Belief Update without Compactness in Non-finitary Languages

Jandson S. Ribeiro[1,2] , **Abhaya Nayak**[1] and **Renata Wassermann**[2]

[1] Macquarie University, Australia
[2] University of São Paulo, Brazil

{jandson, renata}@ime.usp.br, abhaya.nayak@mq.edu.au

Abstract

The main paradigms of belief change require the background logic to be Tarskian and finitary. We look at belief update when the underlying logic is not necessarily finitary. We show that in this case the classical construction for KM update does not capture all the rationality postulates for KM belief update. Indeed, this construction, being fully characterised by a subset of the KM update postulates, is weaker. We explore the reason behind this, and subsequently provide an alternative constructive accounts of belief update which is characterised by the full set of KM postulates in this more general framework.

1 Introduction

Belief change studies the dynamics of a body of knowledge in response to a new piece of information. The two most influential paradigms of belief change are the AGM paradigm of belief revision [Alchourrón *et al.*, 1985] and the KM paradigm of belief update [Katsuno and Mendelzon, 1992], the latter formally developing an idea explored in [Winslett, 1988]. Broadly speaking, the former concerns the process of changing beliefs in the context of *static* worlds, whereas the latter considers knowledge change in *dynamic* worlds.

In both approaches, belief change may be represented as a function f that takes an epistemic state into another one, given a new piece of information. These functions are governed by rationality postulates whose purpose is to prescribe good set of behaviours for belief change.

In both these paradigms, the epistemic state of an agent is represented as a formula, or a set of formulae, in a propositional language, and the background logic is assumed to be, among others, Tarskian and compact. Compactness states that any formula entailed by a set of formulae X should be entailed by a finite subset X' of X. In this work we explore the KM update with respect to a class of logics that have an infinite number of not logically equivalent formulae but are not necessarily compact. We call such logics (or languages) *non-finitary* since no assumption is being made as to if the alphabet in question is finite or not.[1]

Efforts to apply AGM paradigm to non-classical logics were studied in [Restall and Slaney, 1995] and [Wassermann, 2011]. In the context of the AGM contraction (a form of AGM belief change in which an agent relinquishes a piece of information) it was shown in [Ribeiro *et al.*, 2013] that it is not possible to define contraction functions satisfying the AGM contraction postulates in intuitionistic logics. The reason for this is that these logics are not closed under classical negation.

Recently, complete accounts of AGM contraction for non-compact logics have been provided in [Ribeiro *et al.*, 2018], subsequently extended in [Ribeiro *et al.*, 2019]. We study KM update while dispensing with the compactness requirement, therefore *finitaryness*. A number of interesting logics, from both perspective of AI and computing science are not compact. Temporal logics, such as CTL and LTL [Gabbay *et al.*, 1994], which are used in planning and formal methods. Work such as [Zhang and Ding, 2008] proposes to use KM update to repair CTL models that do not comply with the temporal requirements in the context of formal specification and verification of systems. One of the limitation of this work is that it considers only knowledge expressed via a single CTL model.

The question we have tried to answer is: What happens to the standard KM update constructions when compactness is discarded? Interestingly, as in the case of AGM contraction [Ribeiro *et al.*, 2018], the correspondence between the standard KM constructions and the KM rationality postulates breaks down. The source of this problem turns out to be not compactness itself but the inherent *non-finitaryness*. Accordingly, instead of only dispensing with compactness, we address the issue of belief update for non-finitary logics.

Our main contributions are:

(1) a new class of KM update functions that are fully characterised by the KM postulates, in non-finitary logics;

(2) we keep the classical KM update functions, and identify a set of rationality postulate that characterise them.

In Section 2 we briefly outline the notation employed in this paper, and introduce some useful concepts that we will need later. In Section 3, we briefly review the KM rationality postulates and the classical KM update functions. We also show the consequences that discarding compactness has over

---

[1]Propositional modal logic **K** is non-finitary and yet compact, whereas temporal logics like **LTL** and **CTL** are neither finitary nor compact.

these functions. In Section 4, we devise a new class of functions characterised by the set of basic KM postulates. In particular, we extend these constructions to capture the supplementary postulates and provide a fully KM-rational account. Then in Section 5 we show that in non-finitary logics the classical KM update functions become weaker in the sense that they are characterised by a proper subset of the KM postulates. In that same section we consider the issue of total relation in the KM update relational functions, and show how to dispense with them. Finally, we recap the main results obtained and discuss some future directions for our work.

2 Notation and Technical Background

We consider a logic as a pair $\langle L, Cn \rangle$, where L is a language and $Cn: 2^L \to 2^L$ is a logical consequence operator that maps a set of formulae to the set of all formulae that can be inferred from it. For readability, for any formula φ, the set $Cn(\{\varphi\})$ will often simply be written as $Cn(\varphi)$. We limit ourselves to logics that are Tarskian, that is, logics whose consequence operator satisfies the following three properties:

1. **(Monotonicity)**: if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$;

2. **(Idempotence)**: $Cn(A) = Cn(Cn(A))$;

3. **(Inclusion)**: $A \subseteq Cn(A)$;

The KM paradigm was initially proposed for Propositional Logic, whereby the consequence operator Cn satisfies the following properties:

- **(deduction)**: $\varphi \in Cn(A \cup \{\psi\})$ iff $\psi \to \varphi \in Cn(A)$;

- **(supraclassicality)**: if φ is a logical consequence of A in classical propositional logic, then $\varphi \in Cn(A)$;

- **(compactness)**: if $\varphi \in Cn(A)$ then there is a finite subset A' of A such that $\varphi \in Cn(A')$.

Given a logic $\langle L, Cn \rangle$ and a set of structures $\mathcal{I}$, an interpretation or a model of a formula φ (respectively a set of formulae Φ) of L is an element of $\mathcal{I}$ that satisfies φ (or Φ). We will say that a logic $\langle L, Cn \rangle$ is closed under classical negation iff the language L is closed under the negation operator $\neg$ such that, for each formula $\varphi \in L$, $Cn(\varphi) \cap Cn(\neg \varphi) = Cn(\emptyset)$, and $Cn(\{\varphi, \neg \varphi\}) = L$. In other words, the negation is interpreted classically. Analogously, a logic is closed under the disjunction if the associated language is closed under such a connective (classically interpreted, that is, if $\varphi \in Cn(X)$ then $\varphi \lor \psi \in Cn(X)$, for every $\psi \in L$ and $X \subseteq L$).

A pre-order is a binary relation $\leqslant$ that satisfies reflexivity and transitivity:

(reflexivity) $X \leqslant X$;

(transitivity) if $X \leqslant Y$ and $Y \leqslant Z$ then $X \leqslant Z$.

The maximal elements of a set A with respect to a binary relation $\leqslant$ is given by $max_{\leqslant}(A) = \{x \in A \mid \nexists y \in A, x \leqslant y\}$. Analogously, the minimal elements are given by $Min_{\leqslant}(A) = \{x \in A \mid \nexists y \in A, y \leqslant x\}$. A belief set, or simply a theory, K is a set of formulae logically closed under a consequence relation Cn, that is, $K = Cn(K)$. Given a belief set K, the collection of all the interpretations that satisfy K is given by $[K]$. For convenience, we write $[\varphi]$ to denote the interpretations that satisfy the theory $Cn(\varphi)$. Given

a set of models A, $Th(A) = \{\varphi \mid \forall M \in A, M \models \varphi\}$ is the theory of the formulae satisfied by all models in A. A theory K is complete if and only if, for every sentence φ, either $\varphi \in K$ or $\neg \varphi \in K$, and we use T_L to denote all such complete theories that are consistent. The decomposition of a theory K in terms of its consistent complete theories is given by $\omega(K) = \{K' \in T_L \mid K \subseteq K'\}$. Conveniently, we simply write $\omega(\varphi)$ to refer to $\omega(Cn(\varphi))$. A function that maps each pair of theory and formula (K, φ) to another theory K' is a belief change function.

3 Belief Update

In the classic KM approach to belief update, epistemic states are represented as a formula from a propositional language. When dealing with non-finitary logics, however, representing epistemic states in this way may be inconvenient, mainly because in these logics not every belief set can be represented via a single formula or a finite set of formulae. For this reason, we will translate the original KM update functions and rationality postulates in terms of epistemic states as belief sets. We start with the Classical Update (CUP) functions[2]:

$$\text{CUP}: \quad K \diamond \varphi = Th\big(\bigcup_{M \in [K]} Max_{\leqslant_M}([\varphi]).$$

Given a theory K and a formula φ, a CUP operator $\diamond$ behaves as follows. For each model M of K, the operator $\diamond$ chooses from the models of φ those closest to M. The notion of distance between models is usually given by a pre-order $\leqslant_M$. A model M_1 is closer to M than a model M_2, if $M_2 \leqslant_M M_1$. These selected models are then assembled, and the theory of these models corresponds to $K \diamond \varphi$.

In [Katsuno and Mendelzon, 1992] it is shown that CUP functions are characterised precisely by a set of eight postulates U1 to U8. We translate this postulates in terms of epistemic states as belief sets which correspond respectively to the postulates (K$\diamond$2) to (K$\diamond$9) below.

(K$\diamond$1) $K \diamond \varphi$ is a theory.
(K$\diamond$2) $\varphi \in K \diamond \varphi$.
(K$\diamond$3) If $\varphi \in K$, then $K \diamond \varphi = K$.
(K$\diamond$4) $K \diamond \varphi$ is consistent, if both K and φ are consistent.
(K$\diamond$5) If $Cn(\varphi) = Cn(\psi)$, then $K \diamond \varphi = K \diamond \psi$.
(K$\diamond$6) $K \diamond (\varphi \land \psi) \subseteq K \diamond \varphi + \psi$.
(K$\diamond$7) If $\psi \in K \diamond \varphi$ and $\varphi \in K \diamond \psi$, then
$K \diamond \varphi = K \diamond \psi$.
(K$\diamond$8) If K is complete, then
$K \diamond \varphi \lor \psi \subseteq (K \diamond \varphi) \cup (K \diamond \psi)$.
(K$\diamond$9) $(K \cap K') \diamond \varphi = (K \diamond \varphi) \cap (K' \diamond \varphi)$.

As we require an epistemic state to be a belief set, the postulate (K$\diamond$1) is added here to ensure that every update function goes from a belief set to another one. Discussion of the rationality behind these postulates can be found in [Katsuno and Mendelzon, 1992].

---

[2]Originally, a CUP function is defined over minimal elements, here for convenience we will resort to maximal elements, as we will need to deal with infinite chains, and maximal elements make the reasoning task easier.
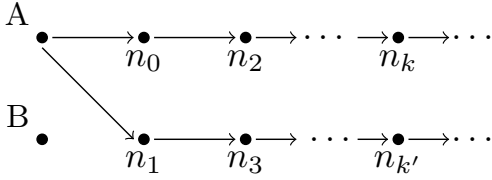
Figure 1: A preorder for the function of Example 1.

Theorem 1. *A belief change function satisfies the KM update postulates iff it is a CUP function.*

We stress that this representation between KM postulates and CUP functions was shown for propositional logic. This connection, however, fails for non-finitary logics. For instance, for non-finitary logics CUP functions do not satisfy postulate (K⋄7), as illustrated in Example 1.

Example 1. *Let ⋄ be a CUP function over the pre-order ⩽ depicted in Figure 1. An edge $x \rightarrow y$ means that $x \leqslant y$. Reflexive and transitive edges are omitted in the figure for convenience. Moreover, let K be a complete consistent theory, and φ and ψ be two formulae such that $[\varphi] = \{A, B, n_0, n_2, \dots\}$, $[\psi] = \{A, B, n_1, n_3, \dots\}$, and $[\varphi \wedge \psi] = \{A, B\}$. According to the relation $\leqslant$, $Max_\leqslant(\varphi) = \{B\}, Max_\leqslant(\psi) = \{B\}$ and $Max_\leqslant(\varphi \wedge \psi) = \{A, B\}$. Therefore,*

$$
\begin{aligned}
K \diamond \varphi &= Th\big(Max_\leqslant([\varphi])\big) &&= Th(\{B\}) \\
K \diamond \psi &= Th\big(Max_\leqslant([\psi])\big) &&= Th(\{B\}) \\
K \diamond \varphi \wedge \psi &= Th\big(Max_\leqslant([\varphi \wedge \psi])\big) &&= Th(\{A, B\}).
\end{aligned}
$$

Though $\leqslant$ is a pre-order, the update function $\diamond$ does not satisfies (K⋄7). The reason for this is that according to (K⋄7) it is required that $K \diamond \varphi = K \diamond \varphi \wedge \psi$, however this does not occur in this example. To see this, note that $\varphi \wedge \psi \in K \diamond \varphi$, as B is model of both φ and ψ. Clearly, $\varphi \in K \diamond \varphi \wedge \psi$. However, $K \diamond \varphi \neq K \diamond \varphi \wedge \psi$ which violates (K⋄7).

The reason why the function $\diamond$ of Example 1 does not satisfy (K⋄7) is related to the characteristic of non-finitary logics having an infinite number of models. In this case, a pre-order may induce an infinite chain in which it is not possible to determine a maximal (resp. minimal) element. To see this, let us consider what happens to the $\diamond$ function of Example 1 if there were no infinity chain. For this, we can easily slightly change that example to consider a finite subset of elements: n_0 to n_3 besides A and B. This gives us that $[\varphi] = \{A, B, n_0, n_2\}$ and $[\psi] = \{A, B, n_1, n_3\}$. Thus, $Max_\leqslant(\varphi) = \{B, n_2\}$ and $Max_\leqslant(\psi) = \{B, n_3\}$. As n_2 is a model of φ but not a model of ψ, we do not have anymore that $\varphi \wedge \psi \in K \diamond \varphi$. Therefore, in this case, (K⋄7) is trivially satisfied.

Example 1 shows us that the connection between the postulates for KM update and CUP functions breaks down in non-finitary logics. It is natural, then, to pose the following question: how can one perform belief update in non-compact logics? Two natural answers are:

1. one can devise a new class of functions characterised by the KM update postulates in non-finitary logics;

2. one can identify which set of postulates follow from CUP functions in non-finitary logics.

We pursue both these paths. We will propose a new class of functions that satisfy all the KM update postulates, and the results we obtain will prove useful to show that CUP functions, for non-finitary logics, are characterised by a proper subset of the KM update postulates.

4 Reconstructing Rational Update Functions

In this section, we provide complete new accounts of rational KM update functions. For convenience, we split the KM postulates in two categories: basic postulates which comprise postulates (K⋄1) to (K⋄5) and (K⋄9); and supplementary postulates which consists of postulates (K⋄6) to (K⋄8). We will call a belief change function that satisfies the basic KM postulates a belief update function. The main difference between our new construction and the CUP functions relies on the underlying preference relations. Unlike the CUP functions, transitivity is not a requirement, instead we identify a new property called *quasi-reflection* which fills up the gap between the supplementary postulates in non-finitary logics.

4.1 KM Basic Rationality

The functions we will devise here will operate over complete consistent theories rather than models, following the approach in [Ribeiro *et al.*, 2018]. This will help us avoid technical issues. For instance, in modal logics, unlike in Propositional logic, two different models (but not consistent complete theories) may satisfy exactly the same set of formulae.

In non-finitary logics, pre-orders are not strong enough to properly indicate which are the best elements to be used in order to update a theory K by a formula φ, as illustrated in Example 1. This means we will need another way to define the 'closeness' criterion between consistent complete theories. We will assume that, instead of a pre-order, an agent has an *appointee function* that judges which complete theories of a formula φ are closest to a complete theory K. This appointee function is used to perform the update of K by φ. An appointee considers only consistent complete theories. We recall from Section 2 that the set of all complete consistent theories comprises T_L.

Definition 1. *An appointee is a function $\delta : T_L \times L \rightarrow 2^{T_L}$ that maps each consistent complete theory and formula φ to a set of consistent complete theories subject to the following conditions:*

(D1) $\delta(S, \varphi) \neq \emptyset$, if S and φ are consistent;

(D2) $\delta(S, \varphi) \subseteq \omega(\varphi)$;

(D3) $\delta(S, \varphi) = \{S\}$, if $\varphi \in S$;

(D4) $\delta(S, \varphi) = \delta(S, \psi)$, if $Cn(\varphi) = Cn(\psi)$.

The purpose of an appointee is to select the best complete consistent theories that satisfy a formula φ (condition D2). This criterion is given locally, and it depends on the fixed complete theory S as background. An appointee is compelled to choose at least one consistent complete theory that contains φ, as long as φ is consistent (condition D1). If a formula φ belongs to a complete theory S, then there is no need to go after another theory, that is, the best theory for φ is S itself (condition D3). Finally, condition (D4) simply determines that an appointee is syntax independent.

The main purpose of defining an appointee function is that it captures the concept of distance between complete consistent theories in a general way. With an appointee in hand, it will become easier to explore which conditions we need to capture all the KM update postulates. We start by defining a new class of update functions, the *splinter functions* which are similar in spirit to CUP functions:

Definition 2. *Given an appointee δ, a splinter is a function $\diamond_\delta$ such that*

1. if K and φ are consistent, then

$$K \diamond_\delta \varphi = \bigcap_{S \in \omega(K)} \left(\bigcap \delta(S, \varphi) \right);$$

2. otherwise, $K \diamond_\delta \varphi = Cn(\bot)$.

We recall from Section 2 that $\omega(K)$ corresponds to the set of all complete consistent theories that contain a theory K.

To understand how a splinter works, we will first need a property that follows from the converse of postulate (K$\diamond$4). To see this, first notice that φ cannot be inconsistent since it is required to be in $K \diamond \varphi$. On the other hand, if K were inconsistent, then we would have $\varphi \in K$, rendering $K \diamond \varphi$ inconsistent. So the only options remaining is that both K and φ are consistent. Conversely, if either K or φ are inconsistent, we get that $K \diamond \varphi$ is inconsistent. This case is captured by the second condition of splinter function definition. For the other cases, a theory K is split into its decomposition of complete consistent theories. An appointee then is used to select for each of these consistent theories the best complete theories from φ. Thereafter they are all intersected and it results in $K \diamond \varphi$.

Observation 1. *Let $\diamond$ be a belief change function that satisfies (K$\diamond$1), (K$\diamond$3) and (K$\diamond$4). Then, $K \diamond \varphi = Cn(\bot)$ iff either φ or K is inconsistent.*

Example 2. *Let $K = Cn(p)$ be a theory such that its decomposition has only two complete consistent theories: K_1 and K_2. Let an appointee δ be as follows: $\delta(K_1, \neg p) = \{Cn(\neg p, q)\}$ and $\delta(K_2, \neg p) = \{Cn(\neg p, \neg q)\}$. Thus,*

$$\begin{aligned} K \diamond_\delta \neg p &= \delta(K_1, \neg p) \cap \delta(K_2, \neg p) \\ &= Cn(\neg p, q) \cap Cn(\neg p, \neg q) \\ &= Cn(\neg p). \end{aligned}$$

A splinter function behaves similar to a CUP function, with the difference of resorting to an appointee instead of preorders to choose elements for the update. It is easy to check that every splinter function satisfies the basic KM postulates.

Theorem 2. *Every splinter function $\diamond_\delta$ satisfies (K$\diamond$1) to (K$\diamond$5) and (K$\diamond$9).*

Not surprisingly, splinter functions not only satisfy the basic KM update postulates, but every update function is a splinter function.

Another interesting property of update functions is that they are monotonic, that is, if $K \subseteq K'$ then $K \diamond \varphi \subseteq K' \diamond \varphi$. Update functions are monotonic due to postulate (K$\diamond$9).

Proposition 1. *Every belief change operation $\diamond$ that satisfies (K$\diamond$9) is monotonic.*

The last piece of the puzzle towards our representation result requires us to look at splinter functions from another perspective. We begin with a special case: *given a complete consistent theory K and a consistent formula φ, how does $K \diamond_\delta \varphi$ look like?* As K is a complete consistent theory, the only complete theory that contains K is K itself. Therefore, $\omega(K) = \{K\}$, which gives us a simpler version of splinter:

$$K \diamond \varphi = \bigcap \delta(K, \varphi).$$

Now let us proceed for a more general case. Imagine that K is a consistent theory, but not complete. Let S be one of the complete consistent theories from the decomposition of K, then have that $S \diamond \varphi = \bigcap \delta(S, \varphi)$. Thus, we can easily see that the update of K by φ consists in taking each $S \diamond \varphi$, for $S \in \omega(K)$, and intersecting all of them. This leads us to the following observation.

Observation 2. *Given a splinter $\diamond_\delta$, a theory K and a formula φ, if both K and φ are consistent then*

$$K \diamond_\delta \varphi = \bigcap_{S \in \omega(K)} S \diamond_\delta \varphi.$$

Observation 2 will be useful to prove the Theorem 3 below. For this, we will also need the following two lemmas from [Ribeiro *et al.*, 2018].

Lemma 1. $[Cn(K \cup K')] = [Cn(K)] \cap [Cn(K')]$, *for all theories K and K'.*

Lemma 2. $Cn(K \cap K') = Th([K] \cup [K'])$, *for all theories K and K'.*

Theorem 3. *Every update function is a splinter.*

Our first representation result follows from Theorems 2 and 3 which jointly show that splinter functions and the basic KM postulates are equivalent.

4.2 Reviewing the Supplementary Postulates

Although splinter functions satisfy the basic KM postulates, they are not strong enough to capture the supplementary postulates. In this section, we introduce a new class of update functions that satisfy the supplementary postulates. The main idea is to make the appointees to resort to binary relations in order to pick the consistent complete theories.

An appointee δ will assign, for each consistent complete theory K, a binary relation $\leqslant_K$. The appointee chooses the maximal elements within $\omega(\varphi)$ modulo $\leqslant_K$. These relations will need to satisfy some requirements in order to yield splinter functions that satisfies the supplementary postulates. The first one is

> **Maximal Cut:** for every consistent formula φ, $\omega(\varphi)$ has a maximal element w.r.t $\leqslant$.

Maximal Cut[3] was introduced in [Ribeiro *et al.*, 2018] and its purpose is to guarantee that for every formula φ, as long as it is consistent, the epistemic agent chooses at least one

---

[3]In [Ribeiro *et al.*, 2018], maximal cut is defined over the complements of a formula φ, since the focus is contraction. Here, we conveniently translate maximal cut in terms of the complete consistent theories of φ.

$$C = Cn(\neg p, q) \longrightarrow D = Cn(\neg p, \neg q)$$
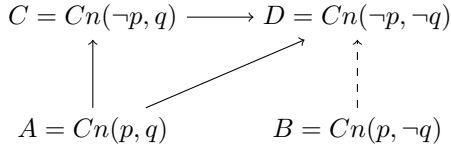
$$A = Cn(p, q) \qquad B = Cn(p, \neg q)$$

Figure 2: A contra-headed relation.

consistent complete theory of φ. This ensures that φ will be successfully incorporated through the update process. A binary relation that satisfies *maximal-cut* is called a *contra-headed* relation[4].

Appointee functions that rely strictly on *contra-headed* relations will be called relational appointees.

Definition 3. *An appointee μ is relational iff for every consistent complete theory K, there is a contra-headed relation $\leqslant_K$ such that $\mu(K, \varphi) = max_{\leqslant_K}(\omega(\varphi))$, for every consistent formula φ.*

The new update functions we introduce are the *royal splinter functions* which consist of splinters functions founded on relational appointees.

Definition 4. *A splinter $\diamond_\mu$ is royal iff its appointee μ is relational.*

Example 3. *Consider the complete theories depicted in Figure 2 and the respective contra-headed relation $\leqslant$ represented by solid arrows. Let μ be a relational appointee such that for the theory A it considers the contra-headed $\leqslant$. Thus, $\mu(A, \neg q) = \{B, D\}$ which implies that*

$$A \diamond_\mu \neg q = B \cap D = Cn(\neg q).$$

At this point, we investigate what we gain by restricting ourselves to royal splinters. The first benefit is that the *contra-headed* relations guarantee satisfaction of postulates (K◇6) and (K◇8) which is shown on Theorem 4 below. Lemma 3 will be of immediate assistance through this section.

Lemma 3. *Let $\diamond_\delta$ be a fully KM rational splinter, and K a complete consistent theory,*

(i) $\delta(K, \varphi) \cap \omega(\psi) \subseteq \delta(K, \varphi \wedge \psi)$);

(ii) $\delta(K, \varphi) \cap \delta(K, \psi) \subseteq \delta(K, \varphi \vee \psi)$.

The conditions (i) and (ii) from Lemma 3 concern properties that follow from fully rational splinters, precisely the appointees of these functions, and are related respectively to postulates (K◇6) and (K◇8).

Theorem 4. *Every royal splinter satisfies (K◇6) and (K◇8).*

We impose one more restriction over the relations considered by a relational appointee in order to capture (K◇7):

Quasi-reflection:[5] if $A \not\leqslant B$ and $B \not\leqslant A$ but $A \leqslant C$ then either (i) $B \leqslant C$, or (ii) if $C \leqslant C'$, then $B \leqslant C'$.

---

[4]This terminology is borrowed from [Ribeiro *et al.*, 2018].

[5]This concept may be intuitively motivated by considering what we would call *the Principle of Graded Indifference* which can be used to represent *equi-preference* at the strongest level, and at weaker levels it represents *ignorance* coupled with certain aspects of equi-preference. This is a work in progress.

To explain quasi-reflection, we will consider conditions (i) and (ii) separately. The main intuition behind (i) is that if an agent has no preference between two elements A and B, then all elements preferable to A should be preferred to B and vice versa. The second condition elevates this mimicking process by one step. For instance, consider the relation $\leqslant$ depicted in Figure 2 by both solid and dashed arrows. In that relation, there is no preference between A and B. Note that $A \leqslant C$ and $C \leqslant D$. However, C is not preferable to B, as condition (i) demands. In this case, instead of forcing $B \leqslant C$, condition (ii) pushes B to mimic one of the elements immediately above A, in this case mimic the preferences over C. Thus, it makes $B \leqslant D$ (depicted as a dashed arrow).

It turns out that quasi-reflection is all we need to capture (K◇7). A relational appointee that considers only *contra-headed* relations that satisfy quasi-refection will be called a *quasi-reflected appointee*. To see how *quasi-reflection* captures (K◇7), we look at a special case of (K◇7) over complete consistent theories.

Lemma 4. *Let $\diamond_\mu$ be a royal splinter, such that μ is quasi-reflected. Given a complete consistent theory K, if $\varphi \in K \diamond_\mu \psi$ and $\psi \in K \diamond_\mu \varphi$, then $K \diamond_\mu \varphi = K \diamond_\mu \psi$.*

The tricky part is to show that quasi-reflection not only captures (K◇7) in complete theories, but also in every theory. We make use of the monotonicity property of update functions to show, via lemma above, that (K◇7) spreads from complete consistent theories to all theories.

Theorem 5. *Given a royal splinter $\diamond_\mu$, if μ is quasi-reflected then $\diamond_\mu$ satisfies (K◇7).*

We get the following corollary from Theorems 4 and 5.

Corollary 1. *A royal splinter whose appointee is quasi-reflected is fully KM rational.*

Corollary 1 provides the first piece of our representation result. The next step is to prove that every fully KM update function is a royal splinter function whose appointee is *quasi-reflected*, in short, a royal reflected splinter.

We begin by showing that fully KM update functions are royal. From Theorem 2 we have that every KM update function is a splinter, which means that all of them have an appointee. All we need to show now is that appointees from fully KM splinters are relational.

This is not a trivial task. We will start with a special case and then extend it to a more general case. This special case regards consistent complete theories. Let K be a complete consistent theory, and $\diamond_\delta$ a splinter fully KM rational. We want to identify a contra-headed relation $\leqslant_K$, such that $\delta(K, \varphi) = Max_{\leqslant_K}(\omega(\varphi))$, for every consistent formula φ. For now, we will say that $\leqslant_K$ is the silhouette of δ within K.

Let us assign to each formula φ a relation $\leqslant_\varphi$ such that the maximal elements of $\omega(\varphi)$ modulo $\leqslant_\varphi$ match $\delta(K, \varphi)$. We achieve this by making $\leqslant_\varphi = \{(A, B) \in \omega(\varphi) \times \omega(\varphi) \mid A \notin \delta(K, \varphi)$ and $B \in \delta(K, \varphi)\}$. Then, we can assemble all $\leqslant_\varphi$ together to reach the general relation $\leqslant_K$. Although this seems a good idea, the relation $\leqslant_K$ constructed in this way is not strong enough to satisfy *quasi-reflection*. We adapt *triangulation* [Ribeiro *et al.*, 2018] for this purpose.

Definition 5. *Given a complete consistent theory K, an appointee δ and two consistent formulae φ and ψ, the $\delta-$triangulation of φ and ψ modulo K is the relation $\nabla(\varphi, \psi) = \{(A, B) \in \omega(\varphi) \times \omega(\varphi) \mid A \in \omega(\varphi) \setminus (\delta(K, \varphi) \cup \omega(\psi))$ and $B \in \delta(K, \psi)\}$.*

The δ-triangulation we employ here is a weaker version of *triangulation* used in [Ribeiro *et al.*, 2018]. Given a theory K, and a appointee δ, a triangulation connects the complete consistent theories of φ that were not chosen by δ to the theories of ψ that were chosen by $\delta(K, \psi)$. This makes all non-chosen elements of φ immediately less preferable to all elements chosen by ψ. The main difference between the triangulation version we use here from the original one is that we are indifferent about the complete theories that entail both φ and ψ. For instance, if a complete theory A entails both φ and ψ then A will not participate in the triangulation $\nabla(\varphi, \psi)$.

Now we go back to our problem of finding the silhouette of a complete consistent theory K within an appointee δ. This can be achieved by combining triangulation and our first naive idea of assembling relations $\leqslant_{\varphi}$.

Definition 6. *Given a complete theory K and an appointee δ, the projection of δ modulo K is the relation $\leqslant_K$ such that $(A, B) \in \leqslant_K$ iff for some formulae φ and ψ*

either $(A, B) \in \leqslant_{\varphi}$ or $(A, B) \in \nabla(\varphi, \psi)$.

To construct a relational appointee, we take the collection of all the projections of the consistent complete theories modulo δ. We say that this is the shadow of δ. It turns out that the projection of δ modulo K is also the silhouette of δ within K, when the splinter function $\diamond_{\delta}$ is fully KM rational.

Proposition 2. *Consider a splinter $\diamond_{\delta}$, and a consistent complete theory K. If $\diamond_{\delta}$ is fully KM rational, then the projection of δ modulo K is the silhouette of δ within K.*

Consequently appointees of fully KM rational splinters are relational. To see this, notice that from the proposition above $\delta(K, \varphi) = max_{\leqslant}(\omega(\varphi))$, for each consistent complete theory K, when $\diamond_{\delta}$ is fully KM rational. And hence fully KM rational splinters are indeed royal.

Theorem 6. *The appointee of a fully KM rational function $\diamond_{\delta}$ is relational.*

Corollary 2. *Fully KM rational functions are royal splinters.*

To show that fully KM rational functions are royal reflected splinters, we show that fully KM rational splinters are reflected. For this, we need to show that the appointee of fully rational splinters are *quasi-reflected*. The triangulation used to construct projection gives us this property in an easy way.

Theorem 7. *Let μ be the appointee from a fully KM rational splinter. Given a complete consistent theory K, the projection $\leqslant_K$ of δ modulo K is quasi-reflected.*

Therefore, we have that every fully KM rational belief change function is a royal reflected splinter. This together with Corollary 1 constitute the representation result between KM postulates and *royal reflected splinters*.

Corollary 3. *An update function is fully KM rational iff it is a royal reflected splinter.*

5   Discussion and Conclusion

Relying solely on preorders does not guarantee the success postulate in absence of the *maximal-cut* condition to guarantee the existence of maximal elements. We will assume a stronger version of the CUP functions: that their preorder are *contra-headed*. Let us call it CUP$^+$functions. It is easy to see that CUP$^+$are a special case of royal splinters. Thus,

Theorem 8. *Every CUP$^+$function satisfies (K$\diamond$1) to (K$\diamond$6) and (K$\diamond$8) and (K$\diamond$9).*

It turns out that transitivity and reflexivity add no substantial change in the royal functions, and every royal splinter function that satisfies (K$\diamond$6) and (K$\diamond$8) is a CUP$^+$function.

Notice that CUP$^+$functions constrained to quasi-reflected relations are royal splinters. Besides the three KM supplementary postulates, the following postulate was proposed to be used in place of postulates (K$\diamond$7) and (K$\diamond$8):

(K$\diamond$10)   if K is complete and $\neg\psi \notin K \diamond \varphi$ then
$\qquad K \diamond \varphi + \psi \subseteq K \diamond \varphi \wedge \psi$.

Postulate (K$\diamond$10) is intimately connected with total preorders as shown in [Katsuno and Mendelzon, 1992], which is taken to be a limitation [Lindström and Rabinowicz, 1991; Rott, 1992]. We show a new class of royal splinter functions characterised by the postulates (K$\diamond$1) to (K$\diamond$6) and (K$\diamond$9) and (K$\diamond$10) that do not depend on total preorders.

As the basic KM postulates and (K$\diamond$6) are satisfied by royal splinter functions, we will still maintain the restriction of appointees to be relational. To capture (K$\diamond$10) we will replace the quasi-reflection by the mirroring condition proposed by [Ribeiro *et al.*, 2018]:

(Mirroring) if $S_1 \not\leqslant S_2$ and $S_2 \not\leqslant S_1$; then for any $S' \in T_L$, if $S_1 \leqslant S'$ then $S_2 \leqslant S'$.

It is worth highlighting that the *mirroring* corresponds to *quasi-reflection* constrained to its condition (i). Thus, every mirrored relation is a quasi-reflected relation.

Theorem 9. *Every mirrored royal splinter satisfies (K$\diamond$10).*

Theorem 10. *Every update function that satisfies (K$\diamond$6) and (K$\diamond$10) is a mirrored royal splinter.*

In this work, we have addressed the issue of KM update in the context of non-finitary logics. We have found both negative and positive results. On the negative side, the CUP functions, characterised by the KM update postulates in Propositional Logic, cannot satisfy all the postulates without being finitary. We showed that, in finitary logics, CUP functions and a slightly stronger version of them (the CUP$^+$) are characterised by a proper subset of the KM rationality postulates. Assuming the KM update postulates as a good guide for belief update, we have devised a new class of update functions, the royal splinter functions, which are (fully) characterised by the KM postulates.

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. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Gabbay *et al.*, 1994] Dov M. Gabbay, Ian Hodkinson, and Mark Reynolds. *Temporal Logic (Vol. 1): Mathematical Foundations and Computational Aspects*. Oxford University Press, Inc., New York, NY, USA, 1994.

[Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In Peter Gärdenfors, editor, *Belief revision*, pages 183–203. Cambridge University Press, 1992.

[Lindström and Rabinowicz, 1991] Sten Lindström and Wlodzimierz Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In *The logic of theory change*, pages 93–126. Springer, 1991.

[Restall and Slaney, 1995] Greg Restall and John Slaney. Realistic belief revision. In *Proceedings of the Second World Conference on Foundations of Artificial Intelligence (WOCFAI-95)*, pages 367–378, 1995.

[Ribeiro *et al.*, 2013] Márcio M. Ribeiro, Renata Wassermann, Giorgos Flouris, and Grigoris Antoniou. Minimal change: Relevance and recovery revisited. *Artificial Intelligence*, 201:59 – 80, 2013.

[Ribeiro *et al.*, 2018] Jandson S. Ribeiro, Abhaya Nayak, and Renata Wassermann. Towards belief contraction without compactness. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, pages 287–296, 2018.

[Ribeiro *et al.*, 2019] Jandson S. Ribeiro, Abhaya Nayak, and Renata Wassermann. Belief change and non-monotonic reasoning sans compactness. In *the Thirty-Third AAAI Conference on Artificial Intelligence, (forthcoming)*, 2019.

[Rott, 1992] Hans Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1(1):45–78, 1992.

[Wassermann, 2011] Renata Wassermann. On AGM for non-classical logics. *Journal of Philosophical Logic*, 40(2):271–294, 2011.

[Winslett, 1988] Marianne Winslett. Reasoning about action using a possible models approach. In *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence*, AAAI'88, pages 89–93. AAAI Press, 1988.

[Zhang and Ding, 2008] Yan Zhang and Yulin Ding. CTL model update for system modifications. *Journal of Artificial Intelligence Research*, 31:113–155, 2008.