On the Logic of Theory Change Iteration of KM-Update, Revised

Liangda Fang[1,3], **Tong Zhu**[1], **Quanlong Guan**[1*], **Junming Qiu**[2*],
Zhao-Rong Lai[1], **Weiqi Luo**[1] and **Hai Wan**[2]
[1] Jinan University, Guangzhou 510632, China
[2] Sun Yat-sen University, Guangzhou 510006, China
[3] Pazhou Lab, Guangzhou 510330, China
{fangld, gql}@jnu.edu.cn, zhutong62@stu2020.jnu.edu.cn, qiujm9@mail2.sysu.edu.cn

Abstract

Belief revision and update, two significant types of belief change, both focus on how an agent modifies her beliefs in presence of new information. The most striking difference between them is that the former studies the change of beliefs in a static world while the latter concentrates on a dynamically-changing world. The famous AGM and KM postulates were proposed to capture rational belief revision and update, respectively. However, both of them are too permissive to exclude some unreasonable changes in the iteration. In response to this weakness, the DP postulates and its extensions for iterated belief revision were presented. Furthermore, Fermé and Gonçalves integrated these postulates in belief update. Unfortunately, some redundant components are included in the definitions of belief states and the faithful assignments for semantic characterizations. Moreover, their approach does not meet the desired property of iterated belief update. They also do not discuss the rationale of any DP postulate within the update context. This paper is intended to address all these shortcomings of Fermé and Gonçalves's approach. Firstly, we present a modification of the original KM postulates based on belief states, and propose the notion of faithful collective assignments of belief states to partial preorders. Subsequently, we migrate several well-known postulates for iterated belief revision to iterated belief update. Moreover, we provide the exact semantic characterizations based on partial preorders for each of the proposed postulates. Finally, we analyze the compatibility between the above iterated postulates and the KM postulates for belief update.

1 Introduction

Belief revision focuses on how an agent changes her beliefs when she encounters new information inconsistent with her initial beliefs. The notable AGM postulates, proposed by Alchourrón *et al.* [1985], became a standard framework to capture the rational behavior of belief revision. Katsuno and Mendelzon [1991b] proposed a characterization of all revision operators that satisfy AGM postulates in terms of total preorders over possible worlds.

Belief update, another significant type of belief change, concentrates on how an agent will modify her beliefs about a dynamically-changing world in view of new information. As in belief revision, Katsuno and Mendelzon [1991a] proposed the KM postulates for regulating rational belief update, which models the process of update as a function of belief sets. Furthermore, they offered a semantic characterization based on partial preorders over possible worlds, and clarified the distinctions between update and revision from the model-theoretic perspective. There is only one total preorder for the belief set $\mathcal{K}$ in belief revision. In contrast, in belief update, a collection of partial preorders is induced by $\mathcal{K}$ where each preorder is associated to each possible world satisfying $\mathcal{K}$.

Although the AGM postulates were considered as a basic framework for belief revision, it is shown to be too permissive to exclude some unreasonable revision operators in the iteration [Darwiche and Pearl, 1997]. The reason can be attributed to the fact that it is comprised of merely a set of one-step postulates, failing to properly deal with the sequential new information in the process of iterated belief revision. To remedy this defect in belief revision, Darwiche and Pearl supplemented the AGM paradigm with four postulates (C1)-(C4) (called DP postulates) and use belief states to denote the belief of an agent instead of belief sets. Two different belief states may have the same belief sets, but not vice versa.

Likewise, we argue that the same problem as above exists for the KM postulates in belief update. In details, the KM framework is unable to regulate the preferences for subsequent updates during the iterated update process, leading to some counter-intuitive results. We use the following example to briefly illustrate the problem.

Example 1. *Let us consider a table with two zones: left and right. There are a book, a cup and a toy on any side on the table. We denote by b (resp. c / t) the proposition "the book (resp. cup / toy) is on the left zone of the table".*

Initially, everything is on the right zone, and hence the initial belief set $\mathcal{K}$ is $\neg b \wedge \neg c \wedge \neg t$. We first instruct a robot to move at least one of the book and the cup to the left zone, which is described as the new information $\phi = b \vee c$. Suppose that updated belief set $\mathcal{K} \diamond \phi$ is $(b \leftrightarrow \neg c) \wedge \neg t$, that is, exactly

*Both are corresponding authors.

one of the book and cup is on the left zone, and the toy is on the right zone. The robot is then given a fresh instruction to place the cup on the left, which is the new information $\varphi = c$. According to the KM postulates, it is acceptable that the updated belief set $(\mathcal{K} \diamond \phi) \diamond \varphi$ becomes $\neg b \wedge c$. However, after updating by ϕ, the agent has believed that the toy was on the right zone. It seems unreasonable to require the agent to give up this belief after putting the cup on the left zone. □

The two updated beliefs $\mathcal{K} \diamond \phi$ and $(\mathcal{K} \diamond \phi) \diamond \varphi$ in the above example are both KM-compatible. Since the KM postulates do not constrain the update strategy after iteration, counter-intuitive results will emerge after multiple iterations.

One topic worth investigation is what will happen if we incorporate the iterated revision postulates above to a belief update scenario. Following the semantic characterization for every iterated postulate in belief revision, Fermé and Gonçalves [2023] integrated these postulates in belief update. However, Fermé and Gonçalves's approach has the following shortcomings. (1) Each belief state defined in [Fermé and Gonçalves, 2023] includes not only the belief set but also an update operator for belief sets, that is, a function mapping a belief set and a sentence to a new belief set. In contrast, a belief state defined in [Darwiche and Pearl, 1997] encodes the preference about possible worlds via faithful assignments and is associated with only the belief set. The revision operator for belief states modifies the preference about possible worlds, and do not rely on any revision operator for belief sets. Hence, the associated update operator of belief states in [Fermé and Gonçalves, 2023] is redundant. (2) The fundamental concept of the semantic definition is faithful complete assignments, that is, a function that maps each belief state to a complete collection of partial preorders. A complete collection of partial preorders is such that there is a unique partial preorder $\leq_w$ for every possible world w. However, in the semantic definition of update operators, it is unnecessary to consider the partial-preorders centered at the possible world that is not a model of the associated belief set. (3) The semantic definition of update operators for belief states [Fermé and Gonçalves, 2023] is based on the initial belief set rather than the most recent belief set. This obviously violates the desired property of iterated belief change including iterated belief revision. (4) DP postulates was defined for iterated belief revision. Fermé and Gonçalves chose DP postulates as the basis for iterated belief update, but did not discuss the rationale of any DP postulate within the update context.

This paper is intended to overcome the above shortcomings of Fermé and Gonçalves's approach. The main contributions are the following. We first adopt the original definition of belief states, proposed by Darwiche and Pearl [1997], within update context and modify the two postulates (U◆4) and (U◆8) defined in [Fermé and Gonçalves, 2023] as the update operator for belief sets occur explicitly in these two postulates. We then propose the notion of faithful collective assignments, that is, a function that maps each belief state to a collection of partial preorders which is centered at one possible world of the associated belief set. Furthermore, we give the new semantic definition of updating belief states, which is based on the possible worlds satisfying the most recent belief sets rather than the initial ones. We also provide the exact semantic characterizations via partial preorders for every DP postulate of iterated belief update. We offer some concrete examples to illustrate the motivation of the usage of DP postulates in iterated belief update. Finally, we analyze the compatibility between the DP iterated postulates and the modified KM postulates for belief update. We identify an update operator that satisfies (CU3) and (CU4). In particular, we show that each of (CU1) and (CU2) is inconsistent with the KM postulates.

2 Formal Preliminaries

Throughout this paper, we fix a finite set $\mathcal{P}$ of *propositional variables*. We define $\mathcal{L}$ to be the *propositional language* built from $\mathcal{P}$, the connectives $\neg$, $\wedge$ and $\vee$, and two Boolean constants $\top$ (truth) and $\bot$ (falsity).

A propositional sentence ψ is *complete*, iff for every sentence $\varphi \in \mathcal{L}$, $\psi \models \varphi$ or $\psi \models \neg\varphi$. A propositional sentence ψ is *consistent*, iff there is no sentence $\varphi \in \mathcal{L}$ s.t. $\psi \models \varphi$ and $\psi \models \neg\varphi$. A *possible world* w is a complete consistent set of literals over $\mathcal{P}$, *i.e.*, for every $p \in \mathcal{P}$, either $p \in w$ or $\neg p \in w$. We use $\mathcal{M}$ to denote the set of all possible worlds, and $[\phi]$ to denote the set of all possible worlds in which ϕ holds. For a finite set of worlds W, $\texttt{Form}(W)$ denotes the sentence $\bigvee_{w \in W} \left(\bigwedge_{l \in w} l\right)$. For ease of representation, we sometimes use $w_1, \cdots, w_n$ to denote the set $\{w_1, \cdots, w_n\}$.

A *(partial) preorder* $\leq$ over $\mathcal{M}$ is a reflexive, transitive binary relation on $\mathcal{M}$. A preorder is *total* if for all $w, w' \in \mathcal{M}$, either $w \leq w'$ or $w' \leq w$. We define $<$ as the strict part of $\leq$, *i.e.*, $w < w'$ iff $w \leq w'$ and $w' \not\leq w$. We define $\approx$ as the symmetric part of $\leq$, *i.e.*, $w \approx w'$ iff $w \leq w'$ and $w' \leq w$.

3 Background

In this section, we briefly review the axiomatic and semantic characterization of belief revision [Alchourrón *et al.*, 1985], iterated belief revision [Darwiche and Pearl, 1997], belief update [Katsuno and Mendelzon, 1991a] and iterated belief update [Fermé and Gonçalves, 2023].

3.1 Belief Revision

The original AGM paradigm models the notion of belief revision as a function that maps a belief set $\mathcal{K}$ and a sentence ϕ to a new belief set $\mathcal{K} \circ \phi$. In this paper, a belief set is defined as a propositional sentence. However, the revision function over belief sets only differentiates what the agent believes and what she does not believe, but does not compare the plausibility degree of different information which the agent does not believe. This leads to improper behaviors of revision functions over belief sets on iterated belief revision [Darwiche and Pearl, 1997]. To fix this defect, Darwiche and Pearl proposed the notion of belief states (also referred to as epistemic states), redefined revision function on belief states, and reformulated the AGM postulates accordingly. Darwiche and Pearl did not provide a standard definition of belief states, and only required that each belief state S is associated with a belief set $\texttt{B}(S)$. We use this abstract representation of the belief state herein.

The following are the modified AGM postulates which shift from belief sets to belief states and are originated from [Darwiche and Pearl, 1997].

(R*1) $\mathtt{B}(S \circ \varphi) \models \varphi$.

(R*2) If $\mathtt{B}(S) \wedge \varphi$ is consistent, then $\mathtt{B}(S \circ \varphi) \equiv \mathtt{B}(S) \wedge \varphi$.

(R*3) If φ is consistent, so is $\mathtt{B}(S \circ \varphi)$.

(R*4) If $S_1 = S_2$ and $\varphi \equiv \phi$, then $\mathtt{B}(S_1 \circ \varphi) \equiv \mathtt{B}(S_2 \circ \phi)$.

(R*5) $\mathtt{B}(S \circ \varphi) \wedge \phi \models \mathtt{B}(S \circ (\varphi \wedge \phi))$.

(R*6) If $\mathtt{B}(S \circ \varphi) \wedge \phi$ is consistent, then $\mathtt{B}(S \circ (\varphi \wedge \phi)) \models \mathtt{B}(S \circ \varphi) \wedge \phi$.

Darwiche and Pearl provided a representation theorem for the modified AGM postulates based on the notion of faithful total preorders and faithful assignments proposed by Katsuno and Mendelzon [1991b]. Given a sentence φ, a total preorder $\leq_\varphi$ over $\mathcal{M}$ is *faithful to* φ, iff (1) $w \approx_\varphi w'$ for every two possible worlds $w, w' \in [\varphi]$; and (2) $w <_\varphi w'$ for every two possible worlds $w \in [\varphi]$ and $w' \in [\neg\varphi]$. A *faithful assignment* is a function that maps each belief state S to a total preorder $\leq_S$ that is faithful to $\mathtt{B}(S)$. The following is the representation theorem for the modified AGM postulates.

Theorem 1 ([Darwiche and Pearl, 1997])**.** *A revision operator $\circ$ satisfies postulates (R*1)-(R*6) iff there is a faithful assignment that maps each belief state S to a total preorder $\leq_S$ s.t.* $[\mathtt{B}(S \circ \varphi)] = \min([\varphi], \leq_S)$.

3.2 Iterated Belief Revision

In the AGM paradigm, there is no guidance on the relationship between the initial revision strategy and the subsequent one. To solve this problem, Darwiche and Pearl proposed DP postulates that describe rational iterated belief revision.

(C1) If $\varphi \models \phi$, then $\mathtt{B}((S \circ \phi) \circ \varphi) \equiv \mathtt{B}(S \circ \varphi)$.

(C2) If $\varphi \models \neg\phi$, then $\mathtt{B}((S \circ \phi) \circ \varphi) \equiv \mathtt{B}(S \circ \varphi)$.

(C3) If $\mathtt{B}(S \circ \varphi) \models \phi$, then $\mathtt{B}((S \circ \phi) \circ \varphi) \models \phi$.

(C4) If $\mathtt{B}(S \circ \varphi) \not\models \neg\phi$, then $\mathtt{B}((S \circ \phi) \circ \varphi) \not\models \neg\phi$.

Darwiche and Pearl proved that postulates (C1)-(C4) correspond to the following semantics constraints (CR1)-(CR4) on possible worlds, respectively.

(CR1) If $w_1, w_2 \in [\phi]$, then $w_1 \leq_S w_2$ iff $w_1 \leq_{S \circ \phi} w_2$.

(CR2) If $w_1, w_2 \in [\neg\phi]$, then $w_1 \leq_S w_2$ iff $w_1 \leq_{S \circ \phi} w_2$.

(CR3) If $w_1 \in [\phi]$ and $w_2 \in [\neg\phi]$, then $w_1 <_S w_2$ only if $w_1 <_{S \circ \phi} w_2$.

(CR4) If $w_1 \in [\phi]$ and $w_2 \in [\neg\phi]$, then $w_1 \leq_S w_2$ only if $w_1 \leq_{S \circ \phi} w_2$.

Theorem 2 ([Darwiche and Pearl, 1997])**.** *Let $\circ$ be a revision operator satisfying (R*1)-(R*6). Then, $\circ$ satisfies (Ci) iff the operator and its corresponding faithful assignment satisfies (CRi) for* $1 \leq i \leq 4$.

3.3 Belief Update

In the belief update literature, an agent is intended to modify her beliefs about a dynamically-changing environment in view of new information. Katsuno and Mendelzon [1991a] clarified the distinction between belief revision and update. Furthermore, following the AGM paradigm, they presented the KM postulates to characterize a family of rational belief update functions, which map a belief set $\mathcal{K}$ and a sentence φ to a new belief set $\mathcal{K} \diamond \varphi$.

(U1) $\mathcal{K} \diamond \varphi \models \varphi$.

(U2) If $\mathcal{K} \models \varphi$, then $\mathcal{K} \diamond \varphi \equiv \mathcal{K}$.

(U3) If both $\mathcal{K}$ and φ are consistent, so is $\mathcal{K} \diamond \varphi$.

(U4) If $\varphi \equiv \phi$, then $\mathcal{K} \diamond \varphi \equiv \mathcal{K} \diamond \phi$.

(U5) $(\mathcal{K} \diamond \varphi) \wedge \phi \models \mathcal{K} \diamond (\varphi \wedge \phi)$.

(U6) If $\mathcal{K} \diamond \varphi \models \phi$ and $\mathcal{K} \diamond \phi \models \varphi$, then $\mathcal{K} \diamond \varphi \equiv \mathcal{K} \diamond \phi$.

(U7) If $\mathcal{K}$ is complete, then $(\mathcal{K} \diamond \varphi) \wedge (\mathcal{K} \diamond \phi) \models \mathcal{K} \diamond (\varphi \vee \phi)$.

(U8) $(\mathcal{K}_1 \vee \mathcal{K}_2) \diamond \varphi \equiv (\mathcal{K}_1 \diamond \varphi) \vee (\mathcal{K}_2 \diamond \varphi)$.

To describe the process of belief update, Katsuno and Mendelzon proposed the notion of *faithful preorder* associated with possible worlds. Formally, given a possible world w, a preorder $\leq_w$ over $\mathcal{M}$ is *faithful to* w, iff for every possible world w', $w' \neq w$ only if $w <_w w'$. A *faithful pointwise assignment* is a function that maps each possible world w to a partial preorder that is faithful to w. The following theorem shows the semantic characterization of KM postulates.

Theorem 3 ([Katsuno and Mendelzon, 1991a])**.** *An update operator $\diamond$ satisfies postulates (U1)-(U8) iff there is a faithful pointwise assignment that maps each possible world w to a partial preorder $\leq_w$ s.t.* $[\mathcal{K} \diamond \varphi] = \bigcup_{w \in [\mathcal{K}]} \min([\varphi], \leq_w)$.

3.4 Iterated Belief Update

Similarly to AGM paradigm, the KM framework ignores iterations in update process. To address this issue, Fermé and Gonçalves [2023] extended KM postulates for belief states so as to capture rational iterated update. In Fermé and Gonçalves's framework, each belief state S is associated to an initial belief set $\mathtt{B}_i(S)$, a most recent belief set $\mathtt{B}(S)$ and a belief update operator $\mathtt{O}(S)$ satisfying the KM postulates (U1)-(U8). The initial belief set $\mathtt{B}_i(S_i \blacklozenge \varphi_1 \cdots \blacklozenge \varphi_n)$ is defined as $\mathtt{B}(S_i)$ where S_i is the initial belief state of $S_i \blacklozenge \varphi_1 \cdots \blacklozenge \varphi_n$. The update operator $\blacklozenge$ for belief states satisfies the following two properties: (1) $\mathtt{B}(S \blacklozenge \varphi) = \mathtt{B}(S) \diamond \varphi$; and (2) if $\mathtt{O}(S_1) = \mathtt{O}(S_2)$, then $\mathtt{O}(S_1 \blacklozenge \varphi) = \mathtt{O}(S_2 \blacklozenge \varphi)$.

(U$\blacklozenge$1) $\mathtt{B}(S \blacklozenge \varphi) \models \varphi$.

(U$\blacklozenge$2) If $\mathtt{B}(S) \models \varphi$, then $\mathtt{B}(S \blacklozenge \varphi) \equiv \mathtt{B}(S)$.

(U$\blacklozenge$3) If both $\mathtt{B}(S)$ and φ are consistent, so is $\mathtt{B}(S \blacklozenge \varphi)$.

(U$\blacklozenge$4) If $\mathtt{B}(S_1) \equiv \mathtt{B}(S_2)$, $\mathtt{O}(S_1) = \mathtt{O}(S_2)$ and $\varphi \equiv \phi$, then $\mathtt{B}(S_1 \blacklozenge \varphi) \equiv \mathtt{B}(S_2 \blacklozenge \phi)$.

(U$\blacklozenge$5) $\mathtt{B}(S \blacklozenge \varphi) \wedge \phi \models \mathtt{B}(S \blacklozenge (\varphi \wedge \phi))$.

(U$\blacklozenge$6) If $\mathtt{B}(S \blacklozenge \varphi) \models \phi$ and $\mathtt{B}(S \blacklozenge \phi) \models \varphi$, then $\mathtt{B}(S \blacklozenge \varphi) \equiv \mathtt{B}(S \blacklozenge \phi)$.

(U◆7) If $\mathtt{B}(S)$ is complete, then $\mathtt{B}(S \blacklozenge \varphi) \wedge \mathtt{B}(S \blacklozenge \phi) \models \mathtt{B}(S \blacklozenge (\varphi \vee \phi))$.

(U◆8) If $\mathtt{B}(S_1) \equiv \mathtt{B}(S_2) \vee \mathtt{B}(S_3)$ and $\mathtt{O}(S_1) = \mathtt{O}(S_2) = \mathtt{O}(S_3)$, then $\mathtt{B}(S_1 \blacklozenge \varphi) \equiv \mathtt{B}(S_2 \blacklozenge \varphi) \vee \mathtt{B}(S_3 \blacklozenge \varphi)$.

A collection $\{\leq_w\}_{w \in W}$ of partial preorders is *complete* iff $W = \mathcal{M}$. A *faithful complete assignment* is a function that maps each belief state S to a complete collection $\{\leq_w^S\}_{w \in \mathcal{M}}$ of partial preorders s.t. each element $\leq_w^S$ is faithful to w.

Theorem 4 ([Fermé and Gonçalves, 2023])**.** *An update operator ◆ satisfies postulates (U◆1)-(U◆8) iff there is a faithful complete assignment that maps each belief state S to a complete collection $\{\leq_w^S\}_{w \in \mathcal{M}}$ of partial preorders s.t.* $[\mathtt{B}(S \blacklozenge \varphi)] = \bigcup_{w \in [\mathtt{B}_i(S)]} \min([\varphi], \leq_w^S)$[1].

The iterated postulates and their semantic characterization can be easily transferred to belief update. To distinguish the iterated postulates for belief update and revision, we use (CUi) for the iterated update version of (Ci) and (CRUi) for the semantic characterization of the postulate (CUi). For example, (CU1) is "If $\varphi \models \phi$, then $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \equiv \mathtt{B}(S \blacklozenge \varphi)$" and (CRU1) is "If $w_1, w_2 \in [\phi]$, then $w_1 \leq_w^S w_2$ iff $w_1 \leq_w^{S \blacklozenge \phi} w_2$."

Theorem 5 ([Fermé and Gonçalves, 2023])**.** *Let ◆ be an update operator satisfying (U◆1)-(U◆8). Then, ◆ satisfies (CUi) iff the operator and its corresponding faithful complete assignment satisfies (CRUi) for $1 \leq i \leq 4$.*

4 A New Approach to Iterated Belief Update

In this section, we propose a new approach to iterated belief update so as to cope with the first three shortcomings of Fermé and Gonçalves's approach mentioned in Introduction.

4.1 Modified Update Postulates on Darwiche and Pearl's Belief States

To overcome the first shortcoming, we directly adopt the original definition of belief states proposed in [Darwiche and Pearl, 1997]. Only (U◆4) and (U◆8) of modified KM postulates utilizes the associated update operator. We therefore redefine these two postulates in the following.

(U◆4*) If $S_1 = S_2$ and $\varphi \equiv \phi$, then $\mathtt{B}(S_1 \blacklozenge \varphi) \equiv \mathtt{B}(S_2 \blacklozenge \phi)$.

(U◆8*) $\mathtt{B}(S \blacklozenge \varphi) \equiv \bigvee_{w_i \in [\mathtt{B}(S)]} \mathtt{B}(S_i \blacklozenge \varphi)$ for some set of belief states $\{S_1, \cdots, S_n\}$ s.t. $\mathtt{B}(S_i) \equiv \mathtt{Form}(w_i)$.

Postulate (U◆4*) is identical with (R*4). If belief states S_1 and S_2 are identical, then they must lead to the same belief state when updating equivalent new information. Postulates (U8) and (U◆8*) both aim to achieve the distributive law of the update operator over disjunction although they look quite different at first glance. This is because the disjunction connective cannot directly apply to belief state. We illustrate the shift from (U8) to (U◆8*) in the following. Since $\mathcal{K} \equiv \bigvee_{w \in [\mathcal{K}]} \mathtt{Form}(w)$, we can obtain postulate (U8') equivalent to (U8) via iteratively applying (U8).

(U8') $\mathcal{K} \diamond \varphi \equiv \bigvee_{w \in [\mathcal{K}]} (\mathtt{Form}(w) \diamond \varphi)$.

According to postulate (U8'), the update of any belief set $\mathcal{K}$ reduces to the update of each $[\mathcal{K}]$-possible world. The update operator ◆ takes a belief state and a sentence as input. In order to describe the update of each possible world w_i of $\mathtt{B}(S)$, we choose a belief state S_i such that the model of its associated belief set $\mathtt{B}(S_i)$ is exactly w_i. Hence, postulate (U◆8*) coincides with (U8) and (U8').

4.2 Semantics of Belief Update over Belief States

We hereafter define the notion of faithful collective assignments that map each belief state S to a collection $\{\leq_w^S\}_{w \in [\mathtt{B}(S)]}$ of partial preorders. We remark that the set $\{\leq_w^S\}_{w \in [\mathtt{B}(S)]}$ is incomplete and contains no partial preorders centered at any possible world $w \notin [\mathtt{B}(S)]$. We therefore use faithful collective assignments instead faithful complete ones as the basis of semantics of belief update so as to solve the second shortcoming. For ease of presentation, we use $\leq^S$ for a collection $\{\leq_w^S\}_{w \in [\mathtt{B}(S)]}$ of partial preorders.

Definition 1. A *faithful collective assignment* is a function that maps each belief state S to a collection $\leq^S$ of partial preorders s.t. for each belief state S and each possible world $w \in [\mathtt{B}(S)]$, we have that

- each preorder $\leq_w^S$ is faithful to w; and
- there is a belief state S' s.t. $\mathtt{B}(S') \equiv \mathtt{Form}(w)$ and $\leq_w^S = \leq_w^{S'}$.

The first condition of faithful collective assignments is the same as the faithful property of faithful pointwise assignment in belief update. The second condition divides a belief state S into a collection of belief states such that the updated belief set $S \blacklozenge \varphi$ is equivalent to the disjunction of the belief set $S' \blacklozenge \varphi$. This ensures postulate (U◆8*) is satisfied.

We point out the difference among faithful assignments for (iterated) belief revision, faithful pointwise assignments for belief update, faithful collective assignments and faithful complete assignments for iterated belief update. A faithful assignment assigns each belief state to a single total preorder and a faithful pointwise assignment maps each possible world to a single partial preorder. In contrast, a faithful collective assignment associates each belief state with a collection of partial preorders. In addition, a faithful complete assignment maps each belief state with a complete set of partial preorders.

Theorem 3 can be extended to the modified KM postulates and faithful collective assignments.

Theorem 6. *An update operator ◆ satisfies postulates (U◆1)-(U◆3), (U◆4*), (U◆5)-(U◆7) and (U◆8*) iff there is a faithful collective assignment that maps each belief state S to*

[1]We remark that, in Theorem 3.4 of [Fermé and Gonçalves, 2023], the semantic definition of the update operator ◆ uses the most recent belief set $\mathtt{B}(S)$ rather than the initial one $\mathtt{B}_i(S)$. However, after carefully examining the proof of the representation theorem for iterated update postulates (Proof of Theorem 3.5 in [Fermé and Gonçalves, 2023]), we confirm that the correct definition is based on $\mathtt{B}_i(S)$, in line with the proof of the Master thesis [Gonçalves, 2015] that is the preliminary version of the paper [Fermé and Gonçalves, 2023].

a collection $\leq^S$ *of partial preorders s.t.* $[\mathtt{B}(S \blacklozenge \varphi)] = \bigcup_{w \in [\mathtt{B}(S)]} \min([\varphi], \leq^S_w)$.

It is observed from Theorem 6 that the semantics employs the most recent belief set $\mathtt{B}(S)$ rather than the initial one $\mathtt{B}_i(S)$, and hence addresses the third shortcoming.

4.3 Semantics Characterization of Iterated Belief Update on New Semantics

We offer the model-theoretic characterization of each DP iterated postulate based on the new semantics of belief update.

(CRU1*) For every $\mathcal{N} \subseteq [\phi]$, the following hold

Forth for every $w \in [\mathtt{B}(S)]$ and $w'' \in \min(\mathcal{N}, \leq^S_w)$, there is $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ s.t. $w'' \in \min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'})$.

Back for every $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ and $w'' \in \min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'})$, there is $w \in [\mathtt{B}(S)]$ s.t. $w'' \in \min(\mathcal{N}, \leq^S_w)$.

(CRU2*) For every $\mathcal{N} \subseteq [\neg\phi]$, the following hold

Forth for every $w \in [\mathtt{B}(S)]$ and $w'' \in \min(\mathcal{N}, \leq^S_w)$, there is $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ s.t. $w'' \in \min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'})$.

Back for every $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ and $w'' \in \min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'})$, there is $w \in [\mathtt{B}(S)]$ s.t. $w'' \in \min(\mathcal{N}, \leq^S_w)$.

(CRU3*) For every $\mathcal{N} \subseteq \mathcal{M}$ s.t. $\min(\mathcal{N}, \leq^S_w) \subseteq [\phi]$ holds for every $w \in [\mathtt{B}(S)]$, we have $\min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'}) \subseteq [\phi]$ for every $w' \in [\mathtt{B}(S \blacklozenge \phi)]$.

(CRU4*) For every $\mathcal{N} \subseteq \mathcal{M}$ s.t. $\min(\mathcal{N}, \leq^S_w) \cap [\phi] \neq \emptyset$ holds for some $w \in [\mathtt{B}(S)]$, we have there is $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ s.t. $\min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'}) \cap [\phi] \neq \emptyset$.

We hereafter illustrate the relationship between postulate (CU1) and its model-theoretic characterization. Postulate (CU1) can be split into two parts:

(CU1Forth) If $\varphi \models \phi$, then $\mathtt{B}(S \blacklozenge \varphi) \models \mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi)$.

(CU1Back) If $\varphi \models \phi$, then $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \models \mathtt{B}(S \blacklozenge \varphi)$.

Suppose that $[\varphi] = \mathcal{N}$ and $w'' \in [\mathtt{B}(S \blacklozenge \varphi)]$. Since $\varphi \models \phi$, $\mathcal{N} \subseteq [\phi]$. According to Postulate (CU1Forth), $w'' \in [\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi)]$. By the semantics defined in Theorem 6, $w'' \in [\mathtt{B}(S \blacklozenge \varphi)]$ iff $w'' \in \min(\mathcal{N}, \leq^S_w)$ for a world $w \in [\mathtt{B}(S)]$. Similarly, $w'' \in [\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi)]$ iff $w'' \in \min(\mathcal{N}, \leq^{S \blacklozenge \phi}_{w'})$ for a world $w' \in [\mathtt{B}(S \blacklozenge \phi)]$. Therefore, postulate (CU1Forth) corresponds to the forth condition of semantics (CRU1*). We can infer that postulate (CU1Back) corresponds to the back condition of (CRU1*) in a similar way. (CRU2*) acts similarly to (CRU1*) except that it focuses on the subset of possible worlds of $\neg\phi$. The semantics (CRU3*) (resp. (CRU4*)) states that for every set $\mathcal{N}$ of possible worlds, if all of the minimal elements of $\mathcal{N}$ w.r.t. the collection $\leq^S$ of preorders satisfy (resp. falsify) ϕ, then such property should be retained after updating by the new information ϕ.

Although each iterated update postulate is identical to the corresponding iterated revision one, the semantics for iterated postulates in belief update is distinct from that in belief revision. In contrast to belief revision, which is based on a single preorder, an update strategy is defined as a collection of preorders. We formalize the iterated update strategy in terms of the minimal elements of preorders. When both of the initial and updated collection of preorders contains a single preorder, the semantics for each iterated update postulate matches with that for the corresponding iterated revision one.

The theorem below provides the correspondence between each of the above postulates and its corresponding semantics, based on faithful collective assignments.

Theorem 7. *Let $\blacklozenge$ be an update operator satisfying (U$\blacklozenge$1)-(U$\blacklozenge$3), (U$\blacklozenge$4*), (U$\blacklozenge$5)-(U$\blacklozenge$7) and (U$\blacklozenge$8*). Then, $\blacklozenge$ satisfies (CUi) iff the operator and its corresponding faithful collective assignment satisfies (CRUi*) for* $1 \leq i \leq 4$.

5 Examples for DP Postulates in Iterated Belief Update

In this section, to avoid the final limitation of Fermé and Gonçalves's approach, we provide several examples to justify the rationale of each DP postulate in iterated belief update.

Example 2 (Postulate (CU1))**.** *Let us continue Example 1. Let $\mathcal{P} = \{b, c, t\}$ and $\mathcal{M} = \{w_0, \ldots, w_7\}$. The definition of each possible world is shown in Table 1, in which an occurrence of $\top$ (resp. $\bot$) indicates that the positive (resp. negative) literal of the corresponding proposition is in the possible world. For example, all of the three cells of the rows "book", "cup" and "toy" and the column "w_0" are $\bot$, meaning that they are all on the right zone of the table in the possible world w_0.*

We provide a KM-compatible update operator characterizing the update manner shown in Example 1 as follows. The initial belief state S that is associated with the belief set $\mathtt{B}(S) \equiv \neg b \wedge \neg c \wedge \neg t \equiv \mathtt{Form}(w_0)$, is assigned to the set $\leq^S$ of partial preorders with one element $\leq^S_{w_0}$.

- $w_0 <^S_{w_0} w_2, w_4 <^S_{w_0} w_1, w_3, w_5, w_6, w_7$.

Then, with the emergence of the new information $\phi = b \vee c$, the updated belief state $S \blacklozenge \phi$ picks the belief set $\mathtt{B}(S \blacklozenge \phi) \equiv (b \leftrightarrow \neg c) \wedge \neg t \equiv \mathtt{Form}(w_2, w_4)$. To be specific, the assigned collection $\leq^{S \blacklozenge \phi}$ of two preorders are as below.

- $w_2 <^{S \blacklozenge \phi}_{w_2} w_4 <^{S \blacklozenge \phi}_{w_2} w_0, w_1, w_3, w_5, w_6, w_7$.
- $w_4 <^{S \blacklozenge \phi}_{w_4} w_3 <^{S \blacklozenge \phi}_{w_4} w_0, w_1, w_2, w_5, w_6, w_7$.

Subsequently, as the robot is informed to place the cup on the left zone, a new information $\varphi = c$ occurs. The two worlds w_2 and w_3 are the minimal element satisfying φ w.r.t. $\leq^{S \blacklozenge \phi}_{w_2}$ and $\leq^{S \blacklozenge \phi}_{w_4}$, respectively. According to the semantics of the update operator (cf. Theorem 6), we have that $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \equiv \mathtt{Form}(w_2, w_3) \equiv \neg b \wedge c$, discarding unjustifiably the belief that the toy is on the right zone.

In contrast, from (CU1), we can deduce that $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \equiv \mathtt{B}(S \blacklozenge \varphi) \equiv \neg b \wedge c \wedge \neg t$, preserving the belief about the toy. □

Example 3 (Postulates (CU2))**.** *We use the previous example to justify postulate (CU2). Let us consider the following KM-compatible update operator. Initially, the belief set $\mathtt{B}(S) \equiv \neg b \wedge \neg c \wedge \neg t \equiv \mathtt{Form}(w_0)$. The associated partial preorder $\leq^S_{w_0}$ of the belief state S is:*

possible worlds	w_0	w_1	w_2	w_3	w_4	w_5	w_6	w_7
book	$\bot$	$\bot$	$\bot$	$\bot$	$\top$	$\top$	$\top$	$\top$
cup	$\bot$	$\bot$	$\top$	$\top$	$\bot$	$\bot$	$\top$	$\top$
toy	$\bot$	$\top$	$\bot$	$\top$	$\bot$	$\top$	$\bot$	$\top$

Table 1: The definition of possible worlds in Example 2.

- $w_0 <^S_{w_0} w_2, w_4 <^S_{w_0} w_1, w_3, w_5, w_6, w_7$.

Now, a robot is instructed to put the book and the cup on the opposite zone and put the toy on the right zone, described as $\phi = (b \leftrightarrow \neg c) \wedge \neg t$. From the semantics of the update operator, the updated belief set $\mathtt{B}(S \blacklozenge \phi)$ is $(b \leftrightarrow \neg c) \wedge \neg t \equiv \mathtt{Form}(w_2, w_4)$, that is, only one of the book and cup is on the left zone, and the toy must be on the right zone. Then, the updated belief state $S \blacklozenge \phi$ is assigned the collection $\leq^{S \blacklozenge \phi}$ of two partial preorders listed below.

- $w_2 <^{S \blacklozenge \phi}_{w_2} w_6, w_7 <^{S \blacklozenge \phi}_{w_2} w_0, w_1, w_3, w_4, w_5$.
- $w_4 <^{S \blacklozenge \phi}_{w_4} w_0, w_1 <^{S \blacklozenge \phi}_{w_4} w_2, w_3, w_5, w_6, w_7$.

Subsequently, the robot is informed to put the book and the cup on the same zone, which is described as new information $\varphi = b \leftrightarrow c$. Clearly, the minimal elements satisfying φ w.r.t. $\leq^{S \blacklozenge \phi}_{w_2}$ (resp. $\leq^{S \blacklozenge \phi}_{w_4}$) are possible worlds w_6 and w_7 (resp. w_0 and w_1). By the semantics of the update operator, the updated belief set $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi)$ becomes $b \leftrightarrow c \equiv \mathtt{Form}(w_0, w_1, w_6, w_7)$. However, as the agent has believed, after the update by ϕ, that the toy was on the right zone, it seems irrational to require the agent to give up this belief.

As $\varphi \models \neg\phi$, according to postulate (CU2), we can deduce that $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \equiv \mathtt{B}(S \blacklozenge \varphi) \equiv \neg b \wedge \neg c \wedge \neg t$. In this way, it preserves the belief about the toy. □

Example 4 (Postulate (CU3))**.** *Consider the scenario where an alarm is in a warehouse. When the alarm sounds, it means a fire breaks out. The proposition a and f denotes that "the alarm sounds" and "the warehouse catches fire", respectively. Let $\mathcal{P} = \{a, f\}$ and $\mathcal{M} = \{w_0, \ldots, w_3\}$. Each possible world is defined in Table 2. Initially, the belief set $\mathtt{B}(S)$ is $\neg a \wedge \neg f \equiv \mathtt{Form}(w_0)$, meaning that neither the alarm sounds nor the warehouse catches fire. The belief state S is associated with the following preorder $\leq^S_{w_0}$.*

- $w_0 <^S_{w_0} w_1 <^S_{w_0} w_3 <^S_{w_0} w_2$.

At some point, some workers smoke, causing a fire in the warehouse, which is described as the new information $\phi = f$. The updated belief set $\mathtt{B}(S \blacklozenge \phi)$ is $\neg a \wedge f \equiv \mathtt{Form}(w_1)$, that is, the alarm does not sound but the fire is in the warehouse. Then, the updated belief state $S \blacklozenge \phi$ is assigned to the set $\leq^{S \blacklozenge \phi}$ of preorders with one element $\leq^{S \blacklozenge \phi}_{w_1}$.

- $w_1 <^{S \blacklozenge \phi}_{w_1} w_0, w_2, w_3$.

Subsequently, the alarm sounds, which is described as new information $\varphi = a$. Obviously, the minimal elements satisfying φ are possible worlds w_2 and w_3. It follows from the semantics of the update operator that the updated belief set $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \equiv a \equiv \mathtt{Form}(w_2, w_3)$. However, after updating by ϕ, we believe that there is a fire in the warehouse. It seems unreasonable to give up this belief now when the alarm is also sounding.

possible worlds	w_0	w_1	w_2	w_3
alarm	$\bot$	$\bot$	$\top$	$\top$
fire	$\bot$	$\top$	$\bot$	$\top$

Table 2: The definition of possible worlds in Example 4.

As mentioned before, the alarm sounds only when there is a fire, hence $\mathtt{B}(S \blacklozenge a) \models f$ holds. We infer from postulate (CU3) that $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \models \phi$, allowing to preserve the belief $\phi = f$ in the belief set $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi)$. □

Example 5 (Postulate (CU4))**.** *Consider a scenario similar to Example 2. The initial belief set $\mathtt{B}(S)$ is $\neg c \wedge \neg t \equiv \mathtt{Form}(w_0, w_4)$. The belief state S is associated with the set $\leq^S$ of partial preorders that contains two elements $\leq^S_{w_0}$ and $\leq^S_{w_4}$.*

- $w_0 <^S_{w_0} w_4 <^S_{w_0} w_1, w_2, w_3, w_5, w_6, w_7$.
- $w_4 <^S_{w_4} w_0 <^S_{w_4} w_1, w_2, w_3, w_5, w_6, w_7$.

A robot first moves the book to the left zone and hence the new information $\phi = b$. Then, it is believed that the book is on the left zone, that is, the updated belief set $\mathtt{B}(S \blacklozenge \phi) \equiv b \wedge \neg c \wedge \neg t \equiv \mathtt{Form}(w_4)$. The associated partial preorder $\leq^{S \blacklozenge \phi}_{w_4}$ of the belief state $S \blacklozenge \phi$ is:

- $w_4 <^{S \blacklozenge \phi}_{w_4} w_0, w_2 <^{S \blacklozenge \phi}_{w_4} w_1, w_3, w_5, w_6, w_7$.

The robot is subsequently instructed to put the cup on the left zone, which is described as the new information $\varphi = c$. Accordingly, we get that the updated belief set $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi)$ becomes $\neg b \wedge c \wedge \neg t \equiv \mathtt{Form}(w_2)$. However, because $\mathtt{B}(S \blacklozenge \phi) \equiv b \wedge \neg c \wedge \neg t$, we have no reason to believe that the book is on the right zone.

Since $\mathtt{B}(S \blacklozenge \varphi) \not\models \neg\phi$, postulate (CU4) can be used to rule out the above irrational behaviour. Postulate (CU4) requires $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \not\models \neg\phi$, forbidding the unreasonable belief $\neg\phi = \neg b$ to be a consequence of the new belief set. □

6 (In)compatibility Results

In this section, based on the new semantics of belief update, we are going to analyze the compatibility between the iterated update and the modified KM postulates.

We hereafter provide a concrete update operator satisfying the modified KM postulates. Given a possible world w and a sentence ϕ s.t. $w \in [\phi]$, we divide the entire set $\mathcal{M}$ of possible worlds into three hierarchies: $\mathcal{H}^{w,\phi}_0 = \{w\}$, $\mathcal{H}^{w,\phi}_1 = [\phi] \setminus \{w\}$ and $\mathcal{H}^{w,\phi}_2 = [\neg\phi]$.

Definition 2. Let S be a belief state associated with a belief set $\mathtt{B}(S)$ and a collection $\leq^S$ of preorders over possible worlds. Let ϕ be a sentence. The operator $\blacklozenge_1$ is defined as:

1. $[\mathtt{B}(S \blacklozenge_1 \phi)] = \bigcup_{w \in [\mathtt{B}(S)]} \min([\phi], \leq^S_w)$;
2. for every $w \in [\mathtt{B}(S \blacklozenge_1 \phi)]$ and $w_1, w_2 \in \mathcal{M}$, $w_1 \leq^{S \blacklozenge_1 \phi}_w w_2$ iff $w_1 \in \mathcal{H}^{w,\phi}_i$, $w_2 \in \mathcal{H}^{w,\phi}_j$ and $i \leq j$.

The above two conditions impose the constraints on the updated belief set $\mathtt{B}(S \blacklozenge_1 \phi)$ and the subsequent update strategy $\leq^{S \blacklozenge_1 \phi}$, respectively. Clearly, condition (1) is the same as the semantics of the update operator. As a result, the operator $\blacklozenge_1$ satisfies postulates (U$\blacklozenge$1)-(U$\blacklozenge$3), (U$\blacklozenge$4*), (U$\blacklozenge$5)-(U$\blacklozenge$7) and (U$\blacklozenge$8*), hence being a KM-compatible update operator. Condition (2) assigns to every possible world w satisfying $\mathtt{B}(S \blacklozenge_1 \phi)$ a partial preorder $\leq_w^{S \blacklozenge_1 \phi}$. To be specific, the preorder $\leq_w^{S \blacklozenge_1 \phi}$ exactly characterizes a binary relation over possible worlds $\mathcal{M}$ divided into three hierarchies: $\mathcal{H}_0^{w,\phi}$, $\mathcal{H}_1^{w,\phi}$ and $\mathcal{H}_2^{w,\phi}$. That is, (1) the possible world w is the most preferable, followed by the ones satisfying the new information ϕ, and finally the remaining ones falsifying ϕ; and (2) every two possible worlds in $[\phi] \backslash \{w\}$ (resp. $[\neg\phi]$) are equally plausible.

Recalling the scenario in Example 5, we illustrate the specific update process of the operator $\blacklozenge_1$ as follows.

Example 6. *The initial belief state S is associated with the belief set $[\mathtt{B}(S)] = \{w_0, w_4\}$ and a collection $\leq^S$ of partial preorders as follows.*

- $w_0 <_{w_0}^S w_4 <_{w_0}^S w_1, w_2, w_3, w_5, w_6, w_7$.
- $w_4 <_{w_4}^S w_0 <_{w_4}^S w_1, w_2, w_3, w_5, w_6, w_7$.

The new information $\phi = b$ corresponds to the set of possible worlds $\{w_4, w_5, w_6, w_7\}$. Clearly, the minimal elements satisfying ϕ w.r.t. $\leq_{w_0}^S$ and $\leq_{w_4}^S$ are both the posssible world w_4. By the semantics of the update operator, we get that $\mathtt{B}(S \blacklozenge_1 \phi) \equiv \mathtt{Form}(w_4) \equiv b \wedge \neg c \wedge \neg t$. According to the definition of $\blacklozenge_1$, the updated belief state $S \blacklozenge_1 \phi$ would be assigned to a collection $\leq^{S \blacklozenge_1 \phi}$ of preorders.

- $w_4 <_{w_4}^{S \blacklozenge_1 \phi} w_5, w_6, w_7 <_{w_4}^{S \blacklozenge_1 \phi} w_0, w_1, w_2, w_3$

As seen above, the possible worlds of $\mathcal{M}$ are splitted into three hierarchies: (1) the most preferable ones $\mathcal{H}_0^{w_4,\phi} = \{w_4\}$ are exactly the current belief set, stating that only the book is on the left zone; (2) the second most preferable ones $\mathcal{H}_1^{w_4,\phi} = \{w_5, w_6, w_7\}$ are the other three possible worlds satisfying ϕ, stating that the book and at least one of the cup and the toy are on the left zone; (3) the least preferable ones $\mathcal{H}_2^{w_4,\phi} = \{w_0, w_1, w_2, w_3\}$ are the possible worlds falsifying ϕ, stating that the book is not on the left zone.

Finally, the newly acquired information $\varphi = c$ corresponds to the set of possible worlds $\{w_2, w_3, w_6, w_7\}$. The minimal elements satisfying φ w.r.t. $\leq_{w_4}^{S \blacklozenge_1 \phi}$ are the two possible worlds w_6 and w_7. Hence, the final updated belief set $\mathtt{B}((S \blacklozenge_1 \phi) \blacklozenge_1 \varphi) = \mathtt{Form}(w_6, w_7)$, which argues that both the book and cup are on the left zone. □

The following theorem confirms the compatibility results between two iterated postulates and the KM postulates.

Theorem 8. *The update operator $\blacklozenge_1$ satisfies postulates (CU3) and (CU4).*

In the seminal paper [Darwiche and Pearl, 1997], it is proved that all of the DP postulates for iterated revision are compatible with the AGM paradigm (a basic framework for belief revision). Unfortunately, we can not draw a similar conclusion in the context of belief update. The following theorem shows that neither postulate (CU1) nor (CU2) is compatible with the KM postulates.

Theorem 9. *There are a belief state S associated with a collection $\leq^S$ of partial preorders and a sentence φ s.t. no update operator $\blacklozenge$ satisfies (U$\blacklozenge$1)-(U$\blacklozenge$3), (U$\blacklozenge$4*), (U$\blacklozenge$5)-(U$\blacklozenge$7) and (U$\blacklozenge$8*) along with (CU1) (resp. (CU2)).*

Proof. We first consider postulate (CU1). We will construct a belief state S associated with a collection $\leq^S$ of partial preorders and a sentence ϕ s.t. there does not exist any collection $\leq^{S \blacklozenge \phi}$ of partial preorders s.t. $\leq^S$ and $\leq^{S \blacklozenge \phi}$ satisfy (CRU1*).

Let $\mathcal{P} = \{p_1, p_2, p_3\}$ and $\mathcal{M} = \{w_1, \ldots, w_8\}$[2]. Let S be a belief state with its associated belief set $[\mathtt{B}(S)] = \{w_1, w_2\}$ and its assigned partial preorders $\leq_{w_1}^S$ and $\leq_{w_2}^S$ as follows.

- $w_1 <_{w_1}^S w_2 <_{w_1}^S w_3 <_{w_1}^S w_4 <_{w_1}^S w_6 <_{w_1}^S w_5 <_{w_1}^S w_7 <_{w_1}^S w_8$.
- $w_2 <_{w_2}^S w_1 <_{w_2}^S w_3 <_{w_2}^S w_5 <_{w_2}^S w_6 <_{w_2}^S w_4 <_{w_2}^S w_7 <_{w_2}^S w_8$.

Let $\phi = \mathtt{Form}(w_3, w_4, w_5, w_6)$. By the semantics of the update operator, $[\mathtt{B}(S \blacklozenge \phi)] = \{w_3\}$. The collection $\leq^{S \blacklozenge \phi}$ contains only one partial preorder $\leq_{w_3}^{S \blacklozenge \phi}$. Assume that the forth part of (CRU1*) holds. That is, for every $\mathcal{N} \subseteq [\phi]$, $w \in [\mathtt{B}(S)]$ and $w'' \in \min(\mathcal{N}, \leq_w^S)$, there is $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ s.t. $w'' \in \min(\mathcal{N}, \leq_{w'}^{S \blacklozenge \phi})$. We will show that the back part of (CRU1*) does not hold, that is, there is $\mathcal{N}_3 \subseteq [\phi]$, $w' \in [\mathtt{B}(S \blacklozenge \phi)]$ and $w'' \in \min(\mathcal{N}_3, \leq_{w'}^{S \blacklozenge \phi})$ s.t. $w'' \notin \min(\mathcal{N}_3, \leq_w^S)$ holds for every $w \in [\mathtt{B}(S)]$.

Let $\mathcal{N}_0 = \{w_4, w_5\}$, $\mathcal{N}_1 = \{w_5, w_6\}$ and $\mathcal{N}_2 = \{w_4, w_6\}$. Since $\mathcal{N}_0 \subseteq [\phi]$, $w_4 \in \min(\mathcal{N}_0, \leq_{w_1}^S)$, by the assumption above, we have $w_4 \in \min(\mathcal{N}_0, \leq_{w_3}^{S \blacklozenge \phi})$. Similarly, as $w_5 \in \min(\mathcal{N}_0, \leq_{w_2}^S)$, we have $w_5 \in \min(\mathcal{N}_0, \leq_{w_3}^{S \blacklozenge \phi})$. It holds that $w_5 \not<_{w_3}^{S \blacklozenge \phi} w_4$ and $w_4 \not<_{w_3}^{S \blacklozenge \phi} w_5$. Therefore, w_4 and w_5 are either equivalently plausible ($w_4 \approx_{w_3}^{S \blacklozenge \phi} w_5$) or incomparable ($w_4 \not\approx_{w_3}^{S \blacklozenge \phi} w_5$ and $w_5 \not\approx_{w_3}^{S \blacklozenge \phi} w_4$) w.r.t. $\leq_{w_3}^{S \blacklozenge \phi}$. Similarly, since $\mathcal{N}_1 \subseteq [\phi]$, $w_5 \in \min(\mathcal{N}_1, \leq_{w_2}^S)$ and $w_6 \in \min(\mathcal{N}_1, \leq_{w_1}^S)$, we get that w_5 and w_6 are either equivalently plausible or incomparable w.r.t. $\leq_{w_3}^{S \blacklozenge \phi}$. In addition, as $\mathcal{N}_2 \subseteq [\phi]$, $w_4 \in \min(\mathcal{N}_2, \leq_{w_1}^S)$ and $w_6 \in \min(\mathcal{N}_2, \leq_{w_2}^S)$, we get that w_4 and w_6 are either equivalently plausible or incomparable w.r.t. $\leq_{w_3}^{S \blacklozenge \phi}$.

Let $\mathcal{N}_3 = \{w_4, w_5, w_6\}$. Clearly, $\mathcal{N}_3 \subseteq [\phi]$. It can be verified that $\min(\mathcal{N}_3, \leq_{w_3}^{S \blacklozenge \phi}) = \{w_4, w_5, w_6\}$, hence $w_6 \in \min(\mathcal{N}_3, \leq_{w_3}^{S \blacklozenge \phi})$. By the back part of (CRU1*), we get that $w_6 \in \min(\mathcal{N}_3, \leq_{w_1}^S)$ or $w_6 \in \min(\mathcal{N}_3, \leq_{w_2}^S)$. However, neither $w_6 \in \min(\mathcal{N}_3, \leq_{w_1}^S)$ nor $w_6 \in \min(\mathcal{N}_3, \leq_{w_2}^S)$, which is a contradiction.

The proof for (CU2) is similar to the above case except that we take into consideration $\phi = \mathtt{Form}(w_3)$. □

The incompatibility is due to the following reason. On the one hand, each of postulates (CU1) and (CU2) requires that $\mathtt{B}((S \blacklozenge \phi) \blacklozenge \varphi) \equiv \mathtt{B}(S \blacklozenge \varphi)$, that is, the two belief states $S \blacklozenge \varphi$ and $(S \blacklozenge \phi) \blacklozenge \varphi$ should associate with the same belief set given a specific new information φ. On the other hand, it is possible that $[\mathtt{B}(S)] \neq [\mathtt{B}(S \blacklozenge \phi)]$. In this case, their assigned collection $\leq^S$ and $\leq^{S \blacklozenge \phi}$ have different numbers of partial

[2]We remark that this proof holds no matter what the truth assignment on each possible world w_i is. Hence, we do not fix a truth assignment on each world.

preorders, resulting in two different belief sets which $S \diamond \varphi$ and $(S \diamond \phi) \diamond \varphi$ have.

7 Discussions

Examples 2-5 justify the rationality of the DP postulates in iterated update scenario. In the following, we provide a counterexamples to (CU2).

Example 7 (Postulate (CU2))**.** *We simplify Example 2, and the only items on the table are a book and a cup. Initially, the book and the cup can be found in any zone of the table, that is, the initial belief set* $\mathtt{B}(S)$ *is* $\top$*. We first instruct a robot to move the book and the cup to the left zone, which is described as the new information* $\phi = b \wedge c$*. The updated belief set* $\mathtt{B}(S \diamond \phi)$ *is therefore* $b \wedge c$*. The new instruction to place the cup on the right zone, which is the new information* $\varphi = \neg c$*, is issued to the robot. There is no reason to abandon the proposition* b *stating that the book is on the left zone, however, imposed by (CU2), we have* $\mathtt{B}((S \diamond \phi) \diamond \varphi) = \mathtt{B}(S \diamond \varphi) = \neg c$*.* □

Example 7 show that (CU2) leads to counterintuitive results which they aim to avoid, respectively. Postulate (CU2), however, are able to rule out inadequate update behaviors as shown in Example 3. We remark that although the DP postulates are controversial in iterated belief revision[3], they still are the cornerstone of iterated belief change. Many subsequent work on iterated belief change are based on the DP postulates. For example, belief contraction is a type of belief change that studies the problem of how to remove a certain information from a belief set. A revision operator can be obtained from a contraction operator via the Levi identity [Levi, 1978] and vice versa via the Harper identity [Harper, 1976]. Chopra *et al.* [2008] proposed four postulates for iterated belief contraction via slightly modifying the DP postulates. Booth and Chandler [2019] considered the four postulates as the benchmark of iterated belief contraction and use them to evaluate their proposed iterated contraction operators. Reasoning about actions, an important topic of knowledge representation and reasoning, studies the change of agents' beliefs due to the effect of action. Belief revision were incorporated into the situation calculus, a well-known framework of reasoning about actions [Shapiro *et al.*, 2011; Fang and Liu, 2013; Schwering *et al.*, 2017]. These works considered the DP postulates as key properties of iterated belief revision and verified the satisfaction of the DP postulates by their approaches so as to demonstrate the advantage of their approaches. Hence, the four postulates (CU1)-(CU4) we propose in this paper is a cornerstone of the subsequent research to iterated belief update.

[3] The criticism comes from convincing counterexamples in which the DP postulates cause counterintuitive results. Meyer [2001] and Stalnaker [2009] put forward counterexamples to postulate (C1). Several counterexamples to postulate (C2) are provided by Cantwell [1999] and Konieczny and Pino Pérez [2000]. In addition, postulate (C2) was discussed by Lehmann [1995], Delgrande *et al.* [2006] and Jin and Thielscher [2007]. Hansson [2016] offered two counterexamples to (C3) and (C4), respectively.

8 Conclusions

In this paper, we have investigated the iteration of belief update. Inspired by Darwiche and Pearl, we have presented a modification of the KM postulates framework over belief states. With the help of belief states, we have migrated the DP postulates for iterated revision to the belief update scenario, contributing to four iterated update postulates. Furthermore, we have offered the exact semantic characterizations based on partial preorders for each of the resulting postulates. At last, the (in)compatibility results between the iterated update postulates and the KM postulates are provided. We have showed that, unlike in revision, each of postulates (CU1) and (CU2) for iterated update is incompatible with the KM postulates.

Despite being the most influential approach to iterated belief revision, the DP postulates are still too liberal to rule out unintended revision operators. To strengthen the DP postulates, some additional postulates are proposed, for example, natural (Nat) postulate [Boutilier, 1996], lexicographic (Lex) postulate [Nayak *et al.*, 2003], and independence (Ind) postulate [Booth and Meyer, 2006; Jin and Thielscher, 2007]. Fermé and Gonçalves [2023] extended these strengthened iterated revision postulates to belief update scenario, yielding iterated update postulates (U-Nat), (U-Lex) and (U-Ind). Based on our proposed update framework, we have studied their semantic characterization and proved the representation theorem. The (in)compatibility results between them and the KM postulates are also analyzed. In particular, we have shown that both (U-Lex) and (U-Ind) are consistent with the KM postulates while (U-Nat) is inconsistent with the KM postulates. Due to space limitation, these results will be presented in a longer version of the paper.

Acknowledgements

This paper was supported by National Natural Science Foundation of China (Nos. 62276114, 62077028 and 62176103), the Guangdong Basic and Applied Basic Research Foundation (No. 2023B1515120064), the Key Laboratory of Smart Education of Guangdong Higher Education Institutes, Jinan University (No. 2022LSYS003), the Guangdong-Macao Advanced Intelligent Computing Joint Laboratory (No. 2020B1212030003), the Science and Technology Planning Project of Guangzhou (No. 2024A04J9896), and the Fundamental Research Funds for the Central Universities, Jinan University (No. 21623202).

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Booth and Chandler, 2019] Richard Booth and Jake Chandler. From iterated revision to iterated contraction: Extending the harper identity. *Artificial Intelligence*, 277:103171:1–103171:20, 2019.

[Booth and Meyer, 2006] Richard Booth and Thomas Meyer. Admissible and Restrained Revision. *Journal of Artificial Intelligence Research*, pages 127–151, 2006.

[Boutilier, 1996] Craig Boutilier. Iterated Revision and Minimal Change of Conditional Beliefs. *Journal of Philosophical Logic*, pages 263–305, 1996.

[Cantwell, 1999] John Cantwell. Some Logics of Iterated Belief Change. *Studia Logica*, 63(1):49–84, 1999.

[Chopra *et al.*, 2008] Samir Chopra, Aditya Ghose, Thomas Andreas Meyer, and Ka-Shu Wong. Iterated belief change and the recovery axiom. *Journal of Philosophical Logic*, 37(5):501–520, 2008.

[Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the Logic of Iterated Belief Revision. *Artificial Intelligence*, 89(1):1–29, 1997.

[Delgrande *et al.*, 2006] James P. Delgrande, Didier Dubois, and Jérôme Lang. Iterated Revision as Prioritized Merging. In *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR-2016)*, pages 210–220, 2006.

[Fang and Liu, 2013] Liangda Fang and Yongmei Liu. Multiagent Knowledge and Belief Change in the Situation Calculus. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence (AAAI-2013)*, pages 304–312, 2013.

[Fermé and Gonçalves, 2023] Eduardo Fermé and Sara Gonçalves. On the logic of theory change iteration of KM-update. *International Journal of Approximate Reasoning*, 161:108984, 2023.

[Gonçalves, 2015] Sara Gonçalves. On the Iteration of Katsuno and Mendelzon Update. Master's thesis, University of Madeira, 2015.

[Hansson, 2016] Sven Ove Hansson. Iterated Descriptor Revision and the Logic of Ramsey Test Conditionals. *Journal of Philosophical Logic*, 45(4):429–450, 2016.

[Harper, 1976] William L. Harper. Rational Conceptual Change. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 2, pages 462–494, 1976.

[Jin and Thielscher, 2007] Yi Jin and Michael Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.

[Katsuno and Mendelzon, 1991a] Hirofumi Katsuno and Alberto O. Mendelzon. On the Difference between Updating a Knowledge Base and Revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR-1991)*, pages 387–394, 1991.

[Katsuno and Mendelzon, 1991b] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence*, 52(3):263–294, 1991.

[Konieczny and Pino Pérez, 2000] Sébastien Konieczny and Ramón Pino Pérez. A framework for iterated revision. *Journal of Applied Non-Classical Logics*, 10:339–367, 2000.

[Lehmann, 1995] Daniel Lehmann. Belief Revision, Revised. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-1995)*, pages 1534–1540, 1995.

[Levi, 1978] Isaac Levi. *Subjunctives, Dispositions and Chances*, pages 303–335. Springer Netherlands, 1978.

[Meyer, 2001] Thomas Meyer. Basic Infobase Change. *Studia Logica*, 67(2):215–242, 2001.

[Nayak *et al.*, 2003] Abhaya C. Nayak, Maurice Pagnucco, and Pavlos Peppas. Dynamic Belief Revision Operators. *Artificial Intelligence*, 146(2):193–228, 2003.

[Schwering *et al.*, 2017] Christoph Schwering, Gerhard Lakemeyer, and Maurice Pagnucco. Belief revision and projection in the epistemic situation calculus. *Artificial Intelligence*, 251:62–97, 2017.

[Shapiro *et al.*, 2011] Steven Shapiro, Maurice Pagnucco, Yves Lespérance, and Hector J. Levesque. Iterated belief change in the situation calculus. *Artificial Intelligence*, 175(1):165–192, 2011.

[Stalnaker, 2009] Robert Stalnaker. Iterated Belief Revision. *Erkenntnis*, pages 189–209, 2009.