e-mail: thealexandrubaltag@gmail.com

Virginie Fiutek
University of Amsterdam, ILLC Amsterdam e-mail: fiutek.virginie@gmail.com

Sonja Smets
University of Amsterdam, ILLC Amsterdam e-mail: S.J.L.Smets@uva.nl

The interaction with these other areas of research gave a boost to the further development of epistemic/doxastic logic and raised the interest in the logical study of belief change and knowledge update.

It is at this point of time in the development of modal doxastic logic that we place the important contributions of Krister Segerberg: his idea was to enhance traditional epistemic and doxastic logics with specific dynamic-modal operators for “belief revision”, thus linking modal logic with Belief Revision Theory (BRT). Looking the other way around, Segerberg’s work provided BRT with a new syntax and formal semantics. Note that traditionally, the work on belief revision [1] focuses on the way in which a given theory (or belief base, consisting of sentences in a given object language) gets revised, but it does *not* treat “belief revision” itself as an ingredient in the object language under study. Segerberg’s work opened up a new perspective by taking the very act of belief revision itself and placing it on an equal (formal) footing with the doxastic attitudes such as “knowledge” and “belief”.

Dynamic doxastic logic (DDL) has been introduced and developed by Krister Segerberg in [19, 20, 21, 22, 24, 23, 25, 26]. The system’s main syntactic construct is the use of a dynamic modal operator $[*\varphi]\psi$ whose intended meaning is that “ ψ holds after (the agent performs a) revision with φ ”. As explained in [26], the main added value of treating belief revision in this way (in contrast with the AGM approach [1]) is that we gain all the well-known advantages provided by working in a modal logic setting. Modal logics have turned into a rich area of investigation with applications to several other domains, hence casting Belief Revision Theory into a modal framework holds a great promise for its future development.

Segerberg distinguished between “basic DDL” and “full DDL”: while basic DDL is about the way an agent revises her beliefs about the world, full DDL deals with the way in which an agent revises her beliefs about the world *and* about her own beliefs. Syntactically, this distinction is captured by restricting all the operators of the basic DDL language to Boolean formulas (while full DDL is not subject to this restriction).

In this paper we take a fresh look at full DDL from the new perspective of “soft DEL” (the belief-revision-friendly version of Dynamic Epistemic Logic [3, 4, 5, 6, 2, 7]), as a modern semantic embodiment of the AGM paradigm. DEL shares with DDL the modal logic approach to belief and belief-revision. However, DEL treats dynamic revision as an “external” operation (representing actions as changes of the current model), while in DDL the dynamics is “internal” to the model (i.e., actions are represented as changes of doxastic structure within the same model). One of our goals in this paper is to show that the DDL approach is at least as powerful as the DEL approach: it can internalize all the recent DEL developments.

We start, in Section 2, by borrowing from DEL the distinction between static and dynamic revision, in order to correct an old conceptual error that plagued all attempts to develop a full DDL: the assumption that the AGM Success Postulate is tenable (and desirable) for dynamic revision. We show that (due to Moore-type paradoxes) a correct version of full DDL must give up the unrestricted dynamic version of the Success Postulate (keeping it only for static revision, or for the restriction of dynamic revision to non-doxastic sentences).

Next, in Section 3, we present an appropriately generalized and simplified version of full DDL. In Section 4 we show that DDL can be considered to be a generalization of the so-called Topo-logic of Moss and Parikh [15, 17]. In Section 5, we deal with static revision in DDL, by adopting the conditional belief logic *CDL* from [3].

Further, in Section 6, we develop and axiomatize (in “constructive”, DEL-style) three versions of DDL, that internalize three of the revision operations considered in the Belief Revision literature. In Section 7, we analyze Segerberg’s constructive treatment of expansion and contraction in DDL, and comment on their non-AGM nature. Next, we introduce and axiomatize three AGM-friendly versions of contraction and expansion in DDL. Based on this work, we argue that (if appropriately generalized), the DDL approach is at least as powerful as the DEL approach. In Section 8, we exemplify this point further by showing that Segerberg’s generalized “hypertheory” version of DDL can internalize, not only belief dynamics, but also the *evidential dynamics* of van Benthem and Pacuit [8].

Finally, in Section 9, we compare three ways of doing dynamic belief revision: DDL, DEL, and PDL/ETL (i.e. the Propositional Dynamic Logic style of modeling belief changes, and its Epistemic Temporal Logic variant). Both PDL/ETL and DDL are ways to “internalize” the doxastic dynamics inside one model, but we argue that the DDL style is the most natural, most elegant and most “economical” way to do this internalization.

1.2 Static versus Dynamic Belief Revision

Before developing full DDL, we first need to correct what we think to be a conceptual mistake of its founder, concerning the validity of the so-called Success Axiom in a dynamic setting. To address this, we follow the DEL literature in *distinguishing between “static” and “dynamic” belief revision*. Though it is often explained in syntactic terms (as referring to two different kinds of behaviour under revision with higher-level doxastic sentences), from a semantic point of view this distinction is in fact related to (though distinct from) the traditional dichotomy between *one-step* revision and *iterated* revision.

To model one-step revision, it is enough to specify, for every proposition P , the result of doxastic revision with P , either syntactically (as a set of sentences) or semantically (as a set of states, the ones that are most plausible after revising with P). *Semantically*, this can be uniformly done in three different ways: by giving a *selection function*, in Stalnaker’s style; by giving a *family of spheres* (in Lewis-Grove style), i.e. an “onion” in the sense of Segerberg (or a “hypertheory”, in his generalized semantics); or by giving a *plausibility relation* (or equivalently, an entrenchment relation). As far as modal (dynamic doxastic) logic can tell, these three semantic styles are equivalent, if considered at an appropriate level of generality.

Syntactically, one can capture static revision by specifying, in AGM-style, a set $T * P$ of revised beliefs, for each original set T of beliefs and each proposition P ; or

alternatively, one can encode static revision using conditional belief operators $B^P Q$ (or $B(Q|P)$), whose meaning is that “after revision with P , the agent will come to believe that Q was the case (*before the revision*)”.

The “static” character of this revision is reflected in the fact that, after the revision, Q is still evaluated according to the original state of affairs; in terms of Grove spheres, this is reflected in the fact that the same onion is used for evaluating Q (though not the same sphere): $B^P Q$ holds iff the smallest sphere in the current onion that intersects P is included in Q .

In contrast, dynamic revision involves a *change of onion*, or a change of plausibility relation (or a change of model). Semantically, it requires a binary relation between onions (in DDL-style), or between states with different plausibility (in PDL/ETL style), or between models (in DEL-style). Again, these three styles of doing doxastic dynamics are equivalent, if considered at an appropriate level of generality. Syntactically, dynamic revision can be captured by the use of dynamic modalities $[*P]Q$. More precisely, $[*P]BQ$ captures the fact that Q is believed to hold after revision with P . The “dynamic” character is reflected in the fact that, after the revision, BQ is evaluated using the *new* onion (to which the old onion is related by the dynamic binary relation R^{*P}).

The “static” character of conditional belief operators $B^P Q$ can be made more explicit by expressing them in terms of dynamic operators (“weakest preconditions”) $[*\phi]\psi$ and their *Galois adjoints*, i.e. the “reversed” dynamic operators $\langle *^{-1}\phi \rangle \psi$ (also known as “strongest postconditions”). While dynamic operators $[*\phi]\psi$ are (in the Segerberg’s onion semantics) the universal (Box) modalities for some binary revision relation $R^{*\phi}$ between onions, the reversed dynamic operators $\langle *^{-1}\phi \rangle \psi$ are the existential (Diamond) modalities for the converse relation $(R^{*\phi})^{-1}$ (going *backwards* in time from of the revised doxastic state to the initial, unrevised doxastic state). It is easy to see that we have the following equivalence:

$$B^\phi \psi \Leftrightarrow [*\phi]B\langle *^{-1}\phi \rangle \psi.$$

This equivalence fully captures our above explanation of static revision $B^\phi \psi$, as reflecting the *revised beliefs* (after a revision with ϕ) about a sentence ψ ’s truth value *before* the revision.

Nevertheless, in our logics we chose *not* to reduce static revision to dynamic revision (and its converse); instead, we take static revision as basic, in the shape of primitive conditional belief operators $B^\phi \psi$, interpreted as *belief-revision plans*: “if *in the future* I ever would have to revise with ϕ , I would then come to believe that ψ was true now”. And we follow the DEL tradition by recursively reducing any instance of dynamic revision to the static revision statements (via so-called Reduction laws, or Recursion laws). We choose this option because we think that, from a semantic point of view, static belief revision is a *simpler* concept than the dynamic one. Indeed, recall that to specify static revision one only needs to give *one onion* (together with a specific way to move between its spheres). While dynamic belief revision is given by a specific type of *onion change* (i.e. a specific way of moving between onions): a relation between onions! So in fact, dynamic belief revision does

not involve only a simple revision of *beliefs*, but rather a *revision of (static) belief revision plans*! Indeed, to syntactically describe in full a given type of dynamic belief revision, we do not need only statements of the form $[*P]BQ$ (describing dynamic revision of beliefs), but rather sentences of the form $[*P]B^RQ$ (describing dynamic revision of static belief-revision plans).

Luckily, this distinction does *not* need to be iterated: since (to use van Benthem’s expression) static belief revision B^RQ “pre-encodes” dynamic belief revision $[*R]BQ$, it is enough to know the behaviour $[*P]B^RQ$ of static revision plans under dynamic revision in order to be able to calculate the result of *iterated dynamic revision* $[*P][*R]BQ$. More generally, for each specific type of doxastic dynamic revision $*$, the statement $[*P]Q$ can be recursively reduced to a statement involving only static revision operators B^RQ : these are the well-known Reduction (or Recursion) Laws, from Dynamic Epistemic Logic.

Thus, dynamic revision, by its very semantic modelling, can be straightforwardly *iterated*; while static belief revision is just a *one-step* revision of (simple) beliefs.

But the distinction of static versus dynamic revision is *not the same* as the distinction between one-step and iterated revision! Dynamic revision fully “keeps up” with the doxastic change, while static revision looks back at the old doxastic state from the perspective of the new one.

To see this, note that *dynamic revision with higher-level doxastic sentences behaves differently than static revision*. Take a Moore sentence, of the form $\phi := p \wedge \neg Bp$. An introspective agent will obviously *not* come to believe ϕ after she learns ϕ ; indeed, believing ϕ would amount to a lack of introspection, since it would mean to believe *both* p *and* the fact that one doesn’t believe p . So, after learning ϕ , an introspective agent will clearly come to believe p , but not ϕ itself. This is correctly reflected by dynamic revision: as we will see, for any reasonable dynamic interpretation of the revision operation $*$ as a binary relation on doxastic states (onions), the formula $[*\phi]B\phi$ is *false* for any Moore sentence ϕ : indeed, even if ϕ was true in the old doxastic state, after revision with ϕ the sentence $B\phi$ is evaluated according to the new doxastic state, in which ϕ is false, and (known to be false, hence) disbelieved. In contrast, static revision with any sentence ϕ will always produce belief in that sentence, since after static revision, the sentence is still evaluated according to the original doxastic state: this is reflected by the conditional-belief validity $B^\phi\phi$, which is a version of the AGM “Success” Postulate $\phi \in T * \phi$.

This distinction is an important one, that DDL needs to learn from DEL, in order to deal correctly with higher-level doxastic sentences. Ignoring this distinction leads to what we think to be a conceptual “mistake”, made by Lindstrom and Rabinowicz in their papers [13, 14] on DDL for introspective agents, as well as by Segerberg himself in [24]. Namely, these authors assume (mistakenly, in our view) that a dynamic version of the Success postulate (in the form of the axiom $[*\phi]B\phi$) is desirable, or even tenable, in full DDL (i.e. when ϕ is itself a doxastic sentence). As we argue below (and as was already argued before in the DEL literature), this assumption is wrong,

We should stress that this conceptual problem affects *only* the *first* solution to the Moore “paradox” proposed by Lindstrom and Rabinowicz (in the first part of

their paper [13]). There, they define a semantics for revision, which together with their (standard PDL-like) semantics for dynamic modalities, can be shown to immediately lead to a *semantic failure of the Success Postulate for any (positively) introspective agent*. Indeed, in a Lindstrom-Rabinowicz model, formulas are evaluated at “total states” x , each coming with an ontic state (“world”) $w(x)$ and a doxastic state $d(x)$. In their turn, doxastic states $d(x)$ are Segerberg “onions” (or more generally hypertheories): these are families of “spheres” (i.e., of closed sets of total states). If we put $b(x) := \bigcap d(x)$ for the “smallest sphere” of the onion $d(x)$, then belief is defined as usually in Grove models: $x \models B\phi$ iff $b(x) \subseteq \|\phi\|$. Take now any Lindstrom-Rabinowicz model M in which the following two conditions are satisfied: (a) the agent is *positively introspective with respect to some specific fact p* (at all the states of the model); and (b) there exists *some* total state x in which *the agent doesn’t believe p and she doesn’t believe $\neg p$* . It seems clear that, no matter what additional restrictions one might want to impose on Lindstrom-Rabinowicz models, situations satisfying (a) and (b) should still be allowed¹. So, even if we add further conditions, a model M of the above kind should still be in the intended class of models. As a consequence of (b), the smallest sphere $b(x) := \bigcap d(x)$ (at total state x) contains both p and $\neg p$ worlds. In this situation, the Moore sentence $\phi := p \wedge \neg Bp$ is semantically consistent with the agent’s (semantic) beliefs; indeed, ϕ is true at all the p -worlds belonging to the smallest sphere: $b(x) \cap \|\phi\| \subseteq \|\phi\|$. Hence, this smallest sphere $b(x)$ has a non-empty intersection $b(x) \cap \|\phi\| = b(x) \cap \|\phi\| \neq \emptyset$ with the extension $\|\phi\|$ of ϕ in this model. The Lindstrom-Rabinowicz semantic conditions most precisely, their postulates on semantic contraction and their Levi-style definition of revision) ensure that in this situation a revision with ϕ is the same as an expansion with ϕ (as is also prescribed by the AGM theory): so, the total state y obtained after revision (i.e. such that $xR^{*\psi}y$) is the same as the state obtained by expansion, i.e. we have $xR^{+\phi}y$. But unlike revision (or contraction), the expansion operation is completely determined by the AGM axioms, which are accepted by Lindstrom and Rabinowicz, who in fact explicitly assume that the expanded state y is the unique total state satisfying the conditions $w(y) = w(x)$ (stability of ontic state) and $d(y) = d(x) + \|\phi\| =: d(x) \cup \{X \cap \|\phi\| : X \in d(x)\}$. This means that the smallest sphere of the new “onion” $d(y)$ must be $b(y) = \bigcap d(y) = \bigcap d(x) \cap \|\phi\| = b(x) \cap \|\phi\| = b(x) \cap \|\phi\| \subseteq \|\phi\|$. As a consequence, in the new total state y , the agent believes p : $y \models Bp$. Since Positive Introspection with respect to p holds in this model, we also have $y \models BBp$. If the Success Postulate would also hold, in its dynamic form $x \models [* \phi] B\phi$, then by the standard PDL semantics for dynamic operators (accepted by Lindstrom and Rabinowicz in this part of their paper), we would have $y \models B\phi$. Using the normality of the operator B (which is another immediate consequence of the Lindstrom-Rabinowicz semantic definition of belief) and the fact that $\phi := p \wedge \neg Bp$, it follows that $y \models B\neg Bp$. So we have that $y \models (BBp \wedge B\neg Bp)$, and by normality again, we conclude that $y \models B(p \wedge \neg Bp)$, which by the semantic definition

¹ Even if one doesn’t accept Positive Introspection as a general axiom, one certainly shouldn’t exclude situations in which the agent *is* introspective, at least with respect to *some particular fact p* .

of B , entails that $b(y) \subseteq \|Bp \wedge \neg Bp\| = \emptyset$. But this contradicts the above-mentioned fact that $b(y) = \cap d(y) = \cap d(x) \cap \|\phi\| = b(x) \cap \|\phi\| = b(x) \cap \|p\| \neq \emptyset$.

Observe that this contradiction is obtained only using the Lindstrom-Rabinowicz semantics for belief and revision, the Success Postulate, and the natural and innocuous assumptions (a) and (b) (i.e. that there occasionally may exist some agent who is introspective with respect to some fact p , while the fact p itself is currently neither believed nor disbelieved by the agent). Since the title of one of the papers presenting their setting is *Belief Change for Introspective Agents* [14], it seems to us that Lindstrom and Rabinowicz do not aim to give up even the mere *possibility* of Positive Introspection (with respect to even just one factual statement). So it follows that they must give up the Success Postulate.

However, Lindstrom and Rabinowicz resist this conclusion. They do prove a “Moore paradox” (namely, that the agent’s beliefs are inconsistent), but only *syntactically* (axiomatically), using an additional (unnecessary) assumption (the so-called Preservation Principle). Their conclusion is that this last-mentioned assumption (rather than the Success axiom) is to be blamed for the paradox. So they propose giving up Preservation, without apparently noticing that the above clear-cut semantic argument shows already (without any use of Preservation, but using only the local assumptions (a) and (b)) that the Success postulate is conceptually incompatible with their dynamic semantics.

It is true that, in the *second* part of their paper [13], Lindstrom and Rabinowicz propose a second solution to the Moore paradox, their so-called bidimensional semantics, which is in fact very close to our (i.e. the DEL) solution. Indeed, their rendering in English of their proposal is essentially the same as our solution: they point out that the Success Postulate makes sense for doxastic sentences ϕ *only* if it is interpreted in terms of *the revised beliefs about ϕ ’s truth value before the revision*. However, they formally package this solution in a different way, in order to maintain the appearance (at a purely syntactic level!) that the Success Postulate is maintained. Namely, they do this by adopting a bidimensional semantics in terms of pairs of states (x, y) , in order to refer to both doxastic states (before and after the revision), and they *radically change the PDL semantics of dynamic operators to a non-standard one*: roughly speaking, their new semantics amounts to evaluating any doxastic expression $B\psi$ that comes in the scope of a dynamic operator $[*\phi]$ as capturing the revised beliefs (after revision with ϕ) about ψ ’s truth value before the revision.

We fully agree with the conceptual analysis underlying the second solution of Lindstrom and Rabinowicz, but we disagree with their non-standard, and completely ad-hoc, modification of the semantics of dynamic operators. We think dynamic modalities should be left to express what they always did: a one-way move in time, from the state before the (revision) action to the state after the action. Instead of twisting the meaning of dynamic operators, we think one should simply recognize the plain, inescapable truth: the Success Postulate does not (and should not) hold for dynamic revision with doxastic sentences.

In most of his papers on DDL, Segerberg himself is cautious not to fall into the above mentioned conceptual mistake, by almost always limiting himself to “basic

DDL”: no revision with doxastic sentences. However, in [24] he proposes an axiomatic system for full DDL. Unfortunately, this converts a conceptual mistake into a logical error: the proposed system is *not sound* with respect to the proposed semantics. The reason is that the proposed Success Axiom $[\ast\phi]B\phi$ is not a validity in this semantics: essentially, the above model provides a counterexample to this. The semantic setting in [24] differs slightly from the version of DDL presented in our paper (since we follow [21, 26]), in that it is actually closer to the Lindstrom-Rabinowicz setting: formulas are evaluated at *states* (called “points”), not at pairs of a state and an onion, and so the dynamics is given via binary relations between states (similarly to the standard *PDL* semantics), rather than via relations between onions. The resulting relational frame is called a revision space. However, in this setting (from [24]), each state is assigned an onion, via an “onion determiner”, which paired with a revision space gives an “onion frame”. Completeness (for an axiomatic system that includes the dynamic version of the Success Axiom) is claimed with respect to the class of “AGM onion frames” (i.e. onion frames satisfying some additional AGM-like semantic conditions). Introspection is *not* assumed by Segerberg in this setting, neither as a semantic condition nor as an axiomatic one. But it is easy to see that (Positive) Introspection is consistent with this setting: there exist AGM onion frames that are positively introspective. More precisely, the above counterexample (an introspective onion model in which neither p nor $\neg p$ are believed) can be easily repackaged as an AGM onion model in the sense of [24]. The dynamic version of the Success Axiom, when instantiated to the Moore sentence $p \wedge \neg Bp$, fails in this model. So this axiom is simply not sound.²

The lesson is that in DDL (as in DEL) we can really make sense of dynamic revision with doxastic sentences by an introspective agent *only* if we drop (the unrestricted, dynamic version of) the Success Postulate. A weakened version of this postulate can be retained either by (a) restricting it to (dynamic revision with) simple, Boolean, *non-doxastic* sentences (as in the AGM literature, as well as in many of Segerberg’s papers), or by (b) interpreting it in terms of *static* revision (i.e. as a conditional-belief statement $B^\phi\phi$).

1.3 General DDL Semantics

We present here a generalized (and simplified) version of the “General Model Theory” for DDL introduced by Segerberg in Section 3 of [21]. The semantics is based on Segerberg’s “hypertheories” (i.e. families of sets of states, called “fallbacks”), which are generalizations of Segerberg’s “onions” (which are families of *nested* sets of states, called “spheres”, in accordance to the Lewis-Grove tradition). As a formal language to describe these models, we use the slightly extended syntax

² While soundness of the given axiomatic system is not explicitly claimed in [24], its completeness is claimed. But from a conceptual point of view, a completeness result (with respect to a class of frames) is of course of no use if the axioms are not sound (with respect to that same class of frames).

for DDL introduced in [26], having (in addition to belief operators B and dynamic modalities) operators K for what Leitgeb and Segerberg call “nonrevisable belief” or “knowledge”. We call this “irrevocable knowledge”, to distinguish it from other, “softer” notions of knowledge considered in the philosophical literature (e.g. defeasible knowledge). To ensure that the K operator is factive (as it is expected from “knowledge”), we make a slight change to the definition of validity, inspired from the Moss-Parikh semantics of Topo-logic: validity is obtained by quantifying only over pairs (s, H) of ontic states and hypertheories such that $s \in H$. We further simplify Segerberg’s setting from [21], by dropping all the topological assumptions (Stone spaces, compactness assumptions), as well as all the closure assumptions on hypertheories (e.g. Lewis’ famous Limit Assumption, or the assumption from [26] of closure under nonempty intersections). The price for this generality is that the definition of belief is more complicated: we adopt the definition of B introduced by van Benthem and Pacuit [8]. But we show that, whenever hypertheories do satisfy closure under intersection, this definition boils down to Segerberg’s notion of belief (which is the same as Grove’s definition: belief equals truth in all the states of the smallest sphere). Moreover, we show that in the special case of onions, this definition amounts to a natural generalization of Grove’s definition (belief equals truth in all the states of all the spheres that are “small enough”), that was already proposed in the Belief Revision literature (and which validates the same modal formulas as Grove’s standard definition). Finally, it is easy to see that, in case of onions satisfying Lewis’ Limit assumption, this definition boils down again to the standard (Grove-Segerberg) notion of belief.

Let U be a set of states (a *universe*). A *hypertheory* in U is a nonempty family $H \subseteq \mathcal{P}(U)$ of nonempty subsets of U , called *fallbacks*. An *onion* (or “sphere system”) in U is a hypertheory $O \subseteq \mathcal{P}(U)$, that is “nested”, i.e. linearly ordered by set-inclusion: $X, Y \in O$ implies that either $X \subseteq Y$ or $Y \subseteq X$. The elements of an onion (its fallbacks) are sometimes called “spheres”.

We think of each $s \in U$ as an “ontic state”: a possible description of all the ontic (i.e. non-doxastic) facts of the world. We think of a hypertheory H as representing the agent’s “doxastic state”. In particular, as we will see in the next section, an onion O will represent a doxastic state that satisfies the AGM postulates (when these postulates are appropriately stated, as axioms about *static* revision).

An onion O is *standard* (or “well-founded”) if there is no infinite descending chain of spheres in O ; i.e. there is no infinite sequence $X_1 \supset X_2 \supset X_3 \supset \dots$, with all $X_i \in O$.

Given a hypertheory $H \subseteq \mathcal{P}(U)$, a family $F \subseteq H$ of fallbacks has the **finite intersection property** (f.i.p.) if every *finite* subfamily $F' \subseteq F$ has a non-empty intersection $\bigcap F' \neq \emptyset$. We say that a family $F \subseteq H$ of fallbacks has the **maximal f.i.p.** if F has the f.i.p. but no proper extension $F \subset G \subseteq H$ does. Observe that, if O is an onion, then O has itself the maximal f.i.p.; and moreover O is the only family $F \subseteq O$ having the maximal f.i.p.

An *A-doxology* is a structure (U, D, R) , where U is a universe, D is a set of hypertheories in U and $R = \{R^\alpha\}_\alpha$ is a set of binary relations $R^\alpha \subseteq D \times D$ on D , labeled with names $\alpha \in A$ coming from a given set A of *action terms*. The elements

$R^\alpha \in R$ of R are called *doxastic actions*, and R itself a *repertoire*. Observe that each R^α is a binary relation between hypertheories (or onions), *not* between states. Intuitively, each R^α describes a specific type of change which may affect the agent's epistemic/doxastic state (but which *does not change the "ontic state"*).

Assume now given any object language \mathcal{L} containing *propositional letters* coming from a set Φ , Boolean connectives, a *belief operator* B , an *irrevocable knowledge operator* K , a set A of *action terms*, as well as and the dynamic modalities $[\alpha]$ ("after action α ") of Propositional Dynamic Logic (one for each action term $\alpha \in A$). Any such language \mathcal{L} is called a *DDL-language*. The *minimal language of ("full") DDL* has *only* the above operators. (But later we will add conditional belief operators, to describe static revision.)

A *DDL model* $M = (U, D, R, V)$ for any *DDL language* \mathcal{L} (with propositional letters in Φ and action terms in A) consists of an A -doxology (U, D, R) together with a valuation V , mapping every propositional letter $p \in \Phi$ to a set $V(p) \subseteq U$ of states. An *onion model* is a *DDL model* (U, D, R, V) in which D consists only of onions. A *static (DDL) model* is a *DDL model* with $R = \emptyset$.

A *semantics* for \mathcal{L} is a map that, for each *DDL model* $M = (U, D, R, V)$ and each hypertheory $H \in D$, assigns to each formula $\phi \in \mathcal{L}$ some set of states $\|\phi\|_{M,H} \subseteq \bigcup H$, and assigns to each action term $\alpha \in A$ some doxastic action $\|\alpha\|_{M,H} \in R$, in such a way that a number of conditions (to be given below) are satisfied. Our restriction to $\bigcup H$ is motivated by the intuition that the states $s \notin \bigcup H$ represent "impossible states": ontic states that are excluded by the doxastic state H . In other words, $\bigcup H$ encompasses the agent's her "hard information" about the world. As a consequence, the operator K (given by quantifying over $\bigcup H$) is *factive* (unlike in the usual setting of DDL): we can think of K as representing the agent's "knowledge", in the absolute sense of infallible, absolutely certain, and absolutely unrevisable knowledge. We use the notation

$$s, H \models_M \phi$$

whenever we have $s \in \|\phi\|_{M,H}$, and we delete the subscript(s) whenever it is possible to do this without ambiguity, writing e.g. $\|\phi\|_H$ and $s, H \models \phi$ when M is fixed, or even $\|\phi\|$ when both M and H are fixed. (Note that $s, H \models \phi$ can only hold for $s \in H$.) A semantics for \mathcal{L} is required to satisfy the following constraints:

$$\begin{aligned} s, H \models p & \quad \text{iff } s \in V(p) \\ s, H \models \neg\phi & \quad \text{iff } s, H \not\models \phi \\ s, H \models \phi \wedge \psi & \quad \text{iff } (s, H \models \phi) \wedge (s, H \models \psi) \\ s, H \models B\phi & \quad \text{iff } \forall \text{maximal f.i.p. } F \subseteq H \exists F' \text{ finite } \subseteq F \forall t \in \bigcap F' (t, H \models \phi) \\ s, H \models K\phi & \quad \text{iff } \forall t \in \bigcup H (t, H \models \phi) \\ s, H \models [\alpha]\phi & \quad \text{iff } \forall H' \in D ((H, H') \in \|\alpha\|_H \wedge s \in \bigcup H' \implies s, H' \models \phi) \end{aligned}$$

For a class \mathcal{C} of (*DDL*) models, we write $\mathcal{C} \models \phi$ and we say that ϕ is *valid* on \mathcal{C} , if $\|\phi\|_{M,H} = U$ for every model $M = (U, D, R, V) \in \mathcal{C}$ and every $H \in D$; equivalently,

iff $s, H \models_M \phi$ holds for all models $M = (U, D, R, V) \in \mathcal{C}$, all hypertheories $H \in D$ and all states $s \in \bigcup H$.

An onion model (U, D, R, V) is *standard* if all the onions $O \in D$ are standard. A weakening of the standardness condition, which has the disadvantage of being language-dependent is the so-called Lewis Limit Assumption: an onion model (U, D, R, V) , together with a semantics $\|\bullet\|$ is said to *satisfy the Limit Assumption* if, for every formula $\phi \in \mathcal{L}$ and every onion $O \in D$, we have that: $\|\phi\| \cap \bigcup O \neq \emptyset$ implies $\bigcap \{X \in O : \|\phi\| \cap X \neq \emptyset\} \in O$.

It is easy to see that *standard onion models always satisfy the Limit Assumption* (for every language \mathcal{L}); but the converse is false. In fact, standard onion models satisfy a stronger condition, that we call the Strong Limit Assumption: for every set $P \subseteq U$ of states and every onion $O \in D$, $P \cap \bigcup O \neq \emptyset$ implies $\bigcap \{X \in O : P \cap X \neq \emptyset\} \in O$. This means that, in a standard model, every onion intersecting a given set P contains a unique smallest sphere intersecting P .

Any fallback H in a DDL model induces a corresponding relation of plausibility between states. We say that *state s is at least as plausible as state t according to H* , and we write $s \leq_H t$, if s belongs to all the fallbacks in H that contain t :

$$s \leq_H t \text{ iff } \forall X \in H (t \in X \Rightarrow s \in X).$$

Obviously, the plausibility relation \leq_H is a *preorder* (reflexive and transitive relation) on the set $\bigcup H$. Moreover, if O is an onion, then \leq_O is a *total* (i.e. connected) preorder on $\bigcup O$: for all $s, t \in \bigcup O$, we have either $s \leq_O t$ or $t \leq_O s$ (or both).

Our definition of irrevocable knowledge K is essentially the same as in [26], except that our modified definition of validity entails the *factivity* of K , making it to behave indeed like a notion of “knowledge” (in contrast to [26]). Our definition of belief B is a generalization of the Grove-Seegerberg definition, due to van Benthem and Pacuit [8]. But it can be simplified in onion models (where it boils down to a widely used generalization of Grove’s), and it can be simplified further when we have either the Limit Condition or closure under intersection (where it boils down to the Grove-Seegerberg definition):

Proposition In DDL models in which the set D of hypertheories is closed under non-empty intersections, ϕ is believed iff it is true in all the “most plausible states” (i.e. the states of the smallest fallback):

$$s, H \models B\phi \text{ iff } \forall t \in \bigcap H (t, O \models \phi).$$

In *onion models*, ϕ is believed iff ϕ is true in all the states that are “plausible enough” (i.e. throughout all the spheres that are “small enough”):

$$s, O \models B\phi \text{ iff } \exists X \in O \forall t \in X (t, O \models \phi).$$

Moreover, in *onion models satisfying the Limit Condition*, this boils down to the usual Grove definition:

$$s, O \models B\phi \text{ iff } \forall t \in \bigcap O (t, O \models \phi).$$

(And, as a consequence, this equivalence holds in standard onion models.)

1.4 DDL as a generalization of Topo-logic

The language of Topo-logic, proposed by Moss and Parikh [15, 17], is a modal logic with two modalities: K (for “knowledge”), and \Box (for “effort”). The box modality modality stands for *stability under information increase*: the sentence $\Box\phi$ means that ϕ (is true and) *stays true no matter what the agent increases her information*. Using its De Morgan dual \Diamond , one can define “learnability” as $\Diamond K\phi$ (i.e. ϕ *might come to be known* after learning some further information). Topo-logic frames (U, \mathcal{T}, V) consists of a universe (set of points, or “states”) U , a family $\mathcal{T} \subseteq \mathcal{P}(U)$ of sets of states (called “opens”)³ and a valuation V for the atomic sentences of the above language. While the points $s \in U$ represent possible ontic states, the opens $V \in \mathcal{T}$ represent possible *information states*: when the agent’s information state is V , this means that the only thing that she knows about the state of the world is that it belongs to V . Sentences are evaluated at pairs (s, V) of an ontic state $s \in U$ and information state $V \in \mathcal{T}$, with the restriction that $s \in V$ (so that “knowledge” is factual: indeed, these are information states, rather than doxastic states!). The semantics is given by putting

$$s, V \models K\phi \quad \text{iff} \quad t, V \models \phi \text{ for every } t \in V,$$

$$s, V \models \Box\phi \quad \text{iff} \quad t, V' \models \phi \text{ for every } V' \in \mathcal{T} \text{ such that } V' \subseteq V.$$

If we think of $V, V' \in \mathcal{T}$ as possible information states, then $V' \subseteq V$ means that V' is a *refinement* of V : it contains at least as much information (about the real state $s \in V'$) as V does. So we can think of the move from V to $V' \subseteq V$ as an increase of information: a form of (correct, accurate, infallible) “learning”.

It is easy to see that Topo-logic is a special case of Generalized DDL: we can reinterpret a *topo-logic model* as a special kind of onion model $M = (U, D, R)$ in which all the onions are singletons ($D = \{\{V\} : V \in \mathcal{T}\}$), each consisting of only one fallback $V \in \mathcal{T}$, and in which the repertoire is a singleton $R = \{R_\Box\}$, where the relation R_\Box is given by:

$$\{V\}R_\Box\{V'\} \text{ iff } V \supseteq V' \quad (\text{for all } V, V' \in \mathcal{T}).$$

We can similarly reinterpret the *language* of topo-logic as simply the minimal DDL language for the above kind of (topo-logic) DDL models. *The distinction between the belief and knowledge operators B and K vanishes* in this case (so that we can follow Moss and Parikh and denote them both by the same letter K), and the (only) dynamic modality is denoted by \Box .

³ Although the family \mathcal{T} of all opens is not in general required to be a topology in the mathematical sense, Moss and Parikh do consider and axiomatize various possible closure conditions on \mathcal{T} , including the ones defining a topology.)

1.5 Complete axiomatization of static revision: the logic CDL

To capture static revision, we follow the DEL tradition, by borrowing from conditional logic a *conditional belief operator* $B^\theta \psi$. Our semantic clauses can be naturally extended to this enlarged language.

But first, following Segerberg [21], we introduce the notation

$$H \cap P := \{X \in H : X \cap P \neq \emptyset\}$$

for all hypertheories $H \in D$ and sets $P \subseteq U$ of states, and moreover we generalize to any families $F \subseteq H$ of fallbacks (of a hypertheory H):

$$F \cap P := \{X \in F : X \cap P \neq \emptyset\}.$$

The *relativization of a family $F \subseteq H$ of fallbacks (of a hypertheory H) to a set $P \subseteq U$ of states* is the family

$$F^P := \{X \cap P : X \in F \cap P\} = \{P \cap X : X \in F, P \cap X \neq \emptyset\}.$$

Of course, this operation can be applied in particular to an hypertheory H or union O , producing a *relativized hypertheory* H^P or *relativized union* O^P . A family $F \subseteq H$ of fallbacks has the **finite intersection property relative to P** (P -f.i.p.) if every *finite* subfamily (of its relativization to P) $F' \subseteq F^P$ has non-empty intersection $\bigcap F' \neq \emptyset$. We say that a family $F \subseteq H$ of fallbacks has the **maximal P -f.i.p.** if F has the P -f.i.p. but no proper extension $F \subset G \subseteq H$ has the P -f.i.p. Observe that, if O is an *union* such that $P \cap O \neq \emptyset$, then O has itself the maximal P -f.i.p.; and moreover O is the only family $F \subseteq O$ having the maximal P -f.i.p.

When $P = \|\phi\|_H$ for some formula ϕ , we write “maximal ϕ -f.i.p.” for “maximal $\|\phi\|_H$ -f.i.p.” and so on. Now we define conditional belief by putting:

$$s, H \models B^\theta \phi \text{ iff } \forall \text{ maximal } \theta\text{-f.i.p. } F \subseteq H \exists F' \text{ finite } \subseteq F^{\|\theta\|_H} \forall t \in \bigcap F' (t, H \models \phi)$$

Proposition In *union models*, ϕ is believed conditional on θ iff ϕ is true in all the most plausible states satisfying θ :

$$s, O \models B^\theta \phi \text{ iff } \exists X \in O^{\|\theta\|_O} \forall t \in X (t, O \models \phi).$$

Moreover, in *union models satisfying the Limit Condition*, this boils down to the usual Grove semantics for static revision:

$$s, O \models B^\theta \phi \text{ iff } \forall t \in \bigcap O^{\|\theta\|_O} (t, O \models \phi).$$

The language of *conditional doxastic logic* CDL is the smallest set of formulas containing the atomic sentences $p \in \Phi$, the tautological formula \top and is closed under conditional belief operators $B^\theta \phi$. It can be considered as a variant of the DDL language, in which we take $R = \emptyset$ (so there are no dynamic modalities), while B and K are defined as abbreviations: indeed, by putting

$$B\phi := B^\top \phi,$$

$$K\phi := B^{-\phi} \perp$$

(where $\perp := \neg\top$), we can easily see that these abbreviations are semantically equivalent to the belief and knowledge operators, as defined in the previous section.

Theorem The following proof system *CDL* for conditional doxastic logic is sound and complete w.r.t. the class of *all onion models*, the class of *standard onion models*, and the class of *finite onion models*:

Necessitation Rule:	From $\vdash \phi$ infer $\vdash B^\psi \phi$
Normality:	$\vdash B^\theta(\phi \rightarrow \psi) \rightarrow (B^\theta \phi \rightarrow B^\theta \psi)$
Truthfulness of Knowledge:	$\vdash K\phi \rightarrow \phi$
Persistence of Knowledge:	$\vdash K\phi \rightarrow B^\psi \phi$
Full Introspection:	$\vdash B^\psi \phi \rightarrow KB^\psi \phi$ $\vdash \neg B^\psi \phi \rightarrow K\neg B^\psi \phi$
Hypotheses are (hypothetically) accepted:	$\vdash B^\phi \phi$
Superexpansion:	$\vdash B^{\phi \wedge \psi} \theta \rightarrow B^\phi(\psi \rightarrow \theta)$
Subexpansion (=Rational Monotonicity):	$\vdash \neg B^\phi \neg \psi \wedge B^{\phi \wedge \psi} \theta \leftrightarrow B^\phi(\psi \rightarrow \theta)$

(where in all the above axioms, K is just the abbreviation $K\phi := B^{-\phi} \perp$).

As a consequence, it is easy to see that *onion models satisfy all the AGM postulates for “static” belief revision*, except for the Vacuity Postulate ($T * \phi = \perp$ iff $\vdash \neg\phi$), which is valid only modulo a natural epistemic restriction: $T * \phi = \perp$ iff $T \vdash K\neg\phi$. This restriction is unavoidable in the presence of any “unrevisable belief” operator K : it seems to us to be the natural *epistemic* version of the Vacuity principle. Hence, the resulting theory was called “epistemic AGM” in [3].

Corollary If we take the initial AGM theory T to be the set $T = \{\psi : s, O \models_M B\psi\}$ of all beliefs held in (a given ontic state s and a given onion O of) an onion model M , and interpret the statically-revised theory $T * \phi$ as the set $T * \phi = \{\psi : s, O \models B^\phi \psi\}$ of all conditional beliefs held (conditional on ϕ) in (the same state s and same onion O of the same model) M , then all the postulates of the “epistemic AGM” theory are satisfied.

In contrast, static revision in general DDL models does *not* satisfy the (epistemic) AGM postulates (since the Subexpansion principle fails in general DDL models). In conclusion, general DDL does *not* support an AGM-type theory of belief revision; but *onion models are the natural AGM-friendly version of DDL*.

1.6 Dynamic revision in DDL: internalizing doxastic upgrades

It is sometimes said that the main difference between Dynamic Epistemic Logic and the traditional Epistemic Temporal Logic approach to information change is that *DEL* is *constructive*, while *ETL* is purely *descriptive*: in *DEL*, one actually defines in a constructive way the new doxastic/epistemic relation after a given doxastic action. But this constructive approach can be internalized in *DDL* models, and in fact

[21] anticipated this! As we will see in the next section, in that paper Segerberg used a *constructive* DDL approach to expansion and contraction.

In this section we will use such a constructive DDL approach to belief *revision*. We give constructive definitions of binary relations between onions, that internalize three different revision operations considered in the literature. We adopt from DEL the method of using Reduction/Recursion laws to give complete axiomatizations of the dynamic logics of these three kinds of revision. Indeed, our laws are identical to the ones considered in the DEL literature: in effect, this section is a concrete example of how the DEL-style of modeling and axiomatizing belief revision can be “internalized” in DDL.

One can think of many ways to change the beliefs of an agent according to the information she receives. She can receive *hard information* (unrevisable and irrevocable, since received from an infallible source), or she can receive *soft information* (fallible and potentially revisable).

Receiving “hard” information φ corresponds to what in the DEL literature [7, 2, 9] is called an *update*⁴ $!\varphi$, and in Belief Revision literature is known as a “radical revision” (or irrevocable revision), with φ . This operation changes the model by eliminating all the $\neg\varphi$ -worlds. The result of this elimination is a submodel only consisting of φ -worlds.

A second, softer kind of revision is given by the DEL operation of *lexicographic upgrade* $\uparrow\varphi$ [6, 5], known in Belief Revision literature as “moderate revision” (or lexicographic revision). It changes the model by making all φ -worlds become more plausible than all $\neg\varphi$ -worlds:

Finally, the DEL operation of *conservative upgrade* $\uparrow\varphi$ [6, 5] is known as “conservative revision” (or natural revision) in the Belief Revision literature. It changes the model by making the most plausible φ -worlds become the most plausible overall (while leaving everything else unchanged).

Dynamic Epistemic Logic *DEL* (in its single-agent version) for the above-mentioned three types of upgrades can now be obtained as a *special case* of Generalized *DDL*. For this, we reuse the “relativized onion” notation

$$\mathcal{O}^P := \{P \cap X : X \in \mathcal{O}, P \cap X \neq \emptyset\}$$

introduced in Section 5, to define binary relations $R^{!P}$ (for update), $R^{\uparrow P}$ (for lexicographic upgrade) and $R^{\uparrow P}$ (for conservative upgrade) between onions $\mathcal{O} \in \mathcal{D}$ (of some onion model (U, D, R)) and sets of sets of states $\mathcal{O}' \subseteq \mathcal{P}(U)$, as follows:

$$(\mathcal{O}, \mathcal{O}') \in R^{!P} \text{ iff } \mathcal{O}' = \mathcal{O}^P \neq \emptyset$$

$$(\mathcal{O}, \mathcal{O}') \in R^{\uparrow P} \text{ iff } \mathcal{O}' = \mathcal{O}^P \cup \{X \cup \bigcup \mathcal{O}^P : X \in \mathcal{O}\}$$

$$(\mathcal{O}, \mathcal{O}') \in R^{\uparrow P} \text{ iff } \mathcal{O}' = \{\bigcap \mathcal{O}^P : \bigcap \mathcal{O}^P \neq \emptyset\} \cup \{X \cup \bigcap \mathcal{O}^P : X \in \mathcal{O}\}$$

⁴ Unfortunately, this terminology diverges from the one in Belief Revision literature, where “update” refers to a completely different type of operation, namely to the Katsuno-Mendelson revision.

To visualize these doxastic actions, see Figures 4,5 and 6 in the Appendix.

Next, we define a *DEL onion model* to be a standard onion model $M = (U, D, R)$ such that

$$R = \{R^{!P} : P \subseteq U\} \cup \{R^{\uparrow P} : P \subseteq U\} \cup \{R^{\uparrow\uparrow P} : P \subseteq U\}$$

and such that D is closed under all the relations in R .

The *language* of (this version of) *DEL* is obtained by adding to *CDL* dynamic modalities for all the above types of upgrades. The *semantics* is obtained by defining the interpretation maps $\|\varphi\|$ and $\|\alpha\|$ by double recursion: the static propositional clauses are as in *CDL*, the semantics of dynamic modalities is as in the generalized *DDL*, while the clauses for $\|\alpha\|$ are given by

$$\|\!|\varphi\|\!| = R^{\!|\varphi\|}$$

$$\|\uparrow\varphi\| = R^{\uparrow\|\varphi\|}$$

$$\|\uparrow\uparrow\varphi\| = R^{\uparrow\uparrow\|\varphi\|}$$

Theorem A sound and complete proof system for *DEL* onion models can be obtained by adding to the above proof system of *CDL* the van Benthem Reduction/Recursion laws [6]. We give here only the reduction laws for conditional belief:

$$[!\varphi]B^\psi R \iff \varphi \Rightarrow B^{\varphi \wedge [!\varphi]\psi} [!\varphi]\theta,$$

$$[\uparrow\varphi]B^\psi \theta \iff B^{\varphi \wedge [\uparrow\varphi]\psi} [\uparrow\varphi]\theta \wedge (K^\varphi [\uparrow\varphi]\neg\psi \Rightarrow B^{[\uparrow\varphi]\psi} [\uparrow\varphi]\theta),$$

$$[\uparrow\uparrow\varphi]B^\psi \theta \iff B^\varphi([\uparrow\varphi]\psi \Rightarrow [\uparrow\varphi]\theta) \wedge (B^\varphi[\uparrow\varphi]\neg\psi \Rightarrow B^{[\uparrow\varphi]\psi}[\uparrow\varphi]\theta),$$

where we used the abbreviation $K^\varphi\psi := K(\varphi \Rightarrow \psi)$.

Strongest Postcondition Modalities The standard dynamic modalities $[\alpha]\varphi$ are known in Computer Science as *weakest preconditions*: indeed, they capture the weakest condition that can be imposed on an input information state (s, H) to ensure that, after performing action α in that state, φ will become true in the output-state. The dual modalities (in the sense of Galois duality, rather than De Morgan duality) are the *strongest postcondition* modalities $\langle\alpha^{-1}\rangle\varphi$, capturing the weakest condition that is ensured to hold in an output-state after performing action α on an input state satisfying φ .

While standard *DEL* cannot represent strongest postconditions⁵, *DDL* models contain enough information to define them, as *existential (Diamond) modalities for the converse relations* R_α^{-1} : equivalently, just put

$$s, H \models \langle\alpha^{-1}\rangle\varphi \text{ iff } \exists H' ((H', H) \in \|\alpha\|_H \wedge s, H' \models \varphi)$$

⁵ But extensions of *DEL* which can define strongest postconditions have been proposed by G.Aucher and H. van Ditmarsch.

It is obvious that these operators are indeed the Galois duals of the standard dynamic modalities, and that the same holds for their corresponding de Morgan duals: i.e. we have the validities

$$\begin{aligned}\varphi &\Rightarrow [\alpha]\langle\alpha^{-1}\rangle\varphi, \\ \varphi &\Rightarrow [\alpha^{-1}]\langle\alpha\rangle\varphi.\end{aligned}$$

Finally, using the strongest postcondition modality for lexicographic upgrade, we can check the semantic equivalence:

$$B^\varphi\psi \iff [\uparrow\varphi]B(\langle\uparrow\varphi\rangle^{-1})\psi.$$

This equivalence confirms our interpretation of conditional beliefs $B^\varphi\psi$ as embodiments of “static revision”: the agent’s *revised beliefs* (after revision with φ) about ψ ’s truth value *before* the revision.

1.7 Expansion and contraction in full DDL

In [21], Krister Segerberg used a constructive approach (similar to the one we used above for revision) for modeling expansion and contraction in DDL. Assuming some additional conditions on the hypertheories (namely that they are closed under non-empty intersections and satisfy the Strong Limit Assumption⁶, Segerberg puts, for hypertheories H and sets $P \subseteq U, X \in H$:

$$\begin{aligned}H/P &:= H \cup \{X \cap P : X \in H, X \cap P \neq \emptyset\}, \\ H|P &:= \{Y \in H : X \subseteq Y\},\end{aligned}$$

and requires the doxology D to be closed under these operations. Using these notations, Segerberg says that a fallback $Z \in H$ is a *contraction with $P \subseteq U$ in H* iff Z is a minimal fallback (with respect to inclusion) in the family $H \cap (U - P)$ (where recall that $H \cap (U - P) := \{X \in H : X \cap (U - P) \neq \emptyset\}$). Note that such a contraction with P in H might not exist⁷, and even if it exists it might not be unique! Segerberg then explicitly defines an *expansion action $+P$* and a *contraction action $-P$* (for any given set $P \subseteq U$ of states), given by the following relation on hypertheories in D :

$$\begin{aligned}(H, H') \in R^{+P} &\text{ iff } H' = H/P, \\ (H, H') \in R^{-P} &\text{ iff } H' = H|Z \text{ for some contraction } Z \text{ with } P \text{ in } H.\end{aligned}$$

⁶ Segerberg calls *LR hypertheories* (from Lindstrom and Rabinowicz) the hypertheories that satisfy these two conditions.

⁷ Though the additional closure assumptions made by Segerberg in [21] do ensure the existence of contractions.

Now, these two operations, as defined by Segerberg, do not fit the *AGM* framework. This was a conscious decision by Segerberg, since his aim in [21] was to give a semantics to the Lindstrom-Rabinowicz theory of contraction (in which contraction is *not* unique), rather than the *AGM* theory. In order to try to accommodate *AGM*, first we have of course to restrict the above definitios to *onion models*: as we saw, these are the *AGM*-friendly models for *DDL*. On onion models, contractions with P (as defined above) might still not exist; but, if they do, then they are unique (as required by *AGM*). To ensure existence, we have to further restrict to *onion models satisfying the Limit condition*; or (for simplicity) to the even more restricted case of *standard onion models*. As we'll see, this restriction does ensure that Segerberg's contraction satisfies the *AGM* principles. But, even in this case, we still have problems with Segerberg's definition of *expansion*: this operation does not preserve the "nestedness" property, so it does not map (standard) onions into onions! Moreover, there is no reasonable additional condition that would ensure that the expansion (in the sense of Segerberg) of an onion O with a set P is an onion whenever $P \cap \bigcup O \neq \emptyset$. Since "onionhood" (i.e. nestedness of the hypertheories) is essential for satisfying *AGM* postulates, this means that one should look for a different definition for *AGM* expansion.

AGM-type Expansion Operations on Standard Onion Models

In fact, any of the known semantic proposals for expansion (as an operation on Grove sphere models) considered in the Belief Revision literature can be internalized in *DDL*. In particular, *for each of the three types of revision defined above there is a corresponding expansion action on standard onion models*:

$$(O, O') \in R^{+!P} \text{ iff } (O, O') \in R^{!P} \text{ and } X \cap P \neq \emptyset \text{ for all } X \in O,$$

$$(O, O') \in R^{+\uparrow P} \text{ iff } (O, O') \in R^{\uparrow P} \text{ and } X \cap P \neq \emptyset \text{ for all } X \in O,$$

$$(O, O') \in R^{+\uparrow P} \text{ iff } (O, O') \in R^{\uparrow P} \text{ and } X \cap P \neq \emptyset \text{ for all } X \in O.$$

See the Appendix (Figures 1, 2 and 3) for visualizations of these operations. Since expansion is a special case of revision (namely the case in which the new information does not contradict any prior beliefs), the corresponding expansion modalities can be reduced to the revision ones, e.g.

$$[+!\varphi]\theta \iff (\neg B\neg\varphi \Rightarrow [!\varphi]\theta).$$

Segerberg Contraction on Standard Onion Models: "Severe Withdrawal"

It is easy to see that, on standard onion models, contractions with P exist and are unique whenever P is not irrevocably known (i.e. whenever $(\bigcup O) \cap (U - P) \neq \emptyset$). Moreover, on standard onion models, Segerberg's definition is equivalent to putting:

$$(O, O') \in R^{-P} \text{ iff } O' = O \cap (U - P) := \{X \in O : X \cap (U - P) \neq \emptyset\}.$$

This semantic contraction operation was called "mild contraction" by Levi [12], "severe withdrawal" by Pagnuco and Rott [16] and "Rott contraction" by Ferme and Rodriguez [10]. See the Appendix (Figure 7) for a picture of severe withdrawal.

Static Withdrawal To axiomatize the dynamic logic of contraction, we need to introduce a “static” contraction modality $B^{-P}Q$ that pre-encodes Segerberg’s dynamic contraction in the same way in which conditional belief B^PQ pre-encodes dynamic revision. We call this operator “withdrawn belief” and we read $B^{-P}Q$ as saying that Q is believed conditionally on the withdrawal of P . We do this by putting:

$$s, O \models B^{-\theta} \phi \text{ iff } \forall \text{maximal } \neg\theta\text{-f.i.p. } F \subseteq O \exists F' \text{ finite } \subseteq F \cap \|\neg\theta\|_O \forall t \in \bigcap F' (t, O \models \phi)$$

Proposition In onion models, *withdrawn belief (after withdrawing P) is the same as truth in all the states of some sphere not included in P :*

$$s, O \models B^{-\theta} \phi \text{ iff } \exists X \in O (X \not\subseteq \|\theta\|_O \text{ and } \forall t \in X (t, O \models \phi)).$$

Moreover, in onion models satisfying the Limit Condition (and in particular, in *standard* onion models), *withdrawn belief (after withdrawing P) is the same as truth in all the states of the smallest sphere not included in P :*

$$s, O \models B^{-\theta} \phi \text{ iff } \forall t \in \bigcap (O \cap \|\neg\theta\|_O) (t, O \models \phi).$$

Observation (“Static” *Levi Identity*) It is easy to see that standard conditional belief can in fact be defined in terms of the withdrawn belief operator, via the following semantic equivalence:

$$B^{\theta} \phi \iff B^{-\theta} (\theta \Rightarrow \phi).$$

The converse is false: *one cannot define withdrawn belief only in terms of conditional belief.*

Open Question Finding a complete axiomatization for (static) withdrawn belief is still an open problem.

Reducing Segerberg’s Dynamic Contraction to Static Withdrawal However, if given such a complete axiomatization for static withdrawal, then we could immediately obtain a complete axiomatization for Segerberg’s dynamic contraction logic, by adding a number of Recursion laws, the most important being:

$$[-\phi]B^{-\psi}\theta \iff K\phi \vee B^{\phi \vee [-\phi]}\psi[-\phi]\psi$$

However, many authors consider severe withdrawal to be a bad candidate for modeling contraction. In addition to *not* satisfying the Recovery principle, it *does* satisfy a highly implausible property, called *Expulsiveness*: for ontic facts p, q , we have that $\neg Bp \wedge \neg Bq$ implies $[-p]Bq \vee [-q]Bp$. This property does not allow unrelated beliefs to be undisturbed by each other’s contraction!

To this objection, we can add another one, based on dynamic logic. Namely, although severe withdrawal satisfies a dynamic version of the so-called *Levi identity* with respect to irrevocable revision (DEL update)

$$R^{-P}; R^{+P} = R^{!P}$$

(where $R;R'$ is relational composition and $P \subseteq U$ is an arbitrary set of states), the corresponding Levi identities for lexicographic revision or minimal revision are *not* satisfied:

$$\begin{aligned} R^{-\neg P}; R^{+\uparrow P} &\neq R^{\uparrow P}, \\ R^{-P}; R^{+\uparrow P} &\neq R^{\uparrow P}. \end{aligned}$$

Since update (irrevocable revision) is a rather implausible operation when dealing to belief change in daily life, this throws more doubt on the appropriateness of Segerberg's definition of contraction.

Other AGM-type Contractions

But severe withdrawal is not the only AGM-friendly semantic contraction operation in the literature. Other options include *conservative contraction* $-_cP$ and *moderate contraction* $-_mP$ (see the pictures). We give below the formal definition over onion models in DDL (but see the pictures for a better intuitive explanation): if we put $O_{-P} := \bigcap O^{U-P}$ (for the smallest non-empty intersection of an O -sphere with $U - P$) whenever $O^{U-P} \neq \emptyset$ (i.e. whenever $\bigcup O \not\subseteq P$), and $O_{-P} := \emptyset$ otherwise, then for any two standard onions $O, O' \in D$ we define

$$(O, O') \in R^{-cP} \text{ iff } O' = \{X \cup O_{-P} : X \in O\},$$

$$(O, O') \in R^{-mP} \text{ iff } O' = \{Y \cup \bigcap O : Y \in O^{U-P}\} \cup \{X \cup \bigcup O^{U-P} : X \in O\}.$$

See the Appendix for visualizations of these operations (Figures 8 and 9). They are much better behaved than severe withdrawal. They satisfy the Recovery postulate, and moreover *they satisfy the dynamic versions of Levi identity for all the above-mentioned revision operators*: e.g. for all sets $P \subseteq U$ of states, we have

$$\begin{aligned} R^{-c\neg P}; R^{+!P} &= R^{!P}, & R^{-m\neg P}; R^{+!P} &= R^{!P}, \\ R^{-c\neg P}; R^{+\uparrow P} &= R^{\uparrow P}, & R^{-m\neg P}; R^{+\uparrow P} &= R^{\uparrow P}, \\ R^{-cP}; R^{+\uparrow P} &= R^{\uparrow P}, & R^{-mP}; R^{+\uparrow P} &= R^{\uparrow P}. \end{aligned}$$

Even better, there is no need to introduce the “static” counterparts of these operators as new primitive operators: *their static versions are definable in terms of conditional beliefs*. This means that the logic of conditional beliefs, conservative contraction and moderate contraction can be directly axiomatized, in a similar way that the logic of (various types of) dynamic revision was axiomatized:

Theorem There exists a sound and complete proof system for conservative contraction and moderate contraction over the class of DDL onion models that are closed under these operations. The system consists of the axioms of conditional doxastic logic *CDL*, together with the following Recursion laws:

$$\begin{aligned}
[-_c\phi]p &\iff p, & [-_m\phi]p &\iff p, \\
[-_c\phi]\neg\theta &\iff \neg[-_c\phi]\theta, & [-_m\phi]\neg\theta &\iff \neg[-_m\phi]\theta, \\
[-_c\phi](\theta \wedge \psi) &\iff [-_c\phi]\theta \wedge [-_c\phi]\psi, & [-_m\phi](\theta \wedge \psi) &\iff [-_m\phi]\theta \wedge [-_m\phi]\psi, \\
[-_c\phi]B^\psi\theta &\iff B([-_c\phi]\psi \Rightarrow [-_c\phi]\theta) \wedge B^{-\phi}([-_c\phi]\psi \Rightarrow [-_c\phi]\theta) \wedge \left(B^{-\phi}[-_c\phi]\neg\psi \Rightarrow B^{[-_c\phi]\psi}[-_c\phi]\theta \right), \\
[-_m\phi]B^\psi\theta &\iff B([-_m\phi]\psi \Rightarrow [-_m\phi]\theta) \wedge B^{-\phi \wedge [-_m\phi]\psi}[-_m\phi]\theta \wedge \left(K^{-\phi}[-_m\phi]\neg\psi \Rightarrow B^{\phi \wedge [-_m\phi]\psi}[-_m\phi]\theta \right)
\end{aligned}$$

1.8 Evidential dynamics in DDL

In the paper [8], van Benthem and Pacuit develop a very interesting extension of DEL aimed to deal with evidential dynamics. Their evidence models are based on the well-known neighbourhood semantics for modal logic, in which the neighbourhoods are interpreted as “evidence sets”: pieces of evidence (possibly false, possibly mutually inconsistent) possessed by the agent. In this section we briefly sketch how their setting can be internalized in DDL.

Belief modality, revisited (in general DDL models) We revert here to general DDL models, based on hypertheories H whose fallbacks are not necessarily nested. But now the fallbacks $X \in H$ is interpreted as “evidence sets”, and each hypertheory H is interpreted as a possible “evidential state” (rather than just a doxastic state): one in which the agent possesses a piece of evidence X iff $X \in H$. Our general definition of belief operators B (and their conditional-belief generalizations) in terms of maximal f.i.p. families is in fact taken from [8]. But now this definition has a clearer justification: when confronted with mutually inconsistent pieces of evidence, a rational agent believes the sentences that are implied by *all* the maximally consistent bodies of available evidence. So belief $B\phi$ is defined as “truth in all the states that are contained in any maximally consistent family of evidence sets”.

Evidence Modality In addition, van Benthem and Pacuit introduce a \square operator, such that $\square\phi$ means that *the agent has evidence for ϕ* . We adopt this notion in general DDL models, except that we denote it by \boxplus (to distinguish from the “effort” modality \square from Topo-logic):

$$s, H \models \boxplus\phi \text{ iff } \exists X \in H \forall t \in X (t, H \models \phi).$$

Note that, by a previous Proposition, $\boxplus\phi$ and $B\phi$ are equivalent in onion models. This is natural: onion models represent situations in which all the available pieces of evidence are mutually consistent; hence, in an onion model belief in ϕ is the same as “having evidence for ϕ ”, while in general these two notions are distinct.

Conditional Evidence Again, we can follow van Benthem and Pacuit and generalize \boxplus to a conditional evidence modality $\boxplus^\theta\phi$, expressing the fact that *the agent has some evidence for ϕ that is compatible with θ* :

$$s, H \models \Box^\theta \varphi \text{ iff } \exists X \in H^{\|\theta\|_H} \forall t \in X (t, H \models \varphi).$$

Evidence Management Actions Further, Van Benthem and Pacuit proceed to formalize a number of “evidence management” actions: they denote *evidence addition* by $+\varphi$, *evidence removal* by $-\varphi$, *evidence upgrade* by $\uparrow\varphi$ and *evidence combination* by $\#$. We briefly sketch here how can these be defined in DDL models. To distinguish the first three of these evidential operations from the doxastic operations that we previously considered, we add a subscript e (from “evidence”).

$$(H, H') \in R^{+eP} \text{ iff } H' = H \cup \{P\},$$

$$(H, H') \in R^{-eP} \text{ iff } H' = H - \{X \in H : X \subseteq P\},$$

$$(H, H') \in R^{\uparrow eP} \text{ iff } H' = \{X \cup P : X \in H\} \cup \{P\},$$

$$(H, H') \in R^\# \text{ iff } H' \text{ is the closure of } H \text{ under non-empty intersections.}$$

An *evidential DDL model* is one whose doxology is closed under these relations. As before, we introduce universal modalities $[+\varphi]$, $[-\varphi]$, $[\uparrow_e \varphi]$, $[\#]$ for the binary relations $R^{+e\|\varphi\|}$ etc. Essentially, evidence addition $+\varphi$ is the action by which φ comes to be accepted as an admissible piece of evidence; evidence removal $-\varphi$ is the action by which all evidence entailing φ is removed; evidence upgrade $\uparrow_e \varphi$ incorporates φ into each piece of available evidence (thus making φ the most important piece of evidence); finally, evidence combination $\#$ is the action by which the agent combines all the mutually consistent pieces of evidence.

Proposition All the Recursion laws for evidence management actions presented in [8] are valid on evidential DDL models.

1.9 Conclusions: comparing DDL, DEL and PDL/ETL

Three Styles of Doing Static Belief Revision As already mentioned, in the literature we encounter three styles of modelling “static” belief revision: Stalnker’s *selection functions*, *plausibility relations*, and the Grove-Lewis *sphere models*. When considered at an appropriate level of generality, these settings are equivalent. In this paper, we followed Segerberg in considering only sphere models. At a syntactic level, we followed the DEL approach (inspired from the Conditional Logic tradition) of using conditional belief operators to express static revision.

Three Styles of Doing Doxastic Dynamics As also mentioned, we are aware of three different styles of modeling doxastic changes. The first is the DEL approach, in which the dynamics is *external* to the models: doxastic actions are seen as model-changing actions, and represented as relations *between models*. The second style is the (doxastic version of the) old PDL (Propositional Dynamic Logic) approach: the dynamics is internalized simply by adding enough states to the model to represent the results of all the possible doxastic actions, which are thus represented *inter-*

nally, as binary relations *between states*. A variant is the ETL style (of Epistemic Temporal Logic), obtained by unravelling PDL models into trees, and by lumping together all the dynamic relations into one single temporal relation (going from a state to the possible next states). Finally, the third style is given by Segerberg’s DDL: this approach keeps the actual states unchanged (as “*ontic states*”) and internalizes the dynamics by representing doxastic actions as binary relations *between doxastic structures* (“onions”, “hypertheories”, “*doxastic states*”) living in a fixed space of possible such structures (the “doxology”).

Again, if considered at an appropriate level of generality, these three approaches are equivalent. However, there are some conceptual (and practical) differences. The DEL approach is the most “open-ended”, well-suited for open systems, in which there are innumerable doxastic actions that might happen. It is also the most “economical”, as only the states and the doxastic structures that are currently epistemically possible are “given”: only they are represented in a given model; hence, the DEL models can be easily visualized and drawn. It is also a “constructive” approach: the doxastic dynamics is *not given* in this approach, but is to be *constructed* (in the form of various model transformers, or “upgrades”).

The PDL/ETL approach has the advantage that it internalizes all the possible dynamics, in a clear way, using an almost flat structure (with only two levels: states, and relations between them). But the price, especially in the ETL version, is that the models are typically huge and quickly risk becoming unmanageable. This is the most un-economical of the three dynamic styles: we cannot actually draw PDL/ETL models for almost any realistic scenario involving iterated belief change.

The DDL style is somewhere in between. It is much more economical than the ETL approach, since it keeps the states fixed and only multiplies the doxastic structure. It also brings conceptual clarity: doxastic changes *are* after all only changes of belief, so they shouldn’t multiply the states of the world. It is an elegant and natural way to internalize doxastic changes. As shown in this paper, it is potentially at least as expressive and powerful as the (single-agent version of the) DEL approach: everything that was ever done in DEL style can be done in DDL style.

However, there is still a price to pay. DDL models are very high-level, involving fourth-order entities (not only states, but also fallbacks as sets of states, and hypertheories as sets of fallbacks, and doxologies as sets of hypertheories..., as well as doxastic actions as binary relations between hypertheories!). As a result, these models are hard to visualize. Their economy of resources is also relative: in complex dynamic-doxastic scenarios, the number of needed hypertheories explodes. Finally, it seems to us that, although in the end equivalent to the DEL style, the DDL approach lacks some of the heuristic value of DEL. As we saw, all the conceptual clarifications and settings that were developed in (single-agent) DEL style (such as the distinction between static and dynamic revision, the use of static conditional-belief modalities to “pre-encode” the dynamics, the axiomatization of various types of belief upgrades, the development of evidential dynamics) *can* be done in DDL style. We are confident that other such developments (such as the doxastic dynamics of questions studied in Interrogative DEL) *can* also be done in DDL style. But there might be a reason for which these developments were first done in DEL style: the

inherent complexity of DDL models, their higher-order nature and the difficulty of visualizing them may reduce their heuristic value and may risk becoming obstacles to the intuitive pursuit of new developments in the field. An open-ended approach such as DEL, which keeps to a minimum the number of entities (and the number of higher-order concepts) in a given model and keeps the dynamics outside the models, may be easier to use when pursuing new developments. But this is just the context of discovery. At a later stage, in the context of presentation and justification, it may again become important, at least for the sake of conceptual clarity, to *re-internalize* these new dynamic developments, using the elegant DDL style. As indeed we attempted here, in this paper.

Acknowledgements Sonja Smets' contribution to this paper was funded by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement nr 283963.

References

1. C.E. Alchourron, P. Gardenfors, D. Makinson. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *The Journal of Symbolic Logic*, **50**, No 2, 510-530.
2. A. Baltag, L.S. Moss and H.P. van Ditmarsch. 2008. Epistemic Logic and Information Update. In (eds.) P. Adriaans and J. van Benthem, *Philosophy of Information*, part of Handbook of the Philosophy of Science, Elsevier, vol. 8, 361-465.
3. A. Baltag and S. Smets. 2006. Conditional Doxastic Models: A Qualitative Approach to Dynamic Belief Revision. In (eds.) G. Mints and R. de Queiroz, Proceedings of WOLLIC 2006, *Electronic Notes in Theoretical Computer Science*, **165**, 5-21.
4. A. Baltag and S. Smets. 2006. Dynamic Belief Revision over Multi-Agent Plausibility Models. In (eds.) G. Bonanno, W. van der Hoek, M. Wooldridge, Proceedings of LOFT 2006, 11-24. University of Liverpool.
5. A. Baltag and S. Smets. 2008. A Qualitative Theory of Dynamic Interactive Belief Revision. In (eds.) G. Bonanno, W. van der Hoek, and M. Wooldridge, *Texts in Logic and Games*, Amsterdam University Press, **3**, 13-60.
6. J. van Benthem. 2007. Dynamic Logic for Belief Change. *Journal of Applied Non-Classical Logics*, **17**, nr. 2, 129-155.
7. J. van Benthem. 2011. *Logical Dynamics of Information and Interaction*. Cambridge University Press.
8. J. van Benthem and E. Pacuit. 2011. Dynamic Logics of Evidence-Based Belief. *Studia Logica*, **99**, nr 1, 61 - 92.
9. H. van Ditmarsch, W. van der Hoek, B. Kooi. 2008. *Dynamic Epistemic Logic*, Springer Coll. Synthese Library: Studies in Epistemology, Logic, Methodology, and Philosophy of Science, Dordrecht, **337**.
10. E. Ferme and R. Rodriguez. 1998. A Brief Note About Rott Contraction, *Logic Journal of the IGPL*, **6**, 835842.
11. J. Hintikka. 1962. *Knowledge and Belief*. Cornell University Press, Ithaca, NY.
12. I. Levi, 2004. *Mild Contraction: Evaluating Loss of Information due to Loss of Belief*, Oxford University Press, Oxford.
13. S. Lindström. and W. Rabinowicz. 1999. DDL Unlimited: Dynamic Doxastic Logic for Introspective Agents. *Erkenntnis* **50**, 353-385.
14. S. Lindström. and W. Rabinowicz. 1999. Belief change for introspective agents. In *Spinning Ideas*, Electronic Essays Dedicated to Peter Gärdenfors on His Fiftieth Birthday.

15. L.S. Moss and R. Parikh. 1992. Topological reasoning and the logic of knowledge. In (ed.) Moses, Y., *Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, 95 - 105.
16. M. Pagnucco, H. Rott. 1999. Severe Withdrawal - and Recovery, *Journal of Philosophical Logic*, **28**,501-547. (Corrected reprint in issue February 2000).
17. R. Parikh, L.S. Moss, C. Steinsvold. 2007, Topology and Epistemic Logic. In (eds.) M. Aiello, I. Pratt-Hartmann, J. van Benthem, *Handbook of Spatial Logics*, Springer, Netherlands.
18. H. Rott. 2006. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In (eds.) D. Makinson, J. Malinowski and H. Wansing, *Towards mathematical philosophy (Trends in Logic)*, vol. IV, Dordrecht: Springer, 269-296.
19. K. Segerberg. 1995. Belief revision from the point of view of doxastic logic. *Bulletin of the IGPL*, **3**, 535-553.
20. K. Segerberg. 1996. A general framework for the logic of theory change. *Bulletin of the section of logic*, **25**, 2-8.
21. K. Segerberg. 1997, Proposal for a theory of belief revision along the lines of Lindström and Rabinowicz. *Fundamenta Informaticae*, **32**, 183 -191.
22. K. Segerberg. 1998. Irrevocable Belief Revision in Dynamic Doxastic Logic. *Notre Dame Journal of Formal Logic*, **39**, Number 3.
23. K. Segerberg. 1999. Two traditions in the logic of belief: bringing them together, In (eds.) H. J. Ohlbach and U. Reyle, *Logic, language and reasoning: essays in honour of Dov Gabbay*, Kluwer, Dordrecht, 135 -147.
24. K. Segerberg. 2001a. A completeness proof in full DDL. *Logic and Logical Philosophy*, **9**, 77-90.
25. K. Segerberg. 2001b. The basic dynamic doxastic logic of AGM. In (eds.) M.A. Williams and H. Rott, *Frontiers in belief revision*, Dordrecht: Kluwer, 57-84.
26. K. Segerberg. and H. Leitgeb. 2007. Dynamic doxastic logic - why, how and where to? *Synthese*, **155**, 167-190.

Appendix: Pictures of the Main Operations on Onions

The pictures drawn here are following Hans Rott’s presentation [18]. The spheres of the initial onion are drawn as usual, as nested circles. The numbers represent the spheres of the new onion, *after* the revision/expansion/contraction: e.g. all regions labeled with 1 form the first sphere of the new onion, the regions labeled with 2 form the second sphere etc. Finally, the regions labeled with \emptyset contain the states that are outside the union of all the spheres of the onion (the “impossible states”).

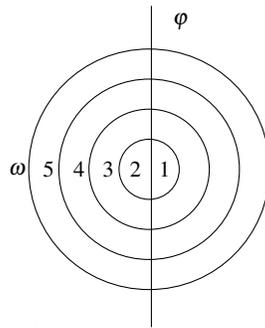


Fig. 1.1 Conservative Expansion ($+\uparrow\varphi$)

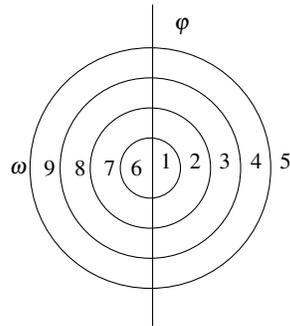


Fig. 1.2 Moderate Expansion ($+\uparrow\varphi$)

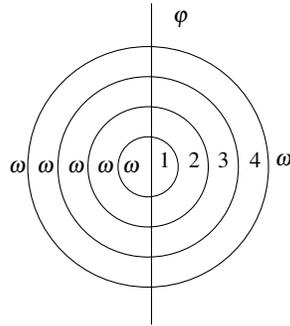


Fig. 1.3 Radical Expansion ($+\! \varphi$)

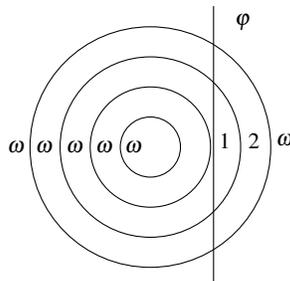


Fig. 1.4 Radical Revision ($\! \varphi$)

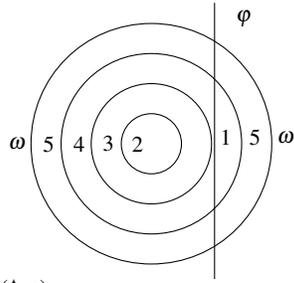


Fig. 1.5 Conservative Revision ($\uparrow \varphi$)

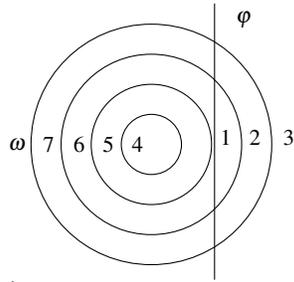


Fig. 1.6 Moderate Revision ($\uparrow \varphi$)

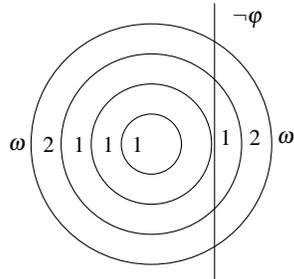


Fig. 1.7 Severe Withdrawal ($-\varphi$)

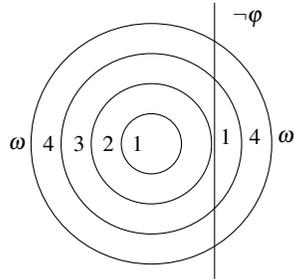


Fig. 1.8 Conservative Contraction ($-c \varphi$)

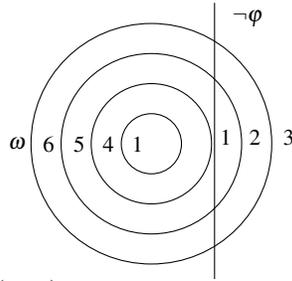


Fig. 1.9 Moderate Contraction ($-_m\varphi$)