

Contraction, Revision, Expansion: Representing Belief Change Operations

Sven Ove Hansson

Abstract The underlying idealizations in Krister Segerberg's Dynamic Doxastic Logic (DDL) are investigated in comparison with other belief revision models. It is argued that the doxastic voluntarism of the proposed interpretation is problematic but can be discarded. The treatment of conditional operators in DDL is discussed, and it is proposed that the use of conditional operators not satisfying the Ramsey test should be further investigated.

1 Introduction

Krister Segerberg's Dynamic Doxastic Logic (DDL) is a major alternative to the AGM model that is the current standard in studies of belief change. In order to investigate its properties we need to have a clear view of the basic idealizations that are common to belief revision theories. That is the subject of Sects. 2 and 3. In Sect. 4, DDL is introduced. After that two of its major features are scrutinized, namely its doxastic voluntarism (Sect. 5) and its treatment of non-truthfunctional connectives such as conditionals (Sect. 6). Finally, some general conclusions are drawn (Sect. 7).

2 Sentences and Epistemic Priorities

Logic is an astoundingly efficient and versatile tool for modelling a wide array of phenomena. However, like any modelling tool it puts emphasis on some aspects of the object of modelling at the expense of others. One of the major characteristics of logical models is that they impose linguistic structure on their subject-matter. This is particularly prominent in logical modelling of belief and knowledge.

S. O. Hansson(✉)
Royal Institute of Technology (KTH), Stockholm, Sweden
e-mail: soh@kth.se

Just a few minutes ago I looked out of the window, saw two roe deer in the garden and believed what I saw. In standard models of belief change, this event is represented by the addition of some sentence p (“There are two roe deer in the garden”) to my set of beliefs. My previous belief state is represented by a set K containing the sentences I believed to be true.¹ When I see the two deer, p is added to K . More precisely, assuming that the resulting set of beliefs is closed under logical consequence, K is exposed to the *input* p and is then replaced by $\text{Cn}(K \cup p)$. This is the simplest form of belief change (expansion), but it involves massive idealizations. Most importantly, if by the input we mean that which makes me believe that there are two deer in the garden, then the input is neither p nor any other sentence or set of sentences; what affected me was a visual impression with no linguistic encoding whatsoever. Furthermore, the resulting belief change may not be perfectly representable by a sentence (or set of sentences). I may have a “mental picture” of how the deer moved around that is not primarily linguistic and may be difficult to translate into words.² This, by the way, is why the police use identity parades, photo-lineups and similar methods in addition to asking witnesses to verbally describe a suspect. A witness may know what the culprit looks like without being able to express this knowledge in words.

But in the belief change literature, both belief states and inputs are taken to be sentential. The totality of the beliefs held by an agent is taken to be represented by a belief set that is a logically closed set of sentences, mostly assumed to be consistent. The inputs refer to a sentence³ that either has to be added to the belief set or removed from it. This gives rise to two basic types of input-based belief changes:

incorporation: The result is that a belief is accepted.

contraction: The result is that a belief is not accepted.

Four basic integrity constraints are usually imposed on the outcome of a belief change operation:

logical closure: The outcome is a logically closed set, just like the original belief set.

consistency preservation: The outcome is consistent, just like the original belief set.

success: (i) A sentence to be incorporated is included in the outcome. (ii) A sentence to be contracted is not included in the outcome.

conservatism: (i) In incorporation, no sentences are removed. (ii) In contraction, no sentences are added.⁴

¹ Or more precisely: the sentences I was committed to believe to be true (Cf. [18]).

² If belief change is interpreted as referring to other belief-holders than individual persons, then the sentential format may be less problematic. One example of this is database management. The contents of databases are more readily representable by sentences than human beliefs or memories. Another example is changes in collectively created and maintained stocks of information or knowledge, such as the corpus of scientific beliefs. Collective processes are usually based on sentential representations since these are needed for interindividual communication.

³ Or set of sentences [9].

⁴ These are the most elementary demands of conservatism. In addition the following, somewhat less precise, demands are common: (iii) In incorporation, no sentences are added unless this is needed to incorporate the input. (ii) In contraction, no sentences are removed unless this is needed to remove the input.

Fig. 1 The two alternative priorities among basic requirements on belief change that standard belief revision theory (such as AGM) vacillates between

Pattern A	Pattern B
logical closure	logical closure
∨	∨
success	success
∨	∨
conservatism	consistency preservation
∨	∨
consistency preservation	conservatism

In contraction, these four requirements are all compatible if the sentence to be removed is non-tautologous. If the sentence is a tautology, then *logical closure* and *success* are incompatible (but each of them is compatible with the other two conditions). The standard solution is to give higher priority to logical closure than to success, i.e. the outcome of contraction by a tautology is a logically closed set and therefore it does not satisfy the success criterion.

In incorporation, all four requirements are compatible if the sentence to be added is consistent with the original belief set (i.e. if $K \cup \{p\}$ is consistent). If p is inconsistent, then *consistency preservation* and *success* cannot both be satisfied. This is traditionally solved by giving priority to *success* (which is compatible with the other two conditions). If p is consistent but inconsistent with the original belief set, then any two of the three conditions *consistency preservation*, *success*, and *conservatism* are compatible, but not all three of them. (*Logical closure* is compatible with each of these combinations). There are two standard solutions to this. One is to give up *consistency preservation*, usually by just letting the outcome be $Cn(K \cup \{p\})$. This form of incorporation is called expansion. The other solution is to give up *conservatism*, and remove enough elements from the original belief set K to ensure that p can be added without giving rise to inconsistency. This type of incorporation is called revision.

Summarizing this, the priorities inherent in these operations can be described as vacillating between the two patterns shown in Fig. 1. The standard operation of contraction is compatible with both patterns, whereas expansion is compatible only with Pattern A and revision with Pattern B.

3 Decomposing Belief Change

The standard framework of belief revision theory originates largely in Isaac Levi’s [18] work from the 1960s and 1970s. He established a framework in which belief states are represented by logically closed belief sets. There are three types of changes: contraction (\div), expansion ($+$), and revision ($*$). Expansions are performed in the

simple way already indicated, i.e. $K + p = \text{Cn}(K \cup \{p\})$. Furthermore, revisions are definable in terms of contractions and expansions through what is now called the Levi identity:

$$K * p = (K \div \neg p) + p. \quad (1)$$

The Levi identity can be seen as based on an underlying assumption of decomposability into simple operations. It can perhaps be defended as follows: Real-life belief change results in new beliefs being added and old ones being removed. Therefore, we can assume without losing generality that all operations of change consist of two suboperations: “pure” contraction that removes beliefs but does not add any new ones, and “pure” incorporation that adds new beliefs but removes no old ones.⁵

Seeger somewhat cautiously endorsed this decomposition principle although in slightly different terms. After discussing straightforward cases in which only removals or additions of beliefs are needed, he said:

There are certainly more complex cases when the agent will go to a new belief-set that is neither weaker nor stronger than the current one; but those can perhaps be seen as derivative, as achievable by a combination of weakenings and strengthenings. ([35], 143)

In the same paper he proposed as a desideratum for belief revision theory “that there be two basic kinds of doxastic action, *basic expansion* and *basic contraction*” (p. 144). Basic contractions (weakenings) of the belief set are representable in possible world models as retreats to sets of worlds that contain the set of worlds that represent the current belief state. Following [21] he called such retreats *fall-backs*. Basic expansions (strengthenings) could analogously be represented by sets of worlds included in the one representing the current belief state. Seeger called them *push-ups*.

It is important to distinguish between two interpretations of the postulated decomposability of all belief changes into contraction and expansion. We can call them the “black box” and the “step-by-step” interpretation. According to the black box interpretation, the decomposition provides us with a convenient method to obtain the desired outcome, but it does not necessarily correspond to how changes in belief actually take place. According to the step-by-step interpretation the decomposition is a representation of how belief change actually takes place, specifying the actual suboperations and the order in which they take place. The black box interpretation is fairly plausible. Irrespective of how a human being goes from a belief set K_1 to another belief set K_2 , in a formal model we can go from K_1 to K_2 by performing first a contraction that takes us from K_1 to some K' such that $K' \subseteq K_1 \cap K_2$, and then an expansion that takes us from K' to K_2 . For this to be feasible it is sufficient that there are two sentences p and q such that contraction of K_1 by p leads us to

⁵ An alternative approach takes as primitive an operation that both removes some sentence(s) and adds some other sentence(s). It is then possible to develop a model of belief change on the basis of one single primitive operation instead of two as in Levi’s model [13].

some K' that is also a subset of K_2 and that expansion of K' by q leads us further on to K_2 .⁶

The step-by-step interpretation of the decomposition is much more problematic. One of the reasons for this is that the required composite operations, “pure” contraction (in which no sentences are added) and “pure expansion” (in which no sentences are removed) do not seem to be matched by actual operations of change. Although contraction is taken for granted as a building-block in belief change theory, it is not easily exemplified. Of course there are belief changes in real life that are driven by a need to give up a certain belief. However, such changes tend to be caused by the acquisition of some new information that is added to the belief set. Not long ago a friend said to me that he was quite sure that the Vatican City State is a member of the United Nations, which I believed it was not. This made me uncertain and induced me to enter a state of hesitation concerning the issue in question. I therefore removed the sentence “The Vatican City State is not a member of the United Nations” from my set of beliefs without adding its negation. In the belief revision literature, this would be treated as a contraction, but in fact it was not since I added the new belief that my friend believes that the Vatican City State is a member of the United Nations. The only credible examples of pure contraction that have been presented in the literature are hypothetical contractions such as contractions for the sake of argument [8, 20].⁷ Pure incorporation, i.e. expansion, is also problematic, as will be seen in Sect. 6.

A crucial step in the theory of belief change was taken by Carlos Alchourrón, Peter Gärdenfors and David Makinson [1] who provided what is now the standard framework of belief revision. Their major invention was a formally precise account of contraction, namely partial meet contraction:

$$K \div p = \bigcap \gamma(K \perp p), \quad (2)$$

where $K \perp p$ is the set of inclusion-maximal subsets of K not implying p and γ is a selection function, such that $\emptyset \neq \gamma(K \perp p) \subseteq K \perp p$ whenever $K \perp p \neq \emptyset$, and $\gamma(K \perp p) = K$ when $K \perp p = \emptyset$. Revision is defined according to the Levi identity, i.e. $K * p = (K \div \neg p) + p$. This framework has turned out to be exceptionally fruitful, and AGM-style belief revision is a rapidly developing research area with a surprising number of ramifications and connections with other areas. [7]

But it should be remembered that the standard framework is the result of a whole series of idealizations and limitations. The belief changes of real-life human beings are often not sentential. Furthermore, even given the choice of sentential representations there are many other options than the three standard ones of contraction,

⁶ For arbitrary K_1 and K_2 , this recipe only works if the language is finite. In a framework with an infinite language, two operations on sentences are not sufficient to take us from a belief set K_1 to any other belief set K_2 . If there is a countably infinite number of logical atoms, then the number of belief sets expressible in the language is uncountable. (This can be shown with a standard diagonalization argument.) On the other hand, there is only a countable number of sequences $\div p + q$ of a contraction by one sentence (p) followed by an expansion by another (q). This problem can be solved by introducing multiple contraction and expansion.

⁷ This refers to the modelling of human beliefs. Pure contraction of databases is unproblematic.

expansion, and revision. Alternative types of operators that may better represent some real-life belief changes include the following:

- *consolidation*, an operation that makes an inconsistent belief state consistent by removing beliefs from it.
- *external revision*, revision by a sentence p that proceeds by first expanding by p and then contracting by $\neg p$, i.e. the two suboperations take place in the reverse order to that of the Levi identity.
- *semi-revision*, an operation that receives a sentence p and weighs it against old information, with no special priority assigned to the new information due to its novelty. The input may be either incorporated or rejected.
- *selective revision*, a generalization of semi-revision in which it is possible for only a part of the input information to be accepted. (Selective revision by $p \& q$ may for instance result in p being incorporated and q rejected.)
- *shielded contraction*, a variant of contraction in which some non-tautological beliefs are not retractable. The agent may hold a non-logical belief p that nothing can make her give up, so that $p \in K \div p$, and presumably also $p \in K \div q$ for all q .
- *lowering* and *raising*, operations in which the belief set is unchanged but the degrees of belief in some of its elements are either decreased or increased, which may have effects on the outcomes of subsequent changes.
- *replacement*, an operation that replaces one sentence by another in a belief set. Excepting limiting cases, the outcome of replacing p by q is a belief set $K \mid_q^p$ such that $p \notin K \mid_q^p$ and $q \in K \mid_q^p$. Replacement can serve as a “Sheffer stroke” for the standard belief revision operators.
- *reconsideration* reintroduces previously removed beliefs if there are no longer any valid reasons for their removal.
- *multiple contraction*, in which a set of sentences, rather than a single sentence, is (simultaneously) removed from the belief set.
- *indeterministic belief change*, in which there are several alternative outcomes of a change operation. In indeterministic contraction, $K \div p$ is typically a set of belief sets that are subsets of K and do not contain p , rather than a single such belief set.

(For references on these operations, see [15].)

In summary, belief revision theory is dominated by an elegant and highly idealized framework (AGM) that only covers some of the many aspects of actual belief change. This is the background against which we should study Krister Segerberg’s contributions to belief revision theory. He has paid much attention to possible extensions of the framework, such as consolidation [33], semi-revision [33], external revision [33], and indeterministic change [34]. But most importantly, he has provided us with an alternative framework in which the very notion of an operation of change is explicated quite differently from the AGM framework.

4 Dynamic Doxastic Logic

Given his background as one of the major contributors to the development of modern modal logic, it should be no surprise that Segerberg took the lead in approaches that employ the resources of modal logic to increase the expressive power of belief change theory. This resulted in *dynamic doxastic logic* (DDL) that includes two major additions to the language that increase its expressive power [32, 35]. (The term “dynamic doxastic logic” is modelled after van Benthem’s “dynamic modal logic”, cf. [32], p. 535.)

The first of these additions is sentence formation with epistemic modal operators of the type introduced by Hintikka [16]. The sentence $B_i p$ denotes that the individual i believes that p . When only one agent is under consideration, the subscript can be deleted, and the operator B can be read “it is believed that” or “the agent believes that” ([32], p. 536).

A major difference between Bp and the formula $p \in K$ of AGM is that the former but not the latter is a sentence in the same language as p . This makes it possible to express in the object language that a sentence is believed. In Segerberg’s own words, he tried to develop belief revision “as a generalization of ordinary Hintikka type doxastic logic”, whereas in contrast “AGM is not really logic; it is a theory about theories” ([35], p. 136). The difference becomes crucial when beliefs about beliefs are introduced. Sentences such as “ i believes that i does not believe that p ” and “ i believes that j believes that p ” are readily expressible in DDL as $B_i \neg B_i p$ respectively $B_i B_j p$. The AGM framework does not have the corresponding resources. (Neither $(p \notin K_i) \in K_i$ nor $(p \in K_j) \in K_i$ is a well-formed formula.)

The other addition is the formalization of belief revision operations (expansion, revision, and contraction) with dynamic modal operators, similar to those used for program execution. This element of DDL was present also in publications by several other authors ([8, 28, 37, 38]). The standard notation used by Segerberg is as follows:

- $[\div p]\alpha$ (α holds after contraction by p)
- $[\ast p]\alpha$ (α holds after revision by p)
- $[+p]\alpha$ (α holds after expansion by p)

The combination of these two elements, belief operators and dynamic operators, provides us with a framework that is in important respects more general than AGM. $[\ast p]Bq$ means that q is believed after revision by p , hence it conveys the same information as the AGM formula $q \in K \ast p$. ([17], 168) Similarly, $[\div p]\neg Bq$ says that q is not believed after contraction by p , i.e. $q \notin K \div p$. But in addition, the combined use of belief operators and dynamic operators makes it possible to express an agent’s beliefs about her own patterns of belief revision. As an example of this, $B([\ast p]Bq)$ means that the agent believes that after revision by p she will believe that q , and the more complex formula $[\ast[\ast p]Bq]Br$ means that the agent will believe r if she revises her belief state by the belief that if she revises by p then she will believe q . ([17], 169)

The success criteria of the three operations are succinctly expressed as follows:

$[*p]Bp$ (Revision success)

$[+p]Bp$ (Expansion success)

If $p \notin \text{Cn}(\emptyset)$ then $[\div p]\neg Bp$ (Contraction success)

According to the Levi identity, $[*p]$ can be read as an abbreviation of $[\div \neg p][+p]$ ([32], p. 357). In the same fashion, iterated operations can be expressed by repetition of the dynamic ($\langle \rangle$) operators, such as $[*p][*q][\div r]$ etc.

The recasting of belief revision theory as modal-style dynamic logic has the important advantage that it “puts at our disposal the rich meta-theory developed in the study of modal and dynamic logic”. Segerberg ([35], 142) As a simple example of this, the analogy between $\langle \circ p \rangle$ (where \circ is any of \div , $+$, and $*$) and \square suggests the introduction of operators of the form $\langle \circ p \rangle$ that stand in the same relation to $\langle \circ p \rangle$ as \diamond to \square , i.e.:

$$\langle \circ p \rangle q \text{ if and only if } \neg[\circ p]\neg q. \quad (3)$$

$\langle \alpha \rangle \beta$ is to be read “after the agent has carried out the action α , it may be the case that β , and consequently $\langle \circ p \rangle Bq$ should be read “after the agent has contracted/expanded/revised by p , it may be the case that the agent believes q ” ([34], pp. 187, 189). In standard, deterministic belief revision models, the extension of the language with $\langle \rangle$ -operators is not of much use. If $\circ p$ has a well-determined outcome, then $\langle \circ p \rangle Bq$ and $[\circ p]Bq$ have the same truth conditions. However, in indeterministic belief revision (that assigns to $\circ p$ a set of possible outcomes, rather than a single outcome) the $\langle \rangle$ -operators provide a highly useful increase in expressive power. (On indeterministic belief revision, see [21].)

It should be mentioned, though, that although DDL has more expressive power than AGM in some respects, there are other respects in which the opposite relation seems to hold. In AGM we can easily express non-prioritized belief changes, i.e. changes in which the input is not always accepted. We can have a semi-revision operation $*$ such that $p \in K * p$ does not hold for all p or a screened contraction operator \div such that $p \notin K \div p$ does not hold for all non-tautologous sentences p . Since $K * p$ simply represents the belief state obtained after receiving the information that p , this does not require any reinterpretation of the formalism. It is less obvious how to interpret $*p$ in DDL if $[*p]Bp$ does not hold; what type of action is then $*p$?

Important contributions to DDL have been made by Segerberg himself, by Sten Lindström and Wlodek Rabinowicz [22] who investigated formulas such $B([*p]Bq)$ that represent introspective agents, and by John Cantwell [6] who explored iterated change. In parallel, largely similar systems have been developed under the name of Dynamic Epistemic Logic, DEL ([5, 27, 40]). The original DEL models referred to belief expansion only, but in later work revision has been included ([3, 5, 37, 39]). A major difference between DDL and DEL is that the latter has mostly been studied in multiagent contexts.

DDL is a major alternative to the AGM-style formalisms that are currently the standard in belief revision theory. Since logical modelling of belief change operates with considerable idealizations, it is a wise strategy to promote the development of

models that put emphasis on different aspects of the subject matter. ([12, 14]) It is also a wise strategy to subject each of these alternative formalisms to critical scrutiny. In what follows, I will discuss two possibly problematic aspects of DDL, namely its concept of doxastic agency and its treatment of non-truthfunctional sentences in the object language.

5 Doxastic Voluntarism

Should belief changes be seen as actions undertaken by the epistemic subject or as uncontrollable effects of external influences? There is one formulation in one of his early papers on the subject where Segerberg kept both options open, describing belief changes as something that the agent can have “undertaken (or undergone)”. ([34], p. 183) However, in his development of DDL he settled for the former option, interpreting the interior of [] and () (for instance $*p$ in $[*p]$) as a representation of action.

“Suppose that you believe that a proposition P is true—thus $X \subseteq P$, where X is your current belief set—but that you have decided that this belief has to be given up. . . This means that you wish to replace X by a belief set Y such that P is no longer believed after the change. Call this operation *contraction* by P . . .

Suppose that you decide to accept the truth of P in the sense of simply adding it to your existing stock of beliefs. Again you are changing your views, you wish to replace your current belief set X by a belief set Y such that after the change you believe P as well as everything you already believe. We call this operation *expansion* by P . ”([34], p. 185)

“Now to expand or revise or contract is to do something. Thus it is possible to think of expansion, revision and contraction as actions of a certain kind—epistemic or doxastic actions.” ([35], p. 137)

In Segerberg’s theory, doxastic actions are a special type of actions. They differ from “real actions” in that they do not change the state of the world, as the latter may do. ([34], p. 187).⁸

Are there any doxastic (epistemic) actions? This is a much debated issue in philosophy, and the standpoint that there are such actions is usually called doxastic (epistemic) voluntarism. As these discussions have shown, it is important to distinguish between different variants of doxastic voluntarism. [26] First of all we need to identify the elements of human behaviour that are candidates for being such actions. Robert Audi [2] provided a useful distinction for this purpose, namely that between

⁸ Heinrich Wansing is another prominent proponent of this view. He has proposed that developing the semantics of belief ascriptions from the viewpoint of doxastic voluntarism can be a way to avoid closure of belief under logical consequence. In his view, a variant of seeing-to-it-that (stit) logic of agency can be used to represent voluntary acquisition and abandonment of belief. [41] In later work he has further specified this as dstit-theory, where dstit stands for “deliberately sees to it that” [42]. He proposes the introduction of a belief formation operator to be read “ α sees to it that α believes that p ” (p. 212).

a behavioural and a genetic version of doxastic voluntarism. According to the behavioural version, believing, i.e. holding a belief, is (or can be) an action-type. According to the genetic version, forming (rather than holding) a belief is (or can be) an action-type.

Both behavioural and genetic doxastic voluntarism can be further subdivided. In what follows I will focus on the variants of genetic doxastic voluntarism. Many authors have referred to the distinction between a weak and a strong variant, but it has often been overlooked that voluntarism can be weak or strong in two senses that give rise to crossing distinctions: Doxastic voluntarism can be either complete or partial. It can also be either direct or indirect. According to complete doxastic voluntarism all beliefs are voluntary, according to partial doxastic voluntarism only some of them. Doxastic voluntarism is direct if it implies that we can make ourselves adopt or give up a belief by just deciding to so. It is indirect if it indicates that we can do so only by performing will-controlled actions that cause, in ways that are not will-controlled, a change in belief. Obviously, both direct and indirect doxastic voluntarism can be either complete or partial.

What type of doxastic voluntarism does Segerberg need? Since his is a logic of belief change, rather than the statics of belief, the behavioural version is irrelevant for his theory. (It is also a version that has very rarely been defended by philosophers.) The doxastic actions that he refers to consist in the adoption or abandonment of beliefs. We should therefore focus on genetic doxastic voluntarism. This gives rise to two further questions: Should a genetic doxastic voluntarism that supports DDL be complete or partial, and should it be direct or indirect?

The answers to both these questions are quite obvious. DDL is a theory of belief change in general, not a theory intended to cover some fraction of the belief changes that epistemic subjects undertake. Therefore a doxastic voluntarism suitable for interpreting DDL will have to be complete. Furthermore, the framework is one of direct causation. $[*p]Bq$ means that the subject performs an action ($*p$) that has Bq as a consequence. Alternatively she may perform some action such as letting another person indoctrinate or hypnotize her to believe that p , but then her own action is not a doxastic action but a “real action” (in Segerberg’s terminology, quoted above) since it changes the state of the world rather than her own beliefs.

In summary then, the type of doxastic voluntarism that we need to support DDL is genetic, complete, and indirect. How credible is such a form of doxastic voluntarism?

If a student comes up to me after a lecture and tells me that my lecture was boring, I acquire the belief that she has said this to me. Reacting to such a sensory impression by questioning its veracity would be a sign of insanity rather than philosophical sophistication. This applies to most of the sensory evidence that we receive in everyday life. This is acknowledged by the majority of doxastic voluntarists. Philosophically credible argumentation in favour of direct doxastic voluntarism tends to stop short of defending the complete version of the thesis that would be necessary for Segerberg’s purposes. Hence, Ronney Mourad [24] concedes that “most beliefs are involuntary” (p. 60) and that this applies in particular when we have conclusive evidence in support of either belief or disbelief (p. 62). Similarly, Philip Nickel acknowledges that “conclusive evidence, when grasped by a doxastic subject, must induce belief” ([25],

p. 313). These and most other defenders of direct doxastic voluntarism do not claim that all beliefs are formed at will, only that some of them are. Unfortunately, this is not sufficient for DDL.

Even the partial version of direct doxastic voluntarism is highly contestable. Proponents often point out that (some) beliefs are not completely determined by evidence. This is incontrovertible. (“[I]t is far from clear that all beliefs of all agents come into being as an inescapable response to some evidence”, [42] p. 211.) However, there are many other influences on our beliefs than volition and evidence. Our beliefs are influenced by factors such as wishful thinking, intellectual sloppiness, and irrational trust in authorities. Influences such as these cannot in general be applied or deactivated at will in order to adopt or give up a particular belief. To substantiate (partial) direct doxastic voluntarism it would seem necessary to exhibit plausible examples of beliefs that are formed by direct volition-driven causation. Such examples do not seem easy to find. (Arguably the best attempts are so-called self-fulfilling beliefs that may arise for instance if someone credibly offers you \$1.000.000 for forming the belief that you are a millionaire. [30], p. 83.)

A strong case can be made in favour of indirect doxastic voluntarism, especially its partial variant. There are things we can do to induce beliefs in ourselves. Someone who wishes to become a believer in a certain religious faith can expose herself to arguments and emotional influence that is expected to make her a believer. Someone who is plagued by her own jealousy may try in different ways to convince herself that her husband is faithful. However, such indirect causation is not always successful, as exemplified by the phenomenon of being plagued by religious doubt.

Ethical arguments have had an important role in argumentation for doxastic voluntarism. There are situations when it is plausible to hold a person responsible for incorrect beliefs that have negative consequences. It may seem as if we can only be responsible for our beliefs if we have some kind of control over them. [23] However, such responsibility can at least in most cases be accounted for in terms of indirect doxastic voluntarism. What we require of persons with wrongful beliefs is that they study the evidence, listen to the experts, and reconsider the issue in a rational fashion. We do not normally demand that they change their belief by *fiat*.

In summary, Segerberg’s explication of DDL seems to require complete, direct doxastic voluntarism, which is an apparently implausible standpoint with very few adherents. There is much to say in favour of partial, indirect doxastic voluntarism, but that is not sufficient for DDL. Is there a way out of this conundrum? Can DDL be saved?

There is indeed a fairly simple way out: The interior of [] and < > need not be interpreted as representing actions. Instead they may be taken to represent external influences, in much the same way as in AGM. $[*p]Bq$ will then be interpreted for instance as “after receiving the information that p , the epistemic subject believes that q . As far as I can see there is nothing in Segerberg’s remarkably versatile formalism that precludes such an interpretation. However, it remains to investigate the more detailed consequences of its adoption.

6 Non-truthfunctional Connectives

Belief revision theory has primarily been concerned with beliefs expressed in an object language that contains no other resources than logical atoms and (the full set of) truth-functional connectives. This is a severe limitation since non-truthfunctional combinations are essential components of our belief systems, without which intentional actions as we know them would not be possible. This applies not least to conditional beliefs. I believe that I can light up the room by turning on the switch, and I also believe that consuming a bottle of wine will make me drunk. Such beliefs can usefully be formalized as conditional sentences. (“If I turn on that switch, then the room will be lit.” “If I drink that bottle of wine then I will be drunk.”) However, in spite of their essential role in our belief systems, such beliefs are disturbingly difficult to express in belief revision theory. In fact, any attempt to include non-truthfunctional expressions into the language seems to have drastic and often unwished-for effects on the formal system.

It is usually assumed that at least a large part of our conditional beliefs satisfy the so-called Ramsey test that is based on a suggestion by Ramsey that has been further developed by Robert Stalnaker [36] and others. The basic idea is that “if p then q ” is taken to be believed by the epistemic subject if and only if she would believe in q after revising her present belief state by p . Let $p \Box \rightarrow q$ denote “if p then q ”, or more precisely: “if p were the case, then q would be the case”. The Ramsey test says:

$$p \Box \rightarrow q \text{ holds if and only if } q \in K * p \quad (4)$$

In AGM, attempts have been made to include sentences of the form $p \Box \rightarrow q$ in the object language, which means that they will be included in the belief set when they are assented to by the agent, thus:

$$p \Box \rightarrow q \in K \text{ if and only if } q \in K * p \quad (5)$$

However, the step from (4) to (5), i.e. the inclusion in belief sets of conditionals that satisfy the Ramsey test, has turned out to require radical changes in the logic of belief change.⁹ As one example of this, contraction cannot then satisfy the inclusion postulate ($K \div p \subseteq K$). The reason for this is that contraction typically generates support for conditional sentences that were not supported by the original belief state. If I give up my belief that John is mentally retarded, then I gain support for the conditional sentence “If John has lived 30 years in London, then he understands the English language” [11].

A famous impossibility theorem by Peter Gärdenfors [10] shows that the Ramsey test is incompatible with a set of plausible postulates for revision. This was shown by Gärdenfors to hold if the underlying logic is (or contains) classical propositional logic. Segerberg [31] generalized this result, showing that it holds whenever the con-

⁹ Issac Levi [19] avoids most of the common difficulties with Ramsey test conditionals by accepting (4) but not (5).

sequence operator Cn of the underlying logic satisfies the three standard conditions $A \subseteq Cn(A)$ (inclusion or reflexivity), If $A \subseteq B$, then $Cn(A) \subseteq Cn(B)$ (monotony or monotonicity), and $Cn(Cn(A)) \subseteq Cn(A)$ (iteration or transitivity).

The crucial part of the proof consists in showing that the Ramsey test implies the following monotonicity condition.

$$\text{If } K \subseteq K' \text{ then } K * p \subseteq K' * p \tag{6}$$

The proof of this is straightforward: Let $K \subseteq K'$ and $q \in K * p$. The Ramsey test yields $p \Box \rightarrow q \in K$, then $K \subseteq K'$ yields $p \Box \rightarrow q \in K'$, and finally one more application of the Ramsey test yields $q \in K' * p$.

(6) is incompatible with the AGM postulates for revision, and it is also easily shown to be implausible.¹⁰ Let K be a belief set in which you know nothing about Ellen’s private life and K' one in which you know that she is a lesbian. Let p denote that she is married and q that she has a husband. Then we can have $K \subseteq K'$ but $q \in K * p$ and $q \notin K' * p$.

DDL was “introduced with the aim of representing the *meta*-linguistically expressed belief revision operator $*$ as an *object*-linguistic sentence operator $[*_]$ in the style of dynamic modal logic” ([17], 167–168). In other words, the driving idea of DDL is that a formula such as $[*p]q$ should be treated on the same level as its components p and q . ([17], p. 171) Furthermore, since the intended semantics of DDL is a possible world semantics, sets of possible worlds will be assigned to formulas such as $[*p]q$, just as this will be done for p and q . Therefore, we should expect the equivalent of (5) to hold in DDL, i.e.:

$$B(p \Box \rightarrow q) \text{ if and only if } [*p]Bq \tag{7}$$

Unfortunately, this gives rise to the same type of problem that Gärdenfors showed to hold in the AGM model. This can be seen from the fact that a conditional property closely related to (6) can be obtained, namely the following:

$$\text{If } [*p]Bq \text{ and } \neg B \neg r \text{ then } [*r][*p]Bq \tag{8}$$

The derivation of (8) is straight-forward.

Postulates

- $B(p \Box \rightarrow q) \leftrightarrow [*p]Bq$ (Ramsey test)
- If $\neg B \neg r$ and Bs then $[*r]Bs$ (preservation)
- Logical equivalence is preserved after substitution of logically equivalent subformulas (extensionality)

¹⁰ However, it is trivially unproblematic in an approach where it holds for all potential belief sets K_1 and K_2 that $K_1 \not\subseteq K_2$. In such a framework there cannot be any pure contraction, or any other operation that takes us from a belief set to one of its proper subsets. Furthermore, there cannot be any pure expansion, or any other operation that takes us from a belief set to one of its proper supersets. Cf. [11].

Derivation

- (1) $[*p]Bq$ (assumption)
- (2) $\neg B\neg r$ (assumption)
- (3) $B(p\Box\rightarrow q)$ ((1) and Ramsey test)
- (4) $[*r](B(p\Box\rightarrow q))$ ((2), (3), and preservation)
- (5) $[*r][*p]Bq$ ((4), Ramsey test, and extensionality)

It is easy to show that (8) is starkly implausible. Consider my beliefs about Rebecca, a casual acquaintance whom I met at a party. Initially I know nothing about her profession or about what musical abilities she may have; in particular I do not know whether or not she is tonedeaf (r). However, if I acquire the belief that she has applied for a position as concertmaster in the Czech Philharmonic (p), then I will also believe that she is a first-rate violinist (q). Hence, $[*p]Bq$ and $\neg B\neg r$. However, it does not hold that $[*r][*p]Bq$. The reason for this is that if I acquire both the beliefs that she is tonedeaf and that she has applied for the position in question, then I may conclude that she is an unqualified but self-conceited fiddler, and thus not believe q to be true.

Leitgeb and Segerberg ([17], 172) mention two ways to avoid difficulties like this. One is to “not allow the derivation, for all p and q , of a formula of the form $B(\chi(p, q)) \leftrightarrow [*p]Bq$, where $\chi(p, q)$ is some formula that is built syntactically from p and q ”.¹¹ There will then be some formulas of the type $[*p]Bq$ that cannot be an argument of the B operator, in other words $B([*p]Bq)$ is not a well-formed formula. This means that there are certain revision patterns that an agent may have but may not believe herself to have. This might have been an acceptable limitation if it only affected beliefs of an uncommon type, or only beliefs with a paradoxical flavour. However, as we have seen it will arise also with respect to seemingly unparadoxical everyday beliefs such as “If I am made to believe that she has applied for a position as concertmaster of the Czech Philharmonic, then I will also believe that she is a first-rate violinist”.

The second option that they mention is that the axioms and rules for $[*]$ may not conform to the AGM postulates. This is a much more plausible option. The AGM postulates (and their counterparts in DDL) have been developed for a restricted language that contains simple factual statements but does not contain conditionals. This is why it can be taken for granted in AGM that for a given belief set K and some sentence p that is compatible with all the factual sentences in K there is some operation that, when applied to K , gives rise to some belief set K' such that $K \cup \{p\} \subseteq K'$ (scilicet an expansion or a revision by a sentence consistent with K). This is plausible as long as K and K' are restricted to the purely factual fraction of the language. However, if K also contains all the conditional beliefs that the agent holds, then the fact that p does not contradict any factual belief in K does not guarantee that it does not contradict any of the conditional beliefs held in K . For instance, suppose that originally I know nothing about John's profession. It seems as if any concrete belief that I can acquire about his profession will lead to the loss of some conditional belief. If I learn that he is a driver by profession, then I will lose my belief that if he

¹¹ The notation in the quoted formulas has been slightly modified.

goes home from work by taxi every day, then he is a rich man. If I learn that he is a policeman, then I will lose my conditional belief that if he drove past several speed cameras at 110 mph last evening then he will be cited for speeding. If I learn that he is a philosopher, then I will lose my belief that if he has spent most of the last two years thinking intensely about the meaning of life, then he is unemployed and depressed. Generally speaking, it is difficult to find a clear example of a belief that can be acquired without the loss of some previously held conditional belief. This has far-reaching implications for the project of developing belief revision models capable of housing conditional sentences. It may for instance be necessary to give up the use of expansion as one of the (idealized) operations by which we try to capture the mechanisms of belief change. [11, 29] This is a problem that seems to affect AGM and DDL alike.

DDL has the major advantage, as compared to AGM, of allowing us to express self-referential beliefs. This applies not only to beliefs about one’s own current beliefs such as BBp or $B \neg B \neg p$ but also to the arguably much more interesting beliefs that refer to how one will change one’s beliefs under certain influences. However, the condition that $B(p \Box \rightarrow q)$ is true if and only if $[*p]Bq$ appears to be in a sense trivializing since it equates two entities between which we have an interesting tension: an agent’s conditional beliefs and her tendencies to change her beliefs.¹² One interesting further development of DDL would be to treat $B(p \Box \rightarrow q)$ and $[*p]Bq$ as separate entities with different truth conditions so that their truth values coincide sometimes but not always.

7 Conclusion

All belief change frameworks are the outcomes of far-reaching idealizations—otherwise they would be much too complex to work with. This applies to DDL as well as to the rival frameworks. In the above reflections on DDL I have focused on some of its limitations, but the remaining impression is that Krister Segerberg has provided us with an unusually versatile framework that is more suitable than most others for the introduction of new formal elements and new interpretations. Above, two such potential additions have been put forward: alternative interpretations in which the interior of $[]$ and $\langle \rangle$ represents the effects of external influences rather than the performance of actions, and models that have separate, non-equivalent representations of conditional beliefs and tendencies to change beliefs. Another development for which DDL is unusually well suited is the introduction of non-sentential inputs, in order to capture some of the properties of actual belief change that are lost in models operating with sentences. We can for instance have a set \mathcal{I} of non-sentential

¹² This conflation is perhaps stimulated by the standard theory of probabilities, in which two notions of degree of belief are merged: the current strength of a belief and the propensity to retain it when it is challenged are represented by the same number. However, this is a limitation that should not necessarily be transferred to other frameworks.

entities called inputs such that for any $\alpha \in \mathcal{I}$, we interpret $[\alpha]Bp$ as saying that after receiving the input α , the sentence p is believed.

References

1. Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50, 510–530.
2. Audi, R. (2001). Doxastic voluntarism and the ethics of belief. In M. Setup (Ed.), *Knowledge, truth, and duty. Essays on epistemic justification, responsibility, and virtue* (pp. 93–111). Oxford: Oxford University Press.
3. Aucher, G. (2003). *A combined system for update logic and belief revision*, ILLC report MoL-2003-03. Amsterdam: ILLC, University of Amsterdam.
4. Baltag, A., Moss, L., & Solecki, S. (1998). The logic of public announcements, common knowledge, and private suspicions. In I. Gilboa (Ed.), *Proceedings of the 7th conference on theoretical aspects of rationality and knowledge (TARK 98)* (pp. 43–56). San Francisco: Morgan Kaufmann.
5. Baltag, A., & Smets, S. (2008). A Qualitative Theory of dynamic interactive belief revision. In G. Bonanno, W. van der Hoek, M. Wooldridge (Eds.) *Logic and the Foundations of Game and Decision Theory (LOFT 7)*, Texts in Logic and Games 3 (pp. 11–58). Amsterdam: Amsterdam University Press.
6. Cantwell, J. (1997). On the logic of small changes in hypertheories. *Theoria*, 63, 54–89.
7. Fermé, E., & Hansson, S. O. (2011). AGM 25 years. Twenty-five years of research in belief change. *Journal of Philosophical Logic*, 40, 295–331.
8. Fuhrmann, A. (1991). Theory contraction through base contraction. *Journal of Philosophical Logic*, 20, 175–203.
9. Fuhrmann, A., & Hansson, S. O. (1994). A survey of multiple contraction. *Journal of Logic, Language, and Information*, 3, 39–76.
10. Gärdenfors, P. (1986). Belief revisions and the Ramsey test for conditionals. *Philosophical Review*, 95, 81–93.
11. Hansson, S. O. (1992). In defense of the Ramsey test. *Journal of Philosophy*, 89, 522–540.
12. Hansson, S. O. (2000). Formalization in philosophy. *Bulletin of Symbolic Logic*, 6, 162–175.
13. Hansson, S. O. (2009). Replacement—a Sheffer stroke for belief revision. *Journal of Philosophical Logic*, 38, 127–149.
14. Hansson, S. O. (2010). Methodological pluralism in philosophy. *Theoria*, 76, 189–191.
15. Hansson, S. O. (2011). Logic of belief revision. In *The Stanford Encyclopedia of Philosophy*. (<http://plato.stanford.edu/entries/logic-belief-revision>)
16. Hintikka, J. (1962). *Knowledge and belief: An introduction to the logic of the two notions*. Ithaca: Cornell University Press.
17. Leitgeb, H., & Segerberg, K. (2007). Dynamic doxastic logic: Why, how, and where to? *Synthese*, 155, 167–190.
18. Levi, I. (1977). Subjunctives, dispositions and chances. *Synthese*, 34, 423–455.
19. Levi, I. (1988). Iteration of conditionals and the Ramsey test. *Synthese*, 76, 49–81.
20. Levi, I. (1991). *The fixation of belief and its undoing*. Cambridge: Cambridge University Press.
21. Lindström, S., & Rabinowicz, W. (1991). Epistemic entrenchment with incomparabilities and relational belief revision. In A. Fuhrmann & M. Morreau (Eds.), *The logic of theory change* (pp. 208–228). Berlin: Springer.
22. Lindström, S., & Rabinowicz, W. (1999). DDL unlimited. Dynamic doxastic logic for introspective agents. *Erkenntnis*, 51, 353–385.
23. Montmarquet, J. (2008). Virtue and voluntarism. *Synthese*, 161, 393–402.
24. Mourad, R. (2008). Choosing to believe. *International Journal for Philosophy of Religion*, 63, 55–69.

25. Nickel, P. J. (2010). Voluntary belief on a reasonable basis. *Philosophy and Phenomenological Research*, 81, 312–334.
26. Nottelmann, N. (2006). The analogy argument for doxastic voluntarism. *Philosophical Studies*, 131, 559–582.
27. Plaza, J. (1989). Logics of public communications. In M. L. Emrich, M. S. Pfeifer, M. Hadzikadic, & Z. W. Ras (Eds.), *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems*, (pp. 201–216). Oak Ridge, Tennessee: Oak Ridge National Laboratory.
28. de Rijke, M. (1994). Meeting some neighbours. In J. van Eijck, & A. Visser (Eds.), *Logic and information flow*, (pp. 170–195). Cambridge: MIT Press.
29. Rott, H. (1989). Conditionals and theory change: Revisions expansions, and additions. *Synthese*, 81, 91–113.
30. Scott-Kakures, D. (1994). On belief and the captivity of the will. *Philosophy and Phenomenological Research*, 54, 77–103. [p. 83, to be checked.].
31. Segerberg, K. (1989). A note on an impossibility theorem of Gärdenfors. *Noûs*, 23, 351–354.
32. Segerberg, K. (1995). Belief revision from the point of view of doxastic logic. *Logic Journal of the IGPL*, 3, 535–553.
33. Segerberg, K. (1996). Three recipes for revision. *Theoria*, 62, 62–73.
34. Segerberg, K. (1997). Proposal for a theory of belief revision along the lines of Lindström and Rabinowicz. *Fundamenta Informaticae*, 32, 183–191.
35. Segerberg, K. (1999). Two traditions in the logic of belief: Bringing them together. In H. J. Ohlbach & U. Reyle (Eds.), *Logic, language and reasoning: Essays in honour of Dov Gabbay* (pp. 135–147). Dordrecht: Kluwer Academic Publishers.
36. Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory. American philosophical quarterly monograph series* (Vol. 2). Oxford: Blackwell.
37. van Benthem, J. (1989). Semantic parallels in natural language and computation. In H.-D. Ebbinghaus, J. Fernandez-Prida, M. Garrido, D. Lascar, M. Rodrigues Artalejo *Logic Colloquium '87* (pp. 331–375). Amsterdam: North-Holland.
38. van Benthem, J. (1995). Logic and the flow of information. In *Proceedings of the 9th International Congress of Logic, Methodology and Philosophy of Science. Studies in Logic and the Foundations of Mathematics* (Vol. 134, pp. 693–724).
39. van Ditmarsch, H. (2005). Prolegomena to dynamic logic for belief revision, *Synthese (Knowledge, Rationality & Action)*, 147, 229–275.
40. van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2007). *Dynamic epistemic logic. Synthese Library*. Dordrecht: Springer.
41. Wansing, H. (2000). A reduction of doxastic logic to action logic. *Erkenntnis*, 53, 267–283.
42. Wansing, H. (2006). Doxastic decisions, epistemic justification, and the logic of agency. *Philosophical Studies*, 128, 201–227.