Actions, Belief Update, and DDL

Jérôme Lang

Abstract Two prominent topics in Krister Segerberg's works are, on the one hand, actions, and on the other hand, belief change. Both topics are connected in multiple ways; one of these connections is via KGM belief update, since, as we argue, belief update is a specific case of feedback-free action progression. We discuss the links between update and action, and, starting from Segerberg's works, discuss further other possible interpretations of belief update, its differences with AGM belief revision, and why it is interesting to develop further KGM-based Dynamic Doxastic Logic.

1 Introduction

Krister Segerberg has introduced and developed a powerful and influential way of dealing with belief change: dynamic doxastic logic (DDL). DDL aims at expressing belief change actions at the same language level as factual sentences, using dynamic modalities $[\star\varphi]$, where $\star\varphi$ is the action of adding φ to the agent's belief. Nesting such belief change modalities allows us to reason about an agent's beliefs about how her beliefs are changed. For instance, borrowing from [25], p. 169, $\mathbf{B}[\star\varphi]B\theta$ expresses that the agent believes that after adding φ to her body of knowledge she will believe θ, and $[\star[\star\varphi]\mathbf{B}\theta]\mathbf{B}\chi$ expresses that the agent believes χ after adding to her belief state the information that adding φ to it would lead to a belief in θ.

In the paragraph above I deliberately avoided using the "to revise", and used the more neutral, but less elegant verbs "to change" or "to add". However, most of the work on DDL assumes that the belief change operation $\star$ corresponds to a *belief revision*, in the sense of Alchourrón, Gärdenfors and Makinson [1]; see for instance

A significant part of this article is a revised version of [23].

J. Lang (✉)
LAMSADE-CNRS, Paris, France
e-mail: Jerome.Lang@irit.fr

R. Trypuz (ed.), *Krister Segerberg on Logic of Actions*,
Outstanding Contributions to Logic 1, DOI: 10.1007/978-94-007-7046-1_11,

[33, 34]. Other parts of this special issue deal with DDL and its relationship with AGM-style belief revision (as well as its iterated versions), and the rôle played by the Ramsey rule and Gärdenfors' impossibility theorem in the development of DDL. Segerberg however noticed that moving belief change actions from the linguistic meta-level to the object level makes also perfectly sense for other paradigms of belief change, be it other operations in AGM-style belief change such as expansion and contraction, and also other non-AGM notions of belief change, the most prominent example being *belief update*, in the sense of Katsuno and Mendelzon [21] and Grahne [13]—which Segerberg calls the KGM paradigm. Developing a KGM version of DDL and highlighting its main differences with the traditional, AGM version of DDL is mentioned first in Lindström and Segerberg [28] and developed further by Leitgeb and Segerberg [25] of which it is one of the main topics.

Now, although many chapters about belief update have been written, including many chapters addressing its differences with belief revision, its precise scope still remains unclear. Part of the reason is that the first generation of chapters on belief update contain a number of vague and ambiguous formulations, such as "belief revision has to do with static worlds, while belief update has to do with dynamic worlds", or "belief update incorporates into a belief base some notification of a change in the world".

Friedman and Halpern [11] were perhaps the first to argue that this is not as simple as that. The issue is also addressed by Leitgeb and Segerberg [25], pages 183 and 184:

> In the literature of belief change the distinction between static and dynamic environments has become important. (...) it seems right to say that that belief change due to new information in an unchanging environment has come to be called *belief revision* (the static case, in the sense that the "world" remains unchanged), while it is fairly accepted to use the term *belief update* for belief change that is due to reported changes in the environment itself (the dynamic case, in the sense that the "world" changes).(...) The established tradition notwithstanding, it would be interesting to see a really convincing argument for tying AGM revision to static environments. (...) But it is also not clear that belief update has to be interpreted as reflecting a proper change in the environment.

Leitgeb and Segerberg also address an important ramification of this major question, which has to do with the role and the meaning of rankings of worlds in revision and in update. They give a very convincing line of argumentation towards the following conclusion: in revision, rankings are subjective and correspond to relative plausibilities (they can be thought of as an ordinal counterpart of subjective probabilities). In belief update, rankings are objective (agent-independent) and correspond to similarity between worlds. Let me quote Leitgeb and Segerberg [25], pp. 184–185:

> (...) Given new evidence, we find that in the case of belief revision the agent tries to change his beliefs in a way such that the worlds that he subsequently believes to be in comprise the *subjectively most plausible deviation* from the worlds he originally believed to inhabit. However, when confronted in the same evidence in belief update, the agent tries to change his beliefs in a way such that the worlds that he subsequently believes to be in are *as objectively similar as possible* to the worlds he originally believed to be the most plausible candidates for being the actual world.

This question about the role of rankings can be pushed even further, as we may even question the need for rankings in belief update. Accordingly, a series of chapters defined and studied families of update operators that, in contrast to the original model by Grahne, Katsuno and Mendelzon, are not based on minimization and thus do not need any rankings at all. This is extensively discussed by Herzig and Rifi [18]. This is in sharp contrast with belief revision, and this may be part of the explanation why the Ramsey test, to which AGM revision does not escape, seems perfectly escapable with belief update. This question of the compatibility of KGM update with the Ramsey test is addressed in detail by Leitgeb and Segerberg, pp. 179–187. It is further linked to the question of iteration, which appear to be much less problematic in belief update in with belief revision.

This chapter addresses all of these questions and develops on them (several more than others; in particular, there will be no emphasize at all on the Ramsey test), and discusses in detail some of the answers given in [25]. It is partly based on a previous conference chapter of mine [23]. The main question of this chapter is the identification of the precise scope of belief update, i.e., the conditions (expressed by properties of the world and of the agent's beliefs) under which update is a suitable process for belief change. After recalling some background on KGM belief update in Sect. 2, we give in Sect. 3 an informal discussion about the role of time in revision and update. In Sect. 4, we relate update to the field of reasoning about action (another issue in which Krister Segerberg is a major contributor). Our main claim is that updating a knowledge base by α corresponds to progressing it by a specific "purely physical", feedback-free action "make α true" whose precise meaning depends on the chosen update operator. This in turn raises the following question, addressed in Sect. 5: if update is progression, are there belief change operators corresponding to regression? In Sect. 6 we discuss another important (and different?) interpretation of belief update, which has to do with counterfactuals and causality; we address the question of whether this interpretation is really different from action progression, or only a variation of it. In Sect. 7 we come back to where the chapter started, namely DDL, and show why it is highly promising to develop further an update-based version of DDL. Further issues are briefly addressed in Sect. 8.

2 Belief Update

Let L^V be the propositional language generated from a finite set of propositional variables V, the usual connectives and the Boolean constants $\top$, $\bot$. $S = 2^V$ is the set of *states* (i.e., propositional interpretations). For any $\varphi \in L^V$, $Mod(\varphi)$ is the set of states satisfying φ. For any $X \subseteq S$, $for(X)$ is the formula of L^V (unique up to logical equivalence) such that $Mod(for(X)) = X$. If $X = \{s\}$, we write $for(s)$ instead of $for(\{s\})$. We use $\varphi \oplus \psi$ as a shorthand for $\varphi \leftrightarrow \neg\psi$.

As in [21], a belief update operator $\diamond$ is as mapping from $L^V \times L^V$ to L^V, i.e., mapping two propositional formulas φ (the initial belief state) and α (the "input") to a propositional formula $\varphi \diamond \alpha$ (the resulting belief state). We recall here the Katsuno-Mendelzon (KM for short) postulates for belief update [21].

U1 $\varphi \diamond \alpha \models \alpha$.
U2 If $\varphi \models \alpha$ then $\varphi \diamond \alpha \equiv \varphi$.
U3 If φ and α are both satisfiable then $\varphi \diamond \alpha$ is satisfiable.
U4 If $\varphi \equiv \psi$ and $\alpha \equiv \beta$ then $\varphi \diamond \alpha \equiv \psi \diamond \beta$.
U5 $(\varphi \diamond \alpha) \wedge \beta \models \varphi \diamond (\alpha \wedge \beta)$.
U6 If $\varphi \diamond \alpha \models \beta$ and $\varphi \diamond \beta \models \alpha$ then $\varphi \diamond \alpha \equiv \varphi \diamond \beta$.
U7 If φ is complete then $(\varphi \diamond \alpha) \wedge (\varphi \diamond \beta) \models \varphi \diamond (\alpha \vee \beta)$.
U8 $(\varphi \vee \psi) \diamond \alpha \equiv (\varphi \diamond \alpha) \vee (\psi \diamond \alpha)$.

Although we have recalled all postulates for the sake of completeness, we should not accept them unconditionally. They have been discussed in several chapters, including [18] in which it was argued that not all these postulates should be required, and that the "uncontroversial" ones (those deeply entrenched in the very notion of update and satisfied by most operators studied in the literature) are (U1), (U3), (U8), and (U4) to a lesser extent. We therefore call a *basic update operator* any operator $\diamond$ from $L^V \times L^V$ to L^V satisfying at least (U1), (U3), (U4) and (U8). In addition, $\diamond$ is said to be *syntax-independent* if it also satisfies (U4), *inertial* if it also satisfies (U2), and $\diamond$ is a *KM update operator* if it satisfies (U1)–(U8).[1] In this chapter we refer to some specific update operators such as the PMA [36]; see [18] for a compendum of belief update operators that date, and [17] for an update on the literature about update since then.

The first goal of this chapter consists in identifying is the exact scope of belief revision and belief update, and more generally belief change operators. To assess the scope of belief change operators, we need to be able to talk about the properties of the system (the world and the available actions) and the properties of the agent's state of knowledge, as in the taxonomy for reasoning about action and change from [31]. However, unlike reasoning about action, belief change processes have never (as far as we know) been analyzed from the point of view of such a taxonomy. A first step is taken towards this direction (for belief revision only) in [11]. We aim at identifying further the precise scope of belief update, i.e., the conditions (expressed by properties of the world and of the agent's beliefs) under which update is a suitable process for belief change.

3 Time, Revision, and Update

As already quoted in the Introduction, Leitgeb and Segerberg write in [25], pp. 183 and 184:

[1] (U5), (U6) and (U7) are much more controversial than the other ones (see [18]); they characterize the specific class of updates based on a similarity-based minimization process (which is known to lead to several counterintuitive results).

> The established tradition notwithstanding, it would be interesting to see a really convincing argument for tying AGM revision to static environments. (...) But it is also not clear that belief update has to be interpreted as reflecting a proper change in the environment.

Their diagnosis is definitely right: the discourse, seen so often, that the difference between the scope of revision and that of update should be seen as an opposition between static and dynamic environments, is wrong indeed. Belief revision, AGM style, has been developed as a qualitative counterpart of probabilistic conditionalisation; tying AGM to "static environments" would thus implicitly mean that the probability calculus does not apply to dynamic environments—which would be absolutely nonsense. And indeed, nothing in the AGM theory of belief revision implies that we should restrict its application to static worlds. Belief revision [10] is meant to map a belief set (a closed logical theory, or equivalently, since the language is finitely generated, a propositional formula[2]) and a new piece of information α (a consistent propositional formula) whose truth is held for sure, into a new belief set $K * \alpha$ taking account of the new piece of information without rejecting too much of the previous beliefs. The initial belief set as well as the new piece of information may talk about the state of an evolving world at different time points. As remarked already by Friedman and Halpern [11], what is essential in belief revision is not that the world is static, but that *the language used to describe the world* is static. Thus, if an evolving world is represented using time-stamped propositional variables of the form v_t (v true at time t), we can perfectly revise a belief set by some new information about the past or the present (or even, sometimes, the future), and infer some new beliefs about the past, the present, or the future.

Example 3.1 On Monday, Alice is the head of the computer science lab while Bob is the head of the math lab. On Tuesday I learned that one of them resigned (but without knowing which one). On Wednesday I learn that Charles is now the head of the math lab, which implies that Bob isn't. (It is implicit that heads of labs tend to keep their position for quite a long time.) What do I believe now?

Example 3.1 contains a sequence of two "changes". Both are detected by observations, and the whole example can be expressed as a revision process (with time-stamped variables). Let us identify Monday, Tuesday and Wednesday by the time stamps 1, 2 and 3. On Monday I believe A_1, B_1, as well as the persistency laws $A_1 \leftrightarrow A_2$, $A_2 \leftrightarrow A_3$, $B_1 \leftrightarrow B_2$ etc., therefore I also believe A_2, B_2 etc.: I expect that Alice and Bob will remain the heads of their respective labs on Tuesday and Wednesday. The revision by $\neg A_2 \vee \neg B_2$ (provided that the revision operator minimizes change) leads me to believe A_1, B_1, $A_2 \oplus B_2$, $A_3 \oplus B_3$ etc.: on Tuesday, I still believe that Alice and Bob were heads of their labs on Monday, and that now *exactly one* of them is. Then the revision by $\neg B_3$ (at time 3) makes me believe A_1, B_1, A_2, $\neg B_2$, A_3, $\neg B_3$: on Wednesday, I understand that Bob was the one to resign on

[2] Our assumption that the language is finite allows us to consider revision operators as acting on propositional formulas as in [22] (instead of belief sets).

Tuesday, and therefore that Alice was still head of the CS lab on Tuesday, and is still now.[3]

Now, the fact that belief revision can deal with (some) evolving worlds suggests that the opposition between revision and update relies on the possibility or not that the state of the world may evolve is not accurate. In particular, claiming that belief update is the right belief change operation for dealing with evolving worlds is unsufficient and ambiguous. The literature on belief update abounds in ambiguous explanations such as "update consists in bringing the knowledge base up to date when the world is described by its changes".[4] Especially, the expressions "describing the world by its changes" and "notification of change", appearing in many chapters, are particularly ambiguous. The problem is not that much, as it has been observed sometimes, that in these expressions "change" has to be understood as "possibility of change" (we'll come back to this point). The main problem is the status of the input formula α. To make things clear, here is an example.

Example 3.2 My initial belief is that either Alice or Bob is in the office (but not both). Both tend to stay in the office when they are in. Now I see Bob going out of the office. What do I believe now?

Trying to use belief update to model this example is hopeless. For all common update operators seen in the literature, updating $A \oplus B$ by $\neg B$ leads to $\neg B$, and not to $\neg A \wedge \neg B$. The culprit is (U8), which, by requiring that all models of the initial belief set be updated separately, forbids us to infer new beliefs about the past from later observations. Indeed, because of (U8), we have $(A \oplus B) \diamond \neg B \equiv [(A \wedge \neg B) \diamond \neg B] \vee [(\neg A \wedge B) \diamond \neg B] \equiv (A \wedge \neg B) \vee (\neg A \wedge \neg B) \equiv \neg B$. The only way to have $\neg A \wedge \neg B$ as the result would be to have $(A \wedge \neg B) \diamond \neg B \equiv \neg A \wedge \neg B$, which can hold only if there is a causal relationship between A and B, such as B becoming false entails A becoming false—which is not the case here.

Example 3.2 definitely deals with an evolving world and contains a "notification of change", and still it cannot be formulated as a belief update process. On the other hand, like Example 3.1, it can be perfectly expressedas is a time-stamped belief revision process.[5]

The key point is (U8) which, by requiring that all models of the initial belief set be updated separately, *forbids us from inferring new beliefs about the past from later observations*: indeed, in Example 3.2, belief update provides no way of eliminating the world $(A, \neg B)$ from the set of *previously* possible worlds, which in turn, does not allow for eliminating $(A, \neg B)$ from the list of possible worlds *after* the update:

---

[3] Note that this scenario is also a case for belief extrapolation [8], which is a particular form of time-stamped revision.

[4] This formulation appears in [21], which may be one of the explanations for such a long-lasting ambiguity.

[5] Note that without time stamps (and in particular within the framework of belief update), we cannot distinguish between "B has become false" (as in "I see Bob go out of the office") and "the world has evolved in such a way that B is now false" (as in "I now see Bob out of his office"). Anyway, for Example 3.2, the expected outcome is the same in both cases (provided that A and B are expected to persist with respect to the granularity of time considered).

if $(A, \neg B)$ is a possible world at time t, then its update by $\neg B$ *must* be in the set of possible worlds at time $t + 1$. In other terms, update fails to infer that Alice wasn't in the office and still isn't.

Belief update fails as well on Example 3.1: updating $A \wedge B \wedge \neg C$ by $\neg A \vee \neg B$ gives the intended result, but only by chance (because the agent's initial belief state is complete). The second step fails: with most common update operators, updating $(A \oplus B) \wedge \neg C$ by $\neg B \wedge C$ leads to $\neg B \wedge C$, while we'd expect to believe A as well.

The diagnosis should now be clear: the input formula α is not a mere observation. An observation made at time $t + 1$ leads to filter out some possible states at time $t + 1$, which in turn leads to filter out some possible states at time t, because the state of the world at time t and the state of the world at time $t + 1$ are correlated (by persistence rules or other dynamic rules.[6]). And finally, the successor worlds (at time $t + 1$) of these worlds at time t that update failed to eliminate can not be eliminated either. Such a backward-forward reasoning needs a proper generalization of update (and of revision), unsurprisingly called *generalized update* [3].

One could try to argue that such scenarios (such as Example 3.1 or 3.2) are both a case for revision and update, depending whether the formulation of the problem uses time-stamped variables or not. This line of argumentation fails: expressing Example 3.2 as a belief update still leads to the counterintuitive results that we do not learn anything about Alice. Besides, several authors remarked that, unless belief bases are restricted to complete bases, a belief update operator cannot be a belief revision operator. For instance, it is shown in [15, 30] that the AGM postulates are inconsitent with U8 as soon as the language contains at least two propositional symbols.

4 Update as Action Progression

We now investigate in further detail the belief change interpretation of belief update. (There is at least one other interpretation, which deals with causality and counterfactuals, on which we shall come back in Sect. 6.) Since standard belief update precludes any possibility of feedback, the input formula α has to be understood as an *action effect*, and certainly not as an observation. If α has to be understood as an action effect, update is a particular form of *action progression* for *feedback-free actions*. Action progression (as considered in the literature of reasoning about action and logic-based planning) consists in determining the belief state obtained from an initial belief state after a given action is performed, this action corresponding to a transition graph (an automaton) between states of the world.

[6] The only case where belief update could be compatible with interpreting α as an observation would therefore be the case where not the faintest correlation exists between the state of the world at different time points; in this case, we would have $\varphi \diamond \alpha \equiv \alpha$ whenever α is consistent—a totally degenerate and uninteresting case.

This connection between belief update and action progression was first mentioned by Del Val and Shoham [5], who argued that updating an initial belief state φ by a formula α corresponds to one particular action; they formalize such actions in a formal theory of actions based on circumscription, and their framework for reasoning action is then used to derive a semantics for belief update. The relationship between update and action progression appears (more or less explicitly) in several other chapters, including [27], who expresses several belief update operators in a specific action language. Still, the relationship between update and action progression still needs to be investigated in more detail.

We first need to give some background on reasoning about action. Generally speaking, an action A has two types of effects: an *ontic (or physical) effect* and an *epistemic effect*. For instance, if the action consists in tossing a coin, its ontic effect is that the next value of the fluent *heads* may change, whereas its epistemic effect is that the new value of the fluent is observed (this distinction between ontic and epistemic effects is classical in most settings). Complex actions (with both kinds of effects) can be decomposed into two actions, one being ontic and feedback-free, the other one being a purely epistemic (sensing) action.

The simplest model for a purely ontic (i.e., feedback-free) action A consists of a *transition graph* R_A on S.[7] $R_A(s, s')$ means that s' is accessible from s after A. $R_A(s) = \{s' \mid R_A(s, s')\}$ is the set of states that can obtain after performing A in s. If $R_A(s)$ is a singleton for all s then A is *deterministic*. If $R_A(s) = \emptyset$ then A is *inexecutable* in s. A is *fully executable* iff $R_A(s) \neq \emptyset$ for every s.

An epistemic action e corresponds to a set of possible *observations*, plus a *feedback function* f_e from S to 2^O, where O is a finite *observation space*. $o \in f_e(s)$ means that observation o may be obtained as feedback when performing e in state s. Observations are of course correlated with states (for instance, an observation can be a propositional formula, or equivalently a set of states.) For the sake of simplicity, we identify O with L_V, that is, we consider that observations are propositional formulas (note however that this implies a loss of generality. The simplest possible epistemic actions are *truth tests*, and correspond to two possible observations, φ and $\neg\varphi$, for some propositional formula φ. An epistemic action e is *truthful* iff for all $s \in S, o \in O, o \in f_e(s)$ implies $s \models o$, *deterministic* iff for all $s \in S$, $f_e(s)$ is a singleton, and *fully executable* iff for all $s \in S$, $f_e(s) \neq \emptyset$.

Let A be a purely ontic action modelled by a transition graph R_A on S. For any formula $\varphi \in L^V$, the *progression of* φ *by* A is the propositional formula (unique up to logical equivalence) whose models are the states that can obtain after performing A in a state of $Mod(\varphi)$: $prog(\varphi, A)$ is defined by

$$prog(\varphi, A) = for\left(\bigcup_{s \models \varphi} R_A(s)\right) \tag{1}$$

[7] More sophisticated models may involve graded uncertainty such as probabilities, delayed effects etc.

Remark that the probabilistic variant of action progression is the well-known action progression operator for stochastic actions: let p is a probability distribution over S and A a stochastic action described by a stochastic matrix $p(.|., A)$, where $p(s'|s, A)$ is the probability of obtaining s' after performing A in s. Then $prog_P(p, A)$ is the probability distribution over S defined by

$$prog_P(p, A)(s') = \sum_{s \in S} p(s)p(s'|s, A)$$

Mapping each probability distribution p into the belief state $B(p) = for(\{s|p(s) > 0\})$ consisting of those states deemed possible by p, i.e., $B(prog_P(p, A)) = prog(B(p), A)$. As argued by Dubois and Prade [7], the probabilistic variant of belief update is Lewis' imaging [26]: $p(.|., \alpha)$ is then defined by

$$p(s'|s, \alpha) = \begin{cases} \frac{1}{|Proj(s,\alpha)|} & \text{if } s' \in proj(s, \alpha) \\ 0 & \text{otherwise} \end{cases}$$

where $proj(s, \alpha)$ is the set of states closest to α (according to some proximity structure).

Lastly, for any action A, $Inv(A)$ is the set of *invariant states* for A, i.e. the set of all states s such that $R_A(s) = \{s\}$.

Clearly enough, (1) is identical to (U8). Therefore, for any update operator (and more generally any operator satisfying (U8)) and any input formula α, *updating by α is an action progression operator*. This raises several questions: (a) Which action is this exactly? (b) What is the class of actions that correspond to updates? (c) If update is progression, are there belief change operators corresponding to regression?

Question (a) first. As argued above, (U8) and (1) mean that the action is feedback-free. Indeed, a feedback would allow us to eliminate some states after the action has been performed, which in turn would lead us to eliminate some states before the action took place (see [3, 8]).[8] This comes down to saying that belief update assumes *unobservability*: the set of possible states after A is performed is totally determined by the set of possible states before it is performed and the transition system corresponding to A. In other words, *what you foresee is what you get* (WYFIWYG): once we have decided to perform A, waiting until it has actually been performed will not bring us any new information. Expressed in a modal language, the WYFIWIG principle is nothing but the (RR) axiom of Grahne [13], of which we give Leitgeb and Segerberg's formulation ([25], p. 181):

$$\mathbf{B}(\varphi \mathbin{\Box\!\!\rightarrow} \psi) \leftrightarrow [\diamond\varphi]\mathbf{B}\psi$$

[8] Unless the state of the world after the action is performed is totally disconnected from the state of the world before the action is performed, which only happens if $R_A(s) = S$ for all s. In this case, a feedback never allows for learning anything about the past state of the world. Clearly, this case is a very degenerated one.

(RR) can be seen the syntactical counterpart of (U8). Leitgeb and Segerberg consider it as the key axiom of KGM, and I do agree.

Note that using update in Example 3.2 would correspond to performing an action whose effect is to make Bob go out of his office (when he is initially not in the office, this action has no effect). Likewise, in Example 3.1, updating $A \oplus B \wedge \neg C$ by $\neg B \wedge C$ corresponds to performing the action "demote Bob from his position and appoint Charles instead".

Therefore, updating by α is a purely ontic (feedback-free) action. Can we now describe this action in more detail? (U1) means that the action of updating by α has to be understood as "make α true". Such actions (or events[9] have been given some attention for long by Segerberg, and are referred to in [25] (pp. 182–183) as

> "resultative" events': events describable in terms of their results (...). The intended meaning of a term $\delta\varphi$ would be "the event resulting in (its being the case that) φ". Accordingly, the intended meaning of a formula $[\delta\varphi]\psi$ would be "after the event resulting in (its being the case that) φ, it is the case that ψ, or more briefly, "after φ has just been realized, ψ."

More precisely, Segerberg studied in [32] a class of actions *bringing about that* α, or simply, *doing* α. In the light of the discussion above, comparing this class of actions *do* α and KGM belief update appears is more than worth doing. One of the main axioms for *do* α *is* $[do\ \alpha]\alpha$, which is obviously equivalent to (U1), modulo reformulation. Axioms (E1) and (E2) ([32], p. 333) are together equivalent to (U4). Where the two frameworks depart is with the last main axiom of *do* α, namely,

$$[do\ \alpha]\beta \rightarrow ([do\ \beta]\gamma \rightarrow [do\ \alpha]\gamma)$$

whose reformulation in the language of belief update is

$$\varphi \diamond \alpha \models \beta \rightarrow ((\varphi \diamond \beta \models \gamma) \rightarrow (\varphi \diamond \alpha \models \gamma))$$

This axiom (which, incidentally, implies the KM axiom (U6)), cannot be satisfied by a belief update operator satisfying (U1) and (U2). Indeed, take $\gamma = \varphi$, $\alpha = \neg\varphi$, and $\beta = \top$. Trivially, $\varphi \diamond \alpha \models \beta$ holds. Due to (U2), we have $\varphi \diamond \beta \equiv \varphi$, thus $\varphi \diamond \beta \models \gamma$ holds. Lastly, due to (U1), $\varphi \diamond \alpha \models \alpha$, which implies that $\varphi \diamond \alpha \models \gamma$ cannot hold. This fact is intriguing, as the axiom seems natural. I leave a deeper discussion for further research, but still, I am convinced that early works by Segerberg on *do* α actions (which appeared several years before the first chapters on belief update)—was very close to belief update, and, probably due to the fact that both streams of work were developed in different communities, very few works mention that.

[9] The distinction between actions and events is mostly irrelevant to our discussion. Actions are usually thought of as agent-trigerred, whereas events don't, or don't necessarily (see for instance [31]). Who triggers what has no impact on our discussion: an action performed consciously and intentionally by an agent, or a nature-trigerred event, or an action performed by another agent, have the same effects on the agent's belief state *provided* that, in all cases, the agent is perfectly aware of the action or the event taking place.

Back to interpreting "updating by α" as "make α true". More precisely, due to the absence of feedback reflected by (U8), updating φ by α could be understood as a dialogue between an agent and a robot: "All I know about the state of the world is that is satisfies φ. Please, go to the real world, see its real state, and whatever this state, act so as to change it into a world satisfying α, following some rules" (given that the robot does not communicate with the agent once it is the real world.) The rules to be followed by the robot are dictated by the choice of the update operator $\diamond$. If $\diamond$ satisfies (U2), then the rules state that if the α is already true then the robot must leave the world as it is. If $\diamond$ is the PMA [36], then the rules are "make α true, without changing more variables than necessary". More generally, when $\diamond$ is a Katsuno-Mendelzon operator, associated with a collection of similarity preorders (one for each world), the robot should make α true by changing s into one of the states that are most similar to it notion (s being closer to s_1 than to s_2 may, in practice, reflect that from s it is easier to go to s_1 than to s_2) and not as an epistemic notion of similarity, as it would be the case for belief revision. When $\diamond$ is a forgetting-based operation, such as WSS [14, 36] or the MPMA [6], then the rules are "make α true, without changing the truth values of a given set of variables (those that do not appear in α, or those that play no role in α)." And so on.

It is the right place to discuss the rôle of minimisation in belief update. It has been remarked already by several authors (see [18] for a synthetic discussion) that requiring minimisation of change is not always the right thing to do, and that many well-behaved update operators do not need it, nor do they need these KM faithful orderings around worlds—which strongly departs with AGM belief revision. These rankings are optional; when relevant, they correspond to *objective similarity* between worlds. Peppas et al. [30], argue that this similarity has be understood as *ontological*, which agrees with our view of $update(\diamond, \alpha)$ as an ontic action. Leitgeb and Segerberg go further in this direction by giving this illuminating argument ([25], pp. 184–185):

> We think that the actual difference between the intended interpretation of revision and update is given by the fact that the former belief change follows a *doxastic* order of "fallback positions" [29] while the latter conforms to a *worldly* similarity order of states of affairs—the one rides on a subjective structure, the other as an objective one. (...) Thus, given new evidence, we find that in the case of belief revision the agent tries to change his beliefs in a way such that he subsequently believes to be in the *subjectively most plausible deviation* from the worlds he originally believed to inhabit. However, confronted with the same evidence in belief update, the agent tries to change his beliefs in a way such that the worlds that he subsequently believes to be are *as objectively similar as possible* to the worlds he originally believed to be the most plausible candidates to be the actual world.

Writing things more formally: given an update operator $\diamond$ and a formula α, let $update(\diamond, \alpha)$ be the ontic action whose transition graph is defined by: for all $s, s' \in S$,

$$s' \in R_{update(\diamond,\alpha)}(s) \quad \text{iff} \quad s' \models for(s) \diamond \alpha$$

The following characterizations are almost straightforward, but worth mentioning, as they shed some light on the very meaning of the KM axioms.

Proposition 4.1 *Let $\diamond$ satisfy (U8).*

1. $\varphi \diamond \alpha \equiv prog(\varphi, update(\diamond, \alpha))$;
2. $\diamond$ *satisfies (U1) if and only if for any formula* $\alpha \in L^V$ *and any* $s \in S$, $R_{update(\diamond,\alpha)}(s) \subseteq Mod(\alpha)$;
3. $\diamond$ *satisfies (U2) if and only if for any formula* $\alpha \in L^V$, $Inv(update(\diamond, \alpha)) \supseteq Mod(\alpha)$;
4. $\diamond$ *satisfies (U3) if and only if for any satisfiable formula* $\alpha \in L^V$, $update(\diamond, \alpha)$ *is fully executable.*

Proof For point 1, (U8) implies that $Mod(\varphi \diamond \alpha) = \cup_{s \models \varphi} for(s) \diamond \alpha$, which, by definition of $update(\diamond, \alpha)$, is equal to $\cup_{s \models \varphi} R_{update(\diamond,\alpha)}(s)$, which, by definition of progression, is equal to $Mod(prog(\varphi, update(\diamond, \alpha)))$.

For point 2, let $\diamond$ satisfying (U1). Then $R_{update(\diamond,\alpha)}(s) = Mod(for(s) \diamond \alpha) \subseteq Mod(\alpha)$. Conversely, if for any α and any $s \in S$, $R_{update(\diamond,\alpha)}(s) \subseteq Mod(\alpha)$ holds, then $Mod(\varphi \diamond \alpha) = \cup_{s \models \varphi} for(s) \diamond \alpha = \cup_{s \models \varphi} R_{update(\diamond,\alpha)}(s) \subseteq Mod(\alpha)$, therefore $\varphi \diamond \alpha \models \alpha$.

For point 3, we have that for all s and α, $for(s) \diamond \alpha = for(s)$ if and only if $R_{update(\diamond,\alpha)}(s) = \{s\}$ if and only if $s \in Inv(update(\diamond, \alpha))$. Now, if $\diamond$ satisfies (U2) then for any α and $s \in Mod(\alpha)$, by (U2) we get $for(s) \diamond \alpha = for(s)$, therefore $s \in Inv(update(\diamond, \alpha))$. Conversely, if $Inv(update(\diamond, \alpha)) \supseteq Mod(\alpha)$ holds then for any φ such that $\varphi \models \alpha$ we have $Mod(\varphi \diamond \alpha) = \cup_{s \models \varphi} R_{update(\diamond,\alpha)}(s) = \cup_{s \models \varphi} s$ (because $for(s) \models \alpha$), therefore $Mod(\varphi \diamond \alpha) = Mod(\varphi)$, hence (U2) is satisfied.

For point 4, let α be a satisfiable formula. For any s, $for(s) \diamond \alpha$ is satisfiable if and only if $R_{update(\diamond,\alpha)}(s) \neq \emptyset$. If $\diamond$ satisfies (U3) then because $for(s)$ is satisfiable, $for(s) \diamond \alpha$ is satisfiable, therefore $R_{update(\diamond,\alpha)}(s) \neq \emptyset$; this being true for all s, $update(\diamond, \alpha)$ is fully executable. Conversely, assume $update(\diamond, \alpha)$ is fully executable, then for any satisfiable φ, $Mod(\varphi \diamond \alpha) = \cup_{s \models \alpha} R_{update(\diamond,\alpha)}(s) \neq \emptyset$; hence $\diamond$ satisfies (U3). □

From point 4 of Proposition 4.1, (U3) corresponds to full executability of $update(\diamond, \alpha)$. We may wonder what new properties of $update(\diamond, \alpha)$ obtain when other postulates are required. (U2) is particularly interesting in this respect. Indeed, the inertia postulate (U2) together with (U1) and (U8), reinterpreted in terms of action progression, means that any state that can be reached by $update(\diamond, \alpha)$ is an invariant state. More precisely:

Proposition 4.2 *Let $\diamond$ satisfying (U1), (U2) and (U8). Then*

$$R_{update(\diamond,\alpha)}(S) = Inv(update(\diamond, \alpha)) \cap Mod(\alpha)$$

Proof By (U1), $update(\diamond, \alpha)$ maps any state to a set of states satisfying α; then by (U2), any of these states is invariant by $update(\diamond, \alpha)$; therefore, $R_{update(\diamond,\alpha)}(S) \subseteq Inv(update(\diamond, \alpha))$. $R_{update(\diamond,\alpha)}(S) \subseteq Mod(\alpha)$ is a direct consequence of (U1). Finally, let $s \in Inv(update(\diamond, \alpha)) \cap Mod(\alpha)$. Then, by (U2), $for(s) \diamond \alpha =$

$for(s)$, hence $R_{update(\diamond,\alpha)}(s) = \{s\}$ and thus $s \in R_{update(\diamond,\alpha)}(S)$, which proves $Inv(update(\diamond, \alpha)) \cap Mod(\alpha) \subseteq Inv(update(\diamond, \alpha)) \cap Mod(\alpha)$. □

Note that if $R_{update(\diamond,\alpha)}(s) \subseteq Inv(update(\diamond, \alpha))$ for all s, then $update(\diamond, \alpha)$ is involutive, i.e., $R_{update(\diamond,\alpha)} \circ R_{update(\diamond,\alpha)} = R_{update(\diamond,\alpha)}$, but the converse fails to hold.

The other postulates do not have any direct effect on the properties of $update(\diamond, \alpha)$ considered *as an isolated action*, but they relate different actions of the form $update(\diamond, \alpha)$. Noticeably, requiring (U4) corresponds to the equality between $update(\diamond, \alpha)$ and $update(\diamond, \beta)$ when α and β are logically equivalent. The characterizations of (U5), (U6) and (U7) in terms of reasoning about action are purely technical and do not present any particular interest.

Let us now consider question (b). Obviously, given a *fixed* update operator $\diamond$ satisfying (U1), (U3), (U4) and (U8), some fully executable actions are not of the form $update(\diamond, \alpha)$. This is obvious because there are 2^{2^n} actions of the form $update(\diamond, \alpha)$ and 2^{n+2^n-1} fully executable actions, where $n = |V|$. Here is another proof, more intuitive and constructive: let $V = \{p\}$, thus $S = \{p, \neg p\}$, and consider the actions $A = switch(p)$, such that $R_A(p) = \{\neg p\}$ and $R_A(\neg p) = \{p\}$. Assume there is a formula α such that $A = update(\diamond, \alpha)$; then $U1$ enforces $\alpha \equiv \top$; therefore, if $A = update(\diamond, \alpha)$ then by (U4), $A = update(A, \top)$. Now, let A' be the identity action; we also have that if A' can be expressed as an update action for $\diamond$, then $A' = update(\diamond, \top)$. Therefore, at most one of A and A' can be expressed as an update action for $\diamond$.

Now, what happens if we allow $\diamond$ to vary? The question now is, what are the actions that can be expressed as $update(\diamond, \alpha)$, for *some* update operator $\diamond$ and some α?

Proposition 4.3 *Let A be a fully executable ontic action such that $R_A(s) \subseteq Inv(A)$ for all $s \in S$. Then there exists a KM-update operator, and a formula α, such that $A = update(\diamond, \alpha)$.*

Proof The proof is constructive. Let us take any formula $\alpha = for(Inv(A))$, and the collection of faithful orderings in the sense of [21] defined by $s_1 <_s s_2$ if and only if $s = s_1 \neq s_2$ or $(s \neq s_1, s \neq s_2, s_1 \in R_A(s), s_2 \notin R_A(s))$; and $s_1 \leq_s s_2$ iff not $(s_2 <_s s_1)$.

Because A is fully executable, $R_A(s) \neq \emptyset$ for any s, therefore $Inv(A) \neq \emptyset$ and α is satisfiable.

Let $s \models \alpha$. Because $\alpha = for(Inv(A))$ we have $R_A(s) = \{s\}$. By (U2), because $for(s) \models \alpha$, we have $for(s) \diamond \alpha = for(s)$, therefore $R_{update(\diamond,\alpha)}(s) = \{s\} = R_A(s)$.

Let $s \models \neg\alpha$. Then $s \notin R_A(s)$, which implies that $Min(\leq_s, Mod(\alpha)) = R_A(s)$, from which we have $for(s) \diamond \alpha = for(R_A(s))$ and $R_{update(\diamond,\alpha)}(s) = R_A(s)$.

We have established that $R_{update(\diamond,\alpha)}(s) = R_A(s)$ holds for all $s \in S$. Because of (U8), $\diamond$ is fully determined by $\{R_{update(\diamond,\alpha)}(s), s \in S\}$, therefore $A = update(\diamond, \alpha)$.

From Propositions 4.1 and 4.3 we get □

Corollary 4.4 *Let A be an ontic action. There exists a KM-update operator $\diamond$, and a formula α such that $A = update(\diamond, \alpha)$, if and only if A is fully executable and $R_A(s) \subseteq Inv(A)$ for all $s \in S$.*

A variant of Proposition 4.3 (and Corollary 4.4) can be obtained by not requiring $R_A(s) \subseteq Inv(A)$: in that case there exists an update operator $\diamond$ satisfying all the KM postulates except (U3), and a formula α such that $A = update(\diamond, \alpha)$. α can be taken as $\top$ and $s \leq_s s_2$ iff $s_1 \in R_A(s)$ or $s_2 \notin R_A(s)$.

Note that if (U2) is not required in Proposition 4.3 then we have the meaningless result that any action is expressible as an update.

5 Reverse Update

Now, question (c). Is there a natural notion which is to action regression what update is to progression? The point is that we do not have one, but two notions of action regression. The *weak* (or *deductive*) *regression* (also called *weak preimage* in the AI planning literature) of ψ by A is the formula whose models are the states from which the execution of A *possibly* leads to a model of ψ, while the *strong* (or *abductive*) *regression* (also called *strong preimage*) of ψ by A is the formula whose models are the states from which the execution of A *certainly* leads to a model of ψ:

$$reg(\psi, A) = form\,(\{s, R_A(s) \cap Mod(\psi) \neq \emptyset\})$$
$$Reg(\psi, A) = form\,(\{s, R_A(s) \subseteq Mod(\psi)\})$$

While weak regression is the suitable operator for *postdiction* (given that ψ now holds and that α has been performed, what can we say about the past state of the world?), strong regression is better understood as *goal regression* (what are the states in which it is guaranteed that performing α will lead to a goal state, i.e. a state satisfying ψ?) See for instance [24] for the interpretation of these two notions of regression in reasoning about action. This naturally leads to two notions of reverse update.

Definition 5.1 Let $\diamond$ be an update operator.

- the *weak reverse update* $\odot$ associated with $\diamond$ is defined by: for all $\psi, \alpha \in L^V$, for all $s \in S$,

$$s \models \psi \odot \alpha \text{ iff } for(s) \diamond \alpha \not\models \neg\psi$$

- the *strong reverse update* $\otimes$ associated with $\diamond$ is defined by: for all $\psi, \alpha \in L^V$, for all $s \in S$,

$$s \models \psi \otimes \alpha \text{ iff } for(s) \diamond \alpha \models \psi$$

Equivalently, $\psi \odot \alpha = for(\{s \mid for(s) \diamond \alpha \not\models \neg\psi\})$ and $\psi \otimes \alpha = for(\{s \mid for(s) \diamond \alpha \models \psi\})$.

Intuitively, weak reverse update corresponds to (deductive) postdiction: given that the action "make α true" has been performed and that we now know that ψ holds, what we can say about the state of the world before the update was performed is that it satisfied $\psi \odot \alpha$. As to strong reverse update, it is an abductive form of postdiction, better interpreted as goal regression: given that a rational agent has a goal ψ, the states of the world in which performing the action "make α true" is guaranteed to lead to a goal states are those satisfying $\psi \otimes \alpha$.

The following result shows that $\odot$ and $\otimes$ can be characterized in terms of $\diamond$:

Proposition 5.2 *1.* $\psi \odot \alpha \models \varphi$ *iff* $\neg\varphi \diamond \alpha \models \neg\psi$*;*
2. $\varphi \models \psi \otimes \alpha$ *iff* $\varphi \diamond \alpha \models \psi$*;*

Proof For point 1, assume $\neg\varphi \diamond \alpha \not\models \neg\psi$. Then there exists s and s' such that $s \models \neg\varphi$, $s' \in R_A(s)$ and $s' \models \psi$. This implies that $for(s) \diamond \alpha \not\models \neg\psi$, *i.e.*, $s \models \psi \odot \alpha$, and since $s \models \neg\varphi$, we have $\psi \odot \alpha \not\models \neg\varphi$. Conversely, assume $\psi \odot \alpha \not\models \varphi$. Then there exists $s' \models \psi$ and $s \models \neg\varphi$ such that $s' \in R_A(s)$, which implies that $\neg\varphi \not\models \neg\psi$. For point 2, assume $\varphi \diamond \alpha \not\models \psi$. Then there exists s' such that $s' \models \varphi \diamond \alpha$, and $s' \models \neg\psi$. This implies that there exists an s such that $s' \in R_A(s)$ and $s \models \varphi$, hence $for(s) \diamond \alpha \not\models \psi$, *i.e.*, $s \not\models \psi \otimes \alpha$. Conversely, assume $\varphi \not\models \psi \otimes \alpha$. Then there exists $s \models \varphi$ such that $s \not\models \psi \otimes \alpha$, *i.e.*, $for(s) \diamond \alpha \not\models \psi$, which implies that there is a s' such that $s' \in R_A(s)$ and $s' \models \neg\psi$, therefore $\varphi \diamond \alpha \not\models \psi$. □

As a consequence of Proposition 5.2, $\psi \odot \alpha$ is the *weakest formula* φ *such that* $\neg\varphi \diamond \alpha \models \neg\psi$*, and* $\psi \otimes \alpha$ *is the strongest formula* φ *such that* $\varphi \diamond \alpha \models \psi$*.*

Example 5.3 Let $\diamond = \diamond_{PMA}$ [36]. Let b and m stand for "the book is on the floor" and "the magazine is on the floor". The action $update(\diamond, b \vee m)$ can be described in linguistic terms by "make sure that the book or the magazine is on the floor". Then $b \odot (b \vee m) \equiv b \vee (\neg b \wedge \neg m) \equiv b \vee \neg m$, which can be interpreted as follows: if we know that the book is on the floor after $update(\diamond, b \vee m)$ has been performed, then what we can say about the previous state of the world is that either the book was already on the floor (in which case nothing changed) or that neither the book nor the magazine was on the floor (and then the update has resulted in the book being on the floor). On the other hand, $b \otimes (b \vee m) \equiv b$: if our goal is to have the book on the floor, the necessary and sufficient condition for the action $update(\diamond, b \vee m)$ to be guaranteed to succeed is that the book is already on the floor (if neither of them is, the update might well leave the book where it is and move the magazine onto the floor).

An interesting question is whether weak and strong reverse update can be characterized by some properties (which then would play the role that the basic postulates play for "forward" update). Here is the answer (recall that a basic update operator satisfies U1, U3, U4 and U8).

Proposition 5.4 $\odot$ *is the weak reverse update associated with a basic update operator* $\diamond$ *if and only if* $\odot$ *satisfies the following properties:*

W1 $\neg\alpha \odot \alpha \equiv \bot$*;*

W3 *if α is satisfiable then $\top \odot \alpha \equiv \top$;*
W4 *if $\psi \equiv \psi'$ and $\alpha \equiv \alpha'$ then $\psi \odot \alpha \equiv \psi' \odot \alpha'$;*
W8 $(\psi \vee \psi') \odot \alpha \equiv (\psi \odot \alpha) \vee (\psi' \odot \alpha)$.

In addition to this, $\diamond$ satisfies (U2) if and only if $\odot$ satisfies

W2 $(\psi \odot \alpha) \wedge \alpha \equiv \psi \wedge \alpha$.

Proof Note first that (W4) and (W8) are *exactly* the same properties as (U4) and (U8), replacing $\diamond$ by $\odot$.

Let $\odot$ be the weak reverse update associated with a basic update operator $\diamond$. Let us show that $\odot$ satisfies (W1), (W3), (W4) and (W8).

From Proposition 5.2, $\neg\alpha \odot \alpha \equiv \bot$ is equivalent to $\top \diamond \alpha \models \alpha$, *i.e.*, for all s, $for(s) \diamond \alpha \models \alpha$, which in turns is equivalent to: for all φ, $\varphi \diamond \alpha \models \alpha$, which is (U1). Therefore, $\odot$ satisfies (W1).

Let α be a satisfiable formula. Assume that $\odot$ does not satisfy (W3), that is, $\top \odot \alpha \not\equiv \top$: from (U8), there is a s such that $s \not\models \top \odot \alpha$, which is equivalent to $\top \odot \alpha \models for(S \setminus \{s\})$, i.e., using Proposition 5.2, $\neg for(S \setminus \{s\}) \diamond \alpha \models \bot$, equivalent to $for(s) \diamond \alpha$ unsatisfiable, which contradicts the assumption that $\diamond$ satisfies (U3). Therefore, $\odot$ satisfies (W3).

Assume $\psi \equiv \psi'$ and $\alpha \equiv \alpha'$. For any s, $s \models \psi \odot \alpha$ holds if only if $for(s) \diamond \alpha \not\models \neg\psi$, which using (U4) is equivalent to $for(s) \diamond \alpha' \not\models \neg\psi'$, therefore $s \models \psi' \odot \alpha'$, which implies that $\odot$ satisfies (W4).

It holds that $s \models (\psi \vee \psi') \odot \alpha$ if and only if $for(s) \diamond \alpha \not\models \neg(\psi \vee \psi')$, which is equivalent to $for(s) \diamond \alpha \not\models \neg\psi$ and $for(s) \diamond \alpha \not\models \neg\psi')$, *i.e.*, to $s \models \psi \odot \alpha$ or $s \models \psi \odot \alpha$, which shows that $\odot$ satisfies (W8).

Conversely, let $\odot$ satisfying (W1), (W3), (W4) and (W8). Let us show that there exists an operator $\diamond$ satisfying satisfies (U1), (U3), (W4) and (U8), such that $\odot$ is the weak reverse update associated with $\diamond$. We first note that definition of $\odot$ from $\diamond$ is symmetric: let us call the *conjugate* of a belief change operator $\star$ the belief change operator $\overline{\star}$ defined by

$$s \models for(s') \overline{\star} for(s) \text{ iff } for(s) \star \alpha for(s')$$

Then we see that if the weak reverse operator $\odot$ associated with $\diamond$ is its conjugate, *i.e.*, $\odot = \overline{\diamond}$, but also *vice versa*: $\diamond = \overline{\odot}$. Therefore, if we define $\diamond$ as the conjugate of $\odot$, $\odot$ is the weak reverse update associated with $\diamond$.

Let us now show that $\diamond = \overline{\odot}$ satisfies (U1), (U3), (U4) and (U8). Since (W4) and (W8) coincide with (U4) and (U8), exchanging $\diamond$ and $\odot$, together with the first half of the proof we immediately get that $\diamond$ satisfies (U4) and (U8).

Recall from above that in presence of (U8), $\diamond$ satifies (U1) if and only if $\odot$ satisfies (W1). Therefore, $\odot$ satisfies (W1).

As to the point concerning (U2) and (W2), assume furthermore that $\diamond$ satisfies (U2). Assume $s \models (\psi \odot \alpha) \wedge \alpha$. Suppose $s \not\models \psi$. Then there exists s' such that $s' \in R_A(s)$ and $s' \models \psi$, which implies $s \neq s'$, therefore $R_A(s) \neq \{s\}$; this, together with $for(s) \models \alpha$, violates (U2). Therefore, $s \models \psi \wedge \alpha$. Now, assume $s \models \psi \wedge \alpha$. By (U2), $R_A(s) = \{s\}$, therefore there is a $s'(= s)$ such that $s' \models \psi$ and $s' \in R_A(s)$, which shows that $s \models \psi \odot \alpha$. Therefore, $\odot$ satisfies (W2). Conversely, assume that $\diamond$ does not satisfy (U2). Then, by (U8), there exist two states s, s' and a formula α such that $s' \neq s$, $s \models \alpha$, and $s' \models for(s) \diamond \alpha$. Take $\psi = for(s')$, we have $s \models (\psi \odot \alpha) \wedge \alpha$ but $s \not\models \psi \wedge \alpha$; therefore $\odot$ does not satisfy (W2). □

Properties (U5), (U6) and (U7) do not seem seem to have meaningful counterparts for $\odot$ (and anyway, as already argued, these three postulates are controversial).

Proposition 5.5 *The strong reverse update $\otimes$ associated with a basic update operator $\diamond$ satisfies the following properties:*

S1 $\alpha \otimes \alpha \equiv \top$*;*
S3 *if α is satisfiable then $\bot \otimes \alpha \equiv \bot$;*
S4 *if $\psi \equiv \psi'$ and $\alpha \equiv \alpha'$ then $\psi \otimes \alpha \equiv \psi' \otimes \alpha'$;*
S8 $(\psi \wedge \psi') \otimes \alpha \equiv (\psi \otimes \alpha) \wedge (\psi' \otimes \alpha)$.

In addition to this, $\diamond$ satisfies (U2) if and only if $\otimes$ satisfies

S2 *if $\psi \models \alpha$ then $\psi \models \psi \otimes \alpha$.*

Note that, unlike weak reverse update, strong reverse update does generally not satisfy modelwise decomposability (U8/W8), but a symmetric, conjunctive decomposability property (S8).

Moreover, if $\diamond$ is a basic update operator then

SIW if α is satisfiable then $\psi \otimes \alpha \models \psi \odot \alpha$

Proof By Proposition 5.2, $\alpha \otimes \alpha \equiv \top$ is equivalent to $\top \diamond \alpha \models \alpha$, which is equivalent to (U1), therefore $\otimes$ satisfies (S1).

Assume $\bot \otimes \alpha \not\equiv \bot$, *i.e.*, $\bot \otimes \alpha$ is satisfiable. Then there exists s such that $s \models \bot \otimes \alpha$, which by Proposition 5.2 implies $for(s) \diamond \alpha \models \bot$, which by (U3) implies that α is unsatisfiable.

Assume $\psi \equiv \psi'$ and $\alpha \equiv \alpha'$. For any φ, by Proposition 5.2, $\varphi \models \psi' \otimes \alpha'$ is equivalent to $\varphi \diamond \alpha' \models \psi'$, which by (U4) is equivalent to $\varphi \diamond \alpha \models \psi$, which again by Proposition 5.2 is equivalent to $\varphi \models \psi \otimes \alpha$. This being true for all φ, we get that $\psi' \otimes \alpha'$ and $\psi \otimes \alpha$ are equivalent: $\otimes$ satisfies (S4).

It is straightforward from the definition of $\otimes$ that $(\psi \wedge \psi') \otimes \alpha \models \psi \otimes \alpha$; therefore, $(\psi \wedge \psi') \otimes \alpha \models (\psi \otimes \alpha) \wedge (\psi' \otimes \alpha)$. Now, let $s \models (\psi \otimes \alpha) \wedge (\psi' \otimes \alpha)$. Then by Proposition 5.2, $for(s) \diamond \alpha \models \psi$ and $for(s) \diamond \alpha \models \psi'$, therefore $for(s) \diamond \alpha \models \psi \wedge \psi'$, which again by Proposition 5.2 is equivalent to $s \models (\psi \wedge \psi') \otimes \alpha$. Hence $\otimes$ satisfies (S8).

Finally, let ψ and α be such that $\psi \models \alpha$. Then by Proposition 5.2, $\psi \models \psi \otimes \alpha$ is equivalent to $\psi \diamond \alpha \models \psi$, which is entailed by (U2). Therefore, if $\diamond$ satisfies (U2) then $\otimes$ satisfies (S2). For the converse, assume $\otimes$ satisfies (S2) and $s \models \psi$. Then $s \models \alpha$, and by (S2) we get $for(s) \models for(s) \otimes \alpha$, which by definition of $\otimes$ is equivalent to $for(s) \diamond \alpha \models for(s)$. Now, by (U3), $for(s) \diamond \alpha \models for(s)$ implies that $for(s) \diamond \alpha \equiv for(s)$, which by (U8) implies $\psi \diamond \alpha \equiv \psi$: $\diamond$ satisfies (U2). □

Note that (SIW) fails without (U3). Example 5.3 shows that the converse implication of (SIW) does not hold in general. Finally, $\otimes$ and $\odot$ coincide if and only if $update(\diamond, \alpha)$ is deterministic.

One may wonder whether reverse update has something to do with erasure [21]. An erasure operator $\blacklozenge$ is defined from an update operator $\diamond$ by $\psi \blacklozenge \alpha \equiv \psi \vee (\psi \diamond \neg\alpha)$. Erasing by α intuitively consists in making the world evolve (following some rules) such that after this evolution, the agent no longer believes α. A quick look suffices to understand that erasure has nothing to do with weak and strong reverse update. Erasure corresponds to action progression for an action $erase(\alpha)$ whose effect is be epistemically negative (*make* α *disbelieved*). This implies in particular that $\top \blacklozenge \top$ is always unsatisfiable ($\top$ cannot be made disbelieved) whereas $\top \odot \top \equiv \top \otimes \top \equiv \top$. To give another short example: if $\diamond = \diamond_{PMA}$, then $(a \leftrightarrow \neg b) \blacklozenge_{PMA} b \equiv (\neg a \vee \neg b)$, whereas $(a \leftrightarrow \neg b) \odot_{PMA} b \equiv (a \leftrightarrow \neg b) \otimes_{PMA} b \equiv \neg a$.

Pursuing the investigation on reverse update does not only have a theoretical interest: weak (deductive) reverse update allows for postdiction, and strong (abuctive) reverse update allows for goal regression (when the actions performed are updates) and is therefore crucial if we want to use an update-based formalism for planning (see [25]).

6 Update as Counterfactual Reasoning

There is another prominent interpretation of belief update, which *a priori* does not seem to be related to feedback-free action progression: counterfactual reasoning and causality. Let me quote Leitgeb and Segerberg [26], pp. 184–185:

> The intended interpretation of the semantics for belief update depends crucially on the manner in which selection functions are interpreted. The standard interpretation is in terms of environmental change; but there is another plausible way of interpreting selection functions, one that enables us to demonstrate that update does not necessarily correspond to environmental changes. Lewis famously considered objective similarity relations between possible worlds to be determinable from the objective spheres systems (...). This, given new evidence, we find that in the case of belief revision the agent tries to change his beliefs in a way such that the worlds that he subsequently believes to be in comprise the *subjectively most plausible deviation* from the worlds he originally believed to inhabit. However, when confronted in the same evidence in belief update, the agent tries to change his beliefs in a way such that the worlds that he subsequently believes to be in are *as objectively similar as possible* to the worlds he originally believed to be the most plausible candidates for being the actual world.

This is in agreement with Grahne's relationship between updates and counterfactuals [13]. Dupin de Saint-Cyr [9] goes further and argues that belief update is the right operation to deal with causality: the fact that α was true (respectively, that some event ϵ took place) at some time point t causes φ to be true at $t' > t$ is equivalent to saying that updating the past of the system by the fact that α was false (respectively, that ϵ did not take place) at t allows to derive that $\neg\varphi$ holds at t'. Updating the past in such a way requires selecting objectively most similar worlds that satisfy the condition part of the counterfactual ($\neg\varphi_t$ or $\neg\epsilon_t$).

Is counterfactual reasoning a radically different interpretation from feedback-free action progression? The traditional view of action progression only involves reasoning about the agent's future beliefs given her current beliefs and the knowledge of the action that is taking place now. Performing an action whose effects take place in the past does not look particularly intuitive at first sight. We argue that updating the past (in order to assess a causality statement) does however correspond to some form of action progression.

Technically, this is clear. The actions involved here act on the whole history. As in [9], consider a time-stamped language generated by propositional variables of the form v_t. A world τ is a full trajectory $\langle s_t, t \in T \rangle$ consisting of a full state at each time point. A temporal formula is a formula α built on the alphabet $\{v_t, t \in T\}$. Updating τ by α is conceptually no different from updating a world by a propositional formula in standard belief update. Updating a temporal formula β by a temporal formula α consists in taking the union of all $\tau \diamond \alpha$ for all trajectories τ satisfying β.

From a philosophical point of view, this is less obvious and we proceed by giving first an analogy between time and space. Consider the following counterfactual statement: if event ϵ had occurred at time point t, would p had been true at time point t'? This is equivalent to check whether (a) $\beta \models \neg p_t$ and (b) $\beta \diamond \epsilon_t \models p_{t'}$. Clearly, the part of the knowledge history β that takes place before t should remain unchanged: for every temporal formula γ involving only time-stamped variables $p_{t''}$ with $t" < t$, we should have $\beta \diamond \epsilon_t \models \gamma$ if and only if $\beta \models \gamma$. Now, consider a series of cells, horizontally connected, with a gate between cell i and cell $i+1$ that can be pushed and opened from i but not from $i+1$: when pushed from the left side towards the right, they open, but when pushed from the right towards the left, they do not. Suppose now that we perform an action in cell i that may increase the pressure, which in turn can lead to increase the pressure in cells $i+1$, $i+2$ etc. and possibly other side effects. Because the doors cannot open from right to left, nothing changes in cells $j < i$. (One can also imagine some information passing between cells that is possible only from the left to the right). It is not difficult to see that the strong left-to-right orientation of space is analogous to the past-to-future orientation of time. Asking whether making α true at cell i results in ψ holding at cell $j > i$ corresponds to asking whether the event of making α true at time t would result in ψ holding at time $t' > t$.

As a second example, consider a fiction writer who has built a scenario for a novel; the temporal formula β represents the beliefs of the reader at each time point (obviously, β is not necessarily complete). We assume here that these beliefs are correct, *i.e.*, the reader is never misled. The author is then asked by the publisher to change the scenario so that a particular temporal formula α be true (and known by the reader). This requires the writer to update β by α. Making α true is an action that can have effects on the whole history, *including maybe at time points earlier to those concerned by* α: it may indeed be simpler for the writer to adapt his novel so that x_t now holds by changing facts at time points $t' < t$. Although this is another example of updating the past, the possible influence from future to past make it radically different from updates used in counterfactual reasoning.

7 Updates and DDL

As developed in length in Krister Segerberg's works on Dynamic Doxastic Logic, there are many reasons why it is tempting to "express doxastic actions such as belief revision on the object language level". This, however, raises a serious issue: the failure of the Ramsey test. Quoting [25], p. 171:

> (...) DDL is bound to face a serious challenge: the danger of getting entangled in the potentially paradoxical of combining belief revision for an object language F with a representation of the revision operator in terms of formulæ in F.
>
> The possibly devastating effects of such a combination first showed up when Gärdenfors considered a doxastic interpretation of conditionals in terms of the so-called Ramsey test for conditionals.
>
> $$\varphi \Rightarrow \theta \quad \text{iff} \quad \theta \in K \star \varphi$$

Indeed, Gärdenfors shows in [12] that as soon as the language contains at least three propositions that are pairwise consistent but jointly inconsistent, the AGM axioms of $\star$ are inconsistent with the Ramsey test for conditionals. The implications of Gärdenfors' impossibility result, to DDL, and the two ways to escape it, are discussed in [25], p. 172. As noticed by Herzig [16] and by Leitgeb and Segerberg [15], Gärdenfors' impossibility result does not carry over to belief update, and indeed, quoting from [15], "most standard systems of conditional logic support update operations". The intuitive reason for this lies in this ([25], p. 186):

> (...) given a body of beliefs [about the ways in which the environment may change] and an initial state of beliefs [about the current state of the environment], in KGM all future beliefs [about the current state of the environment] are determined by reports of what happens. So KGM, unlike basic AGM, is a theory of iterated belief change.

And indeed, iteration in belief update does not cause any particular problem. In the view of the discussion of Sect. 4, this should not be seen as surprising: recall that belief update is a particular kind of action progression, and action progression is naturally iterated. More than that, belief update can, just as action progression, be generalized not only to sequences of updates but also to *conditional updates*, *nondeterministic updates*, and *concurrent updates*. A *nondeterministic update* [4, 16] $\alpha \cup \beta$ corresponds to the nondeterministic choice of the two updates α and β. A *conditional update* [16] $\textsf{if}\ \varphi\ \textsf{then}\ \alpha\ \textsf{else}\ \beta$ corresponds to an update by α if φ holds and by β otherwise. A *concurrent update* [16] $\alpha || \beta$ corresponds to the simulatenous execution of an update by α and an update by β. These constructs, which can be applied recursively, considerable enrich the language of belief update and makes it more suitable to express planning problems.

Now that we know that updates are a specific class of feedback-free actions, associated with transition systems, it makes even more sense to use DDL-KGM for expressing interactions between actions and beliefs, where $\diamond\alpha$ denotes the action of updating by α. As we argued already, the specificity of feedback-free actions is the *what you foresee is what you get* axiom, which is expressed in DDL by

$$\mathbf{B}[\diamond\alpha]\varphi \equiv [\diamond\alpha]\mathbf{B}\varphi$$

which, of course, does not hold for sensing actions or more generally actions that may bring some feedback. Progression and regression can also be expressed in DDL-KGM. The axiom

$$(Prog) \quad (\mathbf{B}\varphi \to \mathbf{B}[\diamond\alpha]\psi) \equiv (prog(\varphi, \alpha) \to \psi)$$

actually gives a *definition* of progression, i.e., a unique characterization of $prog(\varphi, \alpha)$ up to logical equivalence; and similarly for weak and strong regression:

$$(WR) \quad ([\diamond\alpha]\mathbf{B}\psi \to \mathbf{B}\varphi) \to (reg(\psi, \alpha) \to \varphi)$$

$$(SR) \quad (\mathbf{B}\varphi \to [\diamond\alpha]\mathbf{B}\psi) \to (\varphi \to Reg(\psi, \alpha))$$

There is no reason to stop here. For instance, we may integrate DDL-AGM and DDL-KM and express something like that

$$[\star([\diamond\alpha]\mathbf{B}\psi)]\mathbf{B}\varphi$$

expressing that after learning that updating by α would make ψ true, I now believe that it is the case that φ. (As an example, take $\diamond$ to be $\diamond_{PMA}$, and $\alpha = a \vee b$, $\psi = a \leftrightarrow \neg b$, $\varphi = \neg a \vee \neg b$.)

8 Summary and Conclusion

Let us try to summarize what we have said so far. Both revision and update deal with dynamic worlds, but they strongly differ in the nature of the information they process. Belief revision (together with the introduction of time stamps in the propositional language) aims at correcting some initial beliefs about the past, the present, and even the future state of the world by some newly *observed* information about the past or the present state of the world. Belief update is suitable only for (some specific) action progression without feedback: updating φ by α corresponds to progressing (or projecting forward) φ by the action $update(\diamond, \alpha)$, to be interpreted as *make* α *true*. The "input formula" α is the effect of the action $update(\diamond, \alpha)$, and definitely not an observation. Expressed in the terminology of Sandewall [31], the range of applicability of update is the class K_{p}-**IA**: correct knowledge,[10] no observations after the initial time point, inertia if (U2) is assumed, and alternative results of actions.

In complex environments, especially planning under incomplete knowledge, actions are complex and have both ontic and epistemic effects; the belief change

[10] However, this point is somewhat debatable: update would work as well if we don't assume that the agent's initial beliefs is correct—of course, in this case the final beliefs may be wrong as well.

process then is very much like the feedback loop in partially observable planning and control theory: perform an action, project its effects on the current belief state, then get the feedback, and revise the projected belief state by the feedback. Clearly, update allows for projection only. Or, equivalently, if one chooses to separate the ontic and the epistemic effects of actions, by having two disjoint sets of actions (ontic and epistemic), then ontic actions lead to projection only, while epistemic actions lead to revision only. Therefore, if one wants to extend belief update so as to handle feedback, there is no choice but integrating some kind of revision process, as in [3, 19, 20, 35]. Another possibility is to generalize update so that it works in a language that distinguishes facts and knowledge, such as epistemic logic S5: this *knowledge update* process is investigated by Baral and Zhang [2]. Here, effects of sensing actions are handled by *updating* (and not revising) formulas describing the agent's knowledge. Such a framework takes the point of view of a modelling agent O who reasons an the state of knowledge of another agent ag. Thus, for instance, updating a S5 model by $K_{ag}\varphi$ means that the O updates her beliefs about ag's knowledge; considering ag's mental state as part of the *outside world* for agent O, this suits our view of update as a feedback-free action for O (updating by $K_{ag}\varphi$ corresponds as "make $K_{ag}\varphi$ true", which can for instance be implemented by telling ag that φ is true).

Acknowledgments In my conference paper [23], I wrote that I would never have thought of writing that chapter without these years of discussion with Andreas Herzig about the very meaning of belief update. This is still true now, with a few more years in the count.

References

1. Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, *50*(2), 510–530.
2. Baral, C., & Zhang, Y. (2005). Knowledge updates: Semantics and complexity issues. *Artificial Intelligence*, *164*(1–2), 209–243.
3. Boutilier, C. (1998). A unified model of qualitative belief change: A dynamical systems perspective. *Artificial Intelligence*, *98*(1–2), 281–316.
4. Brewka, G., & Hertzberg, J. (1993). How to do things with worlds: On formalizing actions and plans. *Journal of Logic and Computation*, *3*(5), 517–532.
5. del Val, A., & Shoham, Y. (1994). Deriving properties of belief update from theories of action. *Journal of Logic, Language, and Information*, *3*, 81–119.
6. Doherty, P., Łukasziewicz, W., & Madalińska-Bugaj, E. (1998) The PMA and relativizing change for action update. In *Proceedings of the KR'98*, pp. 258–269.
7. Dubois, D., & Prade, H. (1993). Belief revision and updates in numerical formalisms: An overview, with new results for the possibilistic framework. In *Proceedings of IJCAI'93*, pp. 620–625.
8. Dupin de Saint-Cyr, F., & Lang, J. (2011). Belief extrapolation (or how to reason about observations and unpredicted change). *Artificial Intelligence*, *175*(2), 258–269.
9. Dupin de Saint-Cyr, F. (2008). Scenario update applied to causal reasoning. In *Proceedings of the Eleventh International Conference on the Principles of Knowledge Representation and Reasoning (KR 2008)* (pp. 188–197). AAAI Press.

10. Friedman, N., & Halpern, J. Y. (1999). Modelling beliefs in dynamic systems. Part ii: Revision and update. *JAIR*, *10*, 117–167.
11. Friedman, N., & Halpern, J. Y. (1996). Belief revision: A critique. In *Proceedings of the KR'96*, pp. 421–431.
12. Gärdenfors, P. (1986). Belief revisions and the Ramsey test for conditionals. *Philosophical Review*, *95*, 81–93.
13. Grahne, G. (1991). Updates and counterfactuals. In *KR*, pp. 269–276.
14. Herzig, A. (1996). The PMA revisited. In *Proceedings of the KR'96*, pp. 40–50.
15. Herzig, A. (2000). Logics for belief base updating. In D. Gabbay & Ph. Smets (Eds.), *Handbook of defeasible reasoning and uncertainty management systems*. Dordrecht: Kluwer Academic Publishers.
16. Herzig, A., Lang, J., Marquis, P., & Polacsek, T. (2001). Updates, actions, and planning. In *Proceedings of IJCAI'01*, pp. 119–124.
17. Herzig, A., Lang, J., & Marquis, P. (2011). Propositional update operators based on formula/literal dependence. Technical report.
18. Herzig, A., & Rifi, O. (1999). Propositional belief update and minimal change. *Artificial Intelligence*, *115*, 107–138.
19. Hunter, A., & Delgrande, J. (2005). Iterated belief change: A transition system approach. In *Proceedings of IJCAI'05*, pp. xxx–yyy.
20. Jin, Y., & Thielscher, M. (2004). Representing beliefs in the fluent calculus. In *Proceedings of ECAI-04*, pp. 823–827.
21. Katsuno, H., & Mendelzon, A. (1991). On the difference between updating a knowledge base and revising it. In *Proceedings of KR'91*, pp. 387–394.
22. Katsuno, H., & Mendelzon, A. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, *52*, 263–294.
23. Lang, J. (2007). Belief update revisited. In *Proceedings of IJCAI'07*, pp. 2517–2522.
24. Lang, J., Lin, F., & Marquis, P. (2003). Causal theories of action: A computational core. In *Proceedings of IJCAI-03*, pp. 1073–1078.
25. Leitgeb, H., & Segerberg, K. (2007). Dynamic doxastic logic: Why, how, and where to? *Synthese (Knowledge, Rationality and Action)*, *155*, 167–190.
26. Lewis, D. (1973). *Counterfactuals*. Cambridge: Harvard University Press.
27. Liberatore, L. (2000). A framework for belief update. In *Proceedings of JELIA'00*, pp. 361–375.
28. Lindström, S., & Segerberg., K. (2006). *Modal logic and philosophy*. Dordrecht: Elsevier.
29. Lindström, S., & Rabinowicz, W. (1990). Epistemic entrenchment with incomparabilities and relational belief revision. In A. Fuhrmann & M. Morreau (Eds.), *The logic of theory change, Lectures Notes in Artificial Intelligence* (Vol. 645, 93–126). Berlin: Springer
30. Peppas, P., Nayak, A., Pagnucco, M., Foo, N., Kwok, R., & Prokopenko, M. (1996). Revision versus update: Taking a closer look. In *Proceedings of ECAI96*.
31. Sandewall, E. (1995). *Features and fluents*. Oxford: Oxford University Press.
32. Segerberg, K. (1989). Bringing it about. *Journal of Philosphical Logic*, *18*, 327–347.
33. Segerberg, K. (1998). Irrevocable belief revision in dynamic doxastic logic. *Notre Dame Journal of Formal Logic*, *39*, 287–306.
34. Segerberg, K. (2001). *The basic dynamic doxastic logic of AGM* (pp. 57–84). Dordrecht: Kluwer.
35. Shapiro, S., & Pagnucco, M. (2004). Iterated belief change and exogeneous actions in the situation calculus. In *Proceedings of the ECAI04*.
36. Winslett, M. (1990). *Updating logical databases*. Cambridge: Cambridge University Press.