Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Context-Based Belief Revision

Nicolas Schwind
National Institute of Advanced Industrial Science and Technology, Tokyo, Japan
nicolas-schwind@aist.go.jp

Abstract

In credibility-limited (CL) belief revision, an agent may reject new information if it is considered not credible relative to its current beliefs. A core principle of CL revision requires that all consequences of a credible formula must themselves be credible. We propose a new framework, context-based (CB) belief revision, which generalizes CL revision by relaxing this requirement. In CB revision, a formula may be deemed credible because it strengthens one of its non-credible consequences by providing sufficient supporting context, a situation that CL revision does not allow. We introduce an axiomatic framework for CB revision operators, identify specific subclasses, provide representation theorems, and examine the relationships between CB revision operators, their subclasses, and CL revision operators.

1 Introduction

Belief revision operators (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988; Katsuno and Mendelzon 1991; Fermé and Hansson 2011) describe how an agent incorporates new information into its belief state. A central assumption in belief revision is the primacy of update principle, which requires the revised belief state to include the new information and is formalized as the success postulate. The AGM framework assumes belief states are deductively closed sets of formulas (Alchourrón, Gärdenfors, and Makinson 1985), while Katsuno and Mendelzon (KM) adapted this approach to represent belief states as single propositional formulas (Katsuno and Mendelzon 1991). In the KM setting, revision operators satisfy six postulates (R1)-(R6), with (R1) expressing the success postulate that the revised belief state entails the new information.

In many situations, however, agents may refuse to revise their beliefs, for instance when new information is considered not credible. Research on non-prioritized belief revision addressed this by weakening the success postulate (Hansson 1997a; Hansson 1997b; Fermé and Hansson 1999; Hansson et al. 2001; Booth et al. 2012; Schwind, Konieczny, and Marquis 2018; Garapa, Fermé, and Reis 2018; Garapa, Fermé, and Reis 2020; Garapa, Fermé, and Reis 2025). Here, we focus on the credibility-limited (CL) revision framework from Booth et al. (2012), which builds upon the KM approach. CL revision replaces the KM success postulate (R1) with the *relative success* postulate (P1), which requires that the revised beliefs either entail the new information (if credible) or remain unchanged (if not). The CL framework adapts the KM characterization accordingly, comprising six postulates (P1)-(P6) that formalize this relaxation. Our work builds on this setting.

To illustrate CL revision, consider the following example, which will be used throughout the paper:

Example 1. *Alice, who worked from home yesterday, believes that her colleague Bob went to the institute and had lunch with colleagues in the cafeteria. Let this belief state be φ. Later, another colleague tells her that Bob did not have lunch with colleagues but instead had lunch with the prime minister. Denote this information by p, where $p \models \neg\varphi$. Since Alice finds this implausible, she rejects it and keeps her beliefs unchanged: $\varphi \circ p \equiv \varphi$.*

This is a typical CL revision case, where new information is rejected due to lack of credibility.

A key postulate in CL revision is *success monotonicity* (P5) (Booth et al. 2012), which requires that if a formula is credible, then all of its logical consequences must also be credible. This postulate prevents acceptance of information that is credible only in a richer context, as shown by extending the example:

Example 1 (continued)**.** *In addition to p, Alice is informed that the prime minister visited the institute yesterday, denoted by q. This additional context makes the combined statement $p \wedge q$ (i.e., the prime minister visited the institute and Bob had lunch with the prime minister) credible to Alice, even though p alone remains implausible.*

Here, the more informative input $p \wedge q$ is accepted, while its logical consequence p is not credible in isolation. This clearly violates success monotonicity.

An extreme case occurs when Alice receives the input $\alpha = \neg\varphi$, which is exactly the negation of her current beliefs; in other words, she is simply told that she is wrong, without any additional information or justification. Under CL revision, success monotonicity implies that if such a bare contradiction is not credible, then no contradictory information can be credible, regardless of supporting context. This highlights how the principle can be overly restrictive and block reasonable belief changes when additional context is present.

To address this limitation, we introduce context-based (CB) revision, a framework that generalizes CL revision by dropping the success monotonicity postulate (P5) and making the necessary adjustments to other postulates, such as weakening the trichotomy postulate (P6). We also add two new conditions to better capture the role of context in assessing credibility. As a result, CB revision allows situations where a formula is credible even if some of its logical consequences are not.

The remainder of the paper is organized as follows. Section 2 reviews KM revision and CL revision. Section 3 presents the axiomatic characterization of CB revision operators. Section 4 provides a representation result for CB revision operators based on total preorders over worlds defined by a structure called a CB faithful assignment. Section 5 introduces two natural subclasses of CB revision, along with their representation results and interrelations. Proofs of propositions are available online.[1]

2 Preliminaries

We consider a propositional language $\mathcal{L}$ of formulas built up from a finite set of propositional variables P and the usual connectives. A world is a mapping from P to $\{0,1\}$. The set of all worlds is denoted by Ω. A world ω satisfies a formula φ if ω makes φ true; such a world is called a model of φ, and the set of models of φ is denoted by $[\varphi]$. A formula is consistent if $[\varphi] \neq \emptyset$. The set of consistent formulas is denoted by $\mathcal{L}_*$. The symbol $\models$ denotes logical entailment, and $\equiv$ logical equivalence: $\varphi \models \psi$ iff $[\varphi] \subseteq [\psi]$, and $\varphi \equiv \psi$ iff $[\varphi] = [\psi]$. Given a set of worlds W, the notation ϕ_W denotes a formula whose models are exactly W, i.e., $[\phi_W] = W$. When W is written explicitely, the brackets are sometimes omitted: for example, ϕ_ω stands for $\phi_{\{\omega\}}$ and $\phi_{\omega,\omega'}$ for $\phi_{\{\omega,\omega'\}}$. A formula is called complete if it has exactly one model.

Given a set of worlds $W \subseteq \Omega$, a total preorder $(W, \preceq)$ (or simply $\preceq$ on W) is a binary relation on W that is total, reflexive and transitive. Given a total preorder $(W, \preceq)$ and a set $W' \subseteq \Omega$, we define $\min(W', \preceq) = \{\omega \in W \cap W' \mid \forall \omega' \in W \cap W', \omega \preceq \omega'\}$.

A belief change operator $\circ$ models how an agent's belief state φ changes upon receiving new evidence α. Formally, $\circ$ is a mapping $\circ : \mathcal{L} \times \mathcal{L} \to \mathcal{L}$ associating the pair (φ, α) with a formula $\varphi \circ \alpha$ representing the agent's new belief state after revising φ by α. Katsuno and Mendelzon (KM) characterized a class of such operators by a set of postulates, with the central success postulate (R1) requiring the revised belief $\varphi \circ \alpha$ to entail the new evidence α. The complete set of postulates is given below:

Definition 1 (KM Operator (Katsuno and Mendelzon 1991))**.** *A belief change operator $\circ$ is called a* KM operator *if for all formulas $\varphi, \psi, \alpha, \beta \in \mathcal{L}$,*

(R1) $\varphi \circ \alpha \models \alpha$
(R2) *If* $\varphi \wedge \alpha \not\models \bot$ *then* $\varphi \circ \alpha \equiv \varphi \wedge \alpha$
(R3) *If* $\alpha \not\models \bot$ *then* $\varphi \circ \alpha \not\models \bot$
(R4) *If* $\varphi \equiv \psi$ *and* $\alpha \equiv \beta$ *then* $\varphi \circ \alpha \equiv \psi \circ \beta$
(R5) $(\varphi \circ \alpha) \wedge \beta \models \varphi \circ (\alpha \wedge \beta)$
(R6) *If* $(\varphi \circ \alpha) \wedge \beta \not\models \bot$ *then* $\varphi \circ (\alpha \wedge \beta) \models (\varphi \circ \alpha) \wedge \beta$

KM operators can be characterized in terms of total preorders over worlds. This is done by associating each KM operator with a *faithful assignment*:

Definition 2 (Faithful Assignment)**.** *A* faithful assignment *is a mapping associating each formula $\varphi \in \mathcal{L}$ with a total preorder $\preceq_\varphi$ on Ω such that the following conditions are satisfied for all $\varphi, \psi \in \mathcal{L}$:*

1. $[\varphi] = \min(\Omega, \preceq_\varphi)$
2. *If* $\varphi \equiv \psi$, *then* $\preceq_\varphi = \preceq_\psi$

Theorem 1 ((Katsuno and Mendelzon 1991))**.** *A belief change operator $\circ$ is a KM operator if and only if there exists a faithful assignment $\varphi \mapsto \preceq_\varphi$ such that for all $\varphi, \alpha \in \mathcal{L}$:*

$$[\varphi \circ \alpha] = \min([\alpha], \preceq_\varphi)$$

Credibility-limited revision operators generalize KM operators by relaxing the success postulate (R1), and are characterized by the following set of postulates:

Definition 3 (Credibility-Limited Revision Operator (Booth et al. 2012))**.** *A belief change operator $\circ$ is called a* credibility-limited revision operator *(CL operator for short) if for all formulas $\varphi, \psi \in \mathcal{L}_*$, $\alpha, \beta \in \mathcal{L}$,*

(P1) $\varphi \circ \alpha \models \alpha$ *or* $\varphi \circ \alpha \equiv \varphi$
(P2) *If* $\varphi \wedge \alpha \not\models \bot$ *then* $\varphi \circ \alpha \equiv \varphi \wedge \alpha$
(P3) $\varphi \circ \alpha \not\models \bot$
(P4) *If* $\varphi \equiv \psi$ *and* $\alpha \equiv \beta$ *then* $\varphi \circ \alpha \equiv \psi \circ \beta$
(P5) *If* $\varphi \circ \alpha \models \alpha$ *and* $\alpha \models \beta$, *then* $\varphi \circ \beta \models \beta$
(P6) $\varphi \circ (\alpha \vee \beta) \equiv \begin{cases} \varphi \circ \alpha, & \text{or} \\ \varphi \circ \beta, & \text{or} \\ (\varphi \circ \alpha) \vee (\varphi \circ \beta) \end{cases}$

In the CL framework, the success postulate (R1) is replaced by the relative success postulate (P1), which states that if the new input is not credible ($\varphi \circ \alpha \not\models \alpha$), the belief state remains unchanged. Postulates (P2) and (P4) correspond to (R2) and (R4), while (P3) is a strengthening of (R3). The trichotomy postulate (P6), which captures minimal change, corresponds to (R5) and (R6) in the KM framework, given (R1) and (R4) (Booth et al. 2012). Postulate (P5) plays a central role in CL revision: it requires that if a formula is credible (and thus accepted), all its logical consequences must also be credible. This has no direct counterpart in the KM framework, where it is trivially satisfied due to (R1).

The CL postulates are formulated for consistent belief states ($\varphi \in \mathcal{L}_*$) because, in Booth et al. (2012), belief change operators are defined as mappings $\circ : \mathcal{L}_* \times \mathcal{L} \to \mathcal{L}$. Since our work builds on the CL framework, we assume, for the remainder of the paper, that belief change operators are defined on consistent belief states, i.e., $\circ : \mathcal{L}_* \times \mathcal{L} \to \mathcal{L}$, with any comparisons to KM operators restricted accordingly.

CL operators can be semantically characterized by CL faithful assignments, which generalize the standard faithful assignments from the KM framework to accommodate the notion of input credibility:

[1] https://nicolas-schwind.github.io/Schwind-KR25-proofs.pdf

Definition 4 (CL Faithful Assignment)**.** *A* CL faithful assignment *is a mapping associating each formula $\varphi \in \mathcal{L}_*$ with a pair $(C_\varphi, \preceq_\varphi)$, where $[\varphi] \subseteq C_\varphi \subseteq \Omega$, $\preceq_\varphi$ is a total preorder on C_φ, and such that the following conditions are satisfied for all $\varphi, \psi \in \mathcal{L}_*$:*

1. $[\varphi] = \min(C_\varphi, \preceq_\varphi)$
2. *If* $\varphi \equiv \psi$, *then* $(C_\varphi, \preceq_\varphi) = (C_\psi, \preceq_\psi)$

Theorem 2 ((Booth et al. 2012))**.** *A belief change operator $\circ$ is a CL operator if and only if there exists a CL faithful assignment $\varphi \mapsto (C_\varphi, \preceq_\varphi)$ such that for all $\varphi, \alpha \in \mathcal{L}_*$:*

$$[\varphi \circ \alpha] = \begin{cases} \min([\alpha], \preceq_\varphi) & \text{if } [\alpha] \cap C_\varphi \neq \emptyset \\ [\varphi] & \text{otherwise} \end{cases}$$

CL faithful assignments reduce to standard faithful assignments precisely when $C_\varphi = \Omega$ for every $\varphi \in \mathcal{L}_*$, i.e., when all worlds are considered credible. As shown in Booth et al. (2012), this implies that any CL operator satisfying (R1) is in fact a KM operator.

Theorem 3 ((Booth et al. 2012))**.** *If a CL operator $\circ$ satisfies (R1) , then $\circ$ is a KM operator.*

Note that in the presentation of this result in Booth et al. (2012), the restriction to consistent inputs α is applied implicitly. Without this restriction, (R1) and (P3) become incompatible when $\alpha \models \bot$: (R1) requires $\varphi \models \alpha$, which contradicts (P3). This incompatibility is avoided by excluding this case, which is an uninteresting one because it is fully determined in both frameworks. Specifically, KM operators return an inconsistent state (by (R1)), while CL operators return a formula equivalent to the initial belief state φ (by (P1) and (P3)). In this paper, we make the same implicit assumption when comparing some of the newly introduced operator classes with KM operators.

3 Context-Based Revision

Let us first introduce some terminology, which will prove useful to explain the rationale of the postulates we will introduce in this paper:

Definition 5 (Context-Based Credibility)**.** *Let $\circ$ be a belief change operator and $\varphi \in \mathcal{L}_*$. A formula $\alpha \in \mathcal{L}$ is:*

- φ-absolutely credible for $\circ$ *if for each $\beta \in \mathcal{L}$,* $\varphi \circ (\alpha \vee \beta) \models \alpha \vee \beta$
- φ-absolutely non-credible for $\circ$ *if for each $\beta \in \mathcal{L}$,* $\varphi \circ (\alpha \wedge \beta) \not\models \alpha \wedge \beta$
- φ-contextually credible for $\circ$ *if* $\varphi \circ \alpha \models \alpha$ *and there exists $\beta \in \mathcal{L}$ such that* $\varphi \circ (\alpha \vee \beta) \not\models \alpha \vee \beta$
- φ-contextually non-credible for $\circ$ *if* $\varphi \circ \alpha \not\models \alpha$ *and there exists $\beta \in \mathcal{L}$ such that* $\varphi \circ (\alpha \wedge \beta) \models \alpha \wedge \beta$

We will omit "for $\circ$" (respectively, the prefix "φ-") when $\circ$ (respectively, φ) is clear from context. The four cases in Definition 5 are mutually exclusive and collectively exhaustive: every formula α falls into exactly one of them. A formula is said to be *credible* if it is absolutely or contextually credible, and *non-credible* otherwise.

Using this terminology, the success postulate (R1) can be restated as requiring every formula to be absolutely credible. More importantly, (P5) demands that all credible formulas be absolutely credible; equivalently, by contrapositive, all non-credible formulas are absolutely non-credible. In other words, CL revision excludes contextual credibility altogether: no formula can be contextually credible or contextually non-credible. Our goal is to develop a framework for context-based (CB) revision that explicitly accommodates contextual credibility. For instance, our introductory example illustrates a formula that is contextually non-credible (p, Bob had lunch with the prime minister yesterday) and one that is contextually credible ($p \wedge q$, where q means the prime minister visited the institute yesterday). To achieve this goal, we weaken the CL framework by revisiting its postulates.

(P1) governs the outcome of a revision by non-credible inputs: the resulting belief state should remain unchanged. This principle is orthogonal to our generalization, so we maintain this behavior for our CB operators, regardless of whether the non-credibility is absolute or contextual. Similarly, (P3) and (P4), which ensure consistency and syntax irrelevance, are unrelated to our generalization and thus remain unchanged. (P2) captures minimal change when the input is consistent with the agent's beliefs. In CB revision, this means that inputs consistent with the agent's beliefs are absolutely credible, and revision should reduce to expansion. This assumption is natural since such inputs align with the agent's current beliefs, so we retain (P2).[2] (P5) is the key postulate we relax. As previously discussed, it excludes contextual (non-)credibility, so we drop it entirely. Finally, (P6), the trichotomy postulate, reflects a principle of minimal change that we want to preserve in spirit. However, in its current form, it is too strong. For instance, a direct consequence of (P6) is that if two formulas are credible, then their disjunction must also be credible. Stated equivalently, this says:

Proposition 1. *If a belief change operator $\circ$ satisfies (P6), then for all $\varphi, \alpha \in \mathcal{L}_*$, if $\varphi \circ \alpha \not\models \alpha$, then there is a complete formula $\phi_\omega \models \alpha$ such that $\varphi \circ \phi_\omega \not\models \phi_\omega$.*

However, it is easy to imagine a scenario where the agent considers every complete formula credible because each is fully specified or well-supported in context, yet still refuses to revise her beliefs, φ, when faced with a vague or unclear input, such as simply being told being "wrong" ($\alpha = \neg\varphi$).

Example 1 (continued)**.** *Alice finds both q (the prime minister visited the institute) and r (the prime minister visited the town office) individually credible. However, the disjunction $q \vee r$ might be too vague for her to identify a specific plausible scenario. As a result, $q \vee r$ can be contextually non-credible, since it does not specify which situation actually occurred.*

To accommodate such situations, we propose replacing (P6) with the following Relative Trichotomy (RT) postulate:

[2] As briefly discussed at the end of the paper, contextual credibility could be extended to belief expansion. Weakening (P2) to apply relative success in expansion would allow for greater flexibility, but this potential extension is not addressed here.

(RT) If $\varphi \circ (\alpha \vee \beta) \models \alpha \vee \beta$, then

$$\varphi \circ (\alpha \vee \beta) \equiv \begin{cases} \varphi \circ \alpha & \text{or} \\ \varphi \circ \beta & \text{or} \\ (\varphi \circ \alpha) \vee (\varphi \circ \beta) \end{cases}$$

We can observe that in presence of (P1), (P2) and (P3), the postulate (RT) is equivalent to credibility-limited weakenings of (R5) and (R6):

(CR56) If $\varphi \circ \alpha \models \alpha$ and $(\varphi \circ \alpha) \wedge \beta \not\models \bot$, then $\varphi \circ (\alpha \wedge \beta) \equiv (\varphi \circ \alpha) \wedge \beta$

Proposition 2. *If $\circ$ is a belief change operator satisfying (P1), (P2) and (P3), then $\circ$ satisfies (RT) if and only if it satisfies (CR56).*

Notably, (CR56) can be seen as a conjunction of the *guarded subexpansion* postulate from Hansson et al. (2001) ((R6') If $\varphi \circ \alpha \models \alpha$ and $(\varphi \circ \alpha) \wedge \beta \not\models \bot$, then $\varphi \circ (\alpha \wedge \beta) \models (\varphi \circ \alpha) \wedge \beta$) and a similar weakening of (R5). While Booth et al. (2012) showed that (P6) is equivalent to (R5) and (R6') in the presence of the remaining CL postulates, here we need to weaken (R5) similarly, since removing (P5) could lead to situations where $\varphi \circ (\alpha \wedge \beta) \models \alpha \wedge \beta$ and $\varphi \circ \alpha \not\models \alpha$, causing $(\varphi \circ \alpha) \wedge \beta$ ($\equiv \varphi \wedge \beta$ by (P1)) not to entail $\varphi \circ (\alpha \wedge \beta)$.

At this stage, we have removed (P5) and replaced (P6) with (RT). However, the resulting set (P1)-(P4) and (RT) still allows some unexpected behaviors. To refine the notion of contextual credibility, we introduce two core principles for CB revision. The first captures a monotonic behavior of credibility with respect to strengthening and weakening formulas. We call this postulate Contextual Monotonicity (CM):

(CM) If $\varphi \circ \alpha \models \alpha$, $\varphi \circ (\alpha \vee \beta) \not\models \alpha \vee \beta$ and $\varphi \circ (\alpha \wedge \gamma) \not\models \alpha \wedge \gamma$, then $\varphi \circ (\alpha \wedge \gamma \wedge \delta) \not\models \alpha \wedge \gamma \wedge \delta$

According to the terminology of Definition 5, (CM) states that if strengthening a contextually credible formula α results in a non-credible formula $\alpha \wedge \gamma$, this non-credibility persists for all further strengthenings $\alpha \wedge \gamma \wedge \delta$, i.e., $\alpha \wedge \gamma$ is absolutely non-credible. This postulate can also be read, via its contrapositive, as asserting that if weakening a contextually non-credible formula leads to a credible formula, then this credibility persists for all further weakenings; in other words, the formula becomes absolutely credible. This principle can be illustrated with our running example:

Example 1 (continued)**.** *Recall that p (Bob had lunch with the prime minister yesterday) is non-credible, while $p \wedge q$ (where q means the prime minister visited Bob's institute yesterday) is credible for Alice. Now, consider the statement that Bob had dinner with the prime minister yesterday (s). Although $q \wedge s$ might be credible (supported by the prime minister's visit), Alice rejects the idea that Bob had both lunch and dinner with the prime minister, making $p \wedge q \wedge s$ non-credible. This non-credibility is not due to lack of contextual support but because Alice cannot conceive any credible scenario where both events occur. Therefore, $p \wedge q \wedge s$ is absolutely non-credible, meaning any further strengthening remains non-credible.*

Our second postulate is the principle of Contextual Cohesion (CC), which says that if two formulas α and β are individually credible, and if their conjunction $\alpha \wedge \beta$ is also credible, then their disjunction $\alpha \vee \beta$ must be credible:

(CC) If $\varphi \circ \alpha \models \alpha$, $\varphi \circ \beta \models \beta$ and $\varphi \circ (\alpha \wedge \beta) \models \alpha \wedge \beta$, then $\varphi \circ (\alpha \vee \beta) \models \alpha \vee \beta$

Note that if any of the formulas α, β, or $\alpha \wedge \beta$ is absolutely credible, then $\alpha \vee \beta$ is trivially credible by definition. Thus, (CC) mainly addresses the case of contextual credibility: in terms of Definition 5, it requires that if two formulas and their conjunction are contextually credible, then their disjunction must also be contextually credible.

This postulate ensures that credibility remains when considering credible alternatives within the same context. If α and β are contextually credible, and their conjunction $\alpha \wedge \beta$ is also credible, this means there is a credible scenario for $\alpha \wedge \beta$ that satisfies both formulas, each with enough context support. In this case, the disjunction $\alpha \vee \beta$ should also be credible: it does not introduce ambiguity but represents alternative cases within the same supporting context.

Example 1 (continued)**.** *Recall that Alice finds both $p \wedge q$ and $q \wedge s$ credible. Now, unlike before, suppose Alice also finds credible that Bob had both lunch and dinner with the prime minister, i.e., $p \wedge q \wedge s$ is credible. This means that weakening this last statement in either direction (considering either lunch or dinner during the prime minister's visit) does not reduce its contextual support for Alice. Hence, the disjunction $p \vee s$, combined with q, remains credible because it preserves that support without introducing ambiguity.*

With the two additional axioms (CM) and (CC) in place, we are now ready to introduce our definition of context-based (CB) revision operator:

Definition 6 (Context-Based Revision Operator)**.** *A belief change operator $\circ$ is called a* context-based revision operator *(CB operator for short) if it satisfies (P1)-(P4), (RT), (CM) and (CC).*

It is important to note that CB revision generalizes CL revision. Specifically, adding (P5) to CB revision yields exactly a CL revision operator.

Proposition 3. *A belief change operator is a CL operator if and only if it is a CB operator that satisfies (P5).*

We can show that the generalization is strict by presenting a concrete example of CB operators that are not CL operators in general. Given two formulas $\mu, \sigma \in \mathcal{L}$, let $\bullet_{\mu,\sigma}$ be the belief change operator defined for each $\varphi \in \mathcal{L}_*$, $\alpha \in \mathcal{L}$ as:

$$\varphi \bullet_{\mu,\sigma} \alpha = \begin{cases} \varphi \wedge \alpha & \text{if } \varphi \wedge \alpha \not\models \bot \\ \mu \wedge \alpha & \text{if } \varphi \wedge \alpha \models \bot, \mu \wedge \alpha \not\models \bot, \\ & \text{and } \alpha \models \sigma \\ \varphi & \text{in the remaining case} \end{cases}$$

Proposition 4. *For each $\mu, \sigma \in \mathcal{L}$, $\bullet_{\mu,\sigma}$ is a CB operator.*

This operator is defined by two formulas, μ and σ. Intuitively, μ characterizes the set of credible worlds that are not already models of the beliefs φ, while σ represents a (single) context. For the operator $\bullet_{\mu,\sigma}$, a formula α is considered credible if either (1) α is consistent with the agent's

current beliefs ($\varphi \wedge \alpha \not\models \bot$), or (2) it has some credible models according to μ ($\mu \wedge \alpha \not\models \bot$) and is supported by the context σ ($\alpha \models \sigma$).

When the formula σ characterizing the context is valid ($\top \models \sigma$), every credible formula is absolutely credible. In this case, $\bullet_{\mu,\sigma}$ behaves as a CL operator. On the other hand, $\bullet_{\mu,\sigma}$ is not a CL operator in the general case, particularly when σ is not valid and $\mu \wedge \sigma$ is consistent. For instance, if we take $\mu = \sigma = p \wedge q$, it is easy to verify that $\neg p \bullet_{\mu,\sigma} p \equiv \neg p$ and $\neg p \bullet_{\mu,\sigma} (p \wedge q) \equiv p \wedge q$. This example shows that (P5) is not satisfied by $\bullet_{\mu,\sigma}$. In fact, we can show that:

Proposition 5. *For each $\mu, \sigma \in \mathcal{L}$, the operator $\bullet_{\mu,\sigma}$ is:*

1. *a CL operator if and only if (σ is valid or $\mu \wedge \sigma \models \bot$)*
2. *a KM operator if and only if both μ and σ are valid*

The class of operators $\bullet_{\mu,\sigma}$ thus provides a genuine example of CB revision that lies outside the CL framework.

4 A Representation Result

We now aim to demonstrate how CB operators can be semantically represented, following the standard approach in belief change, using total preorders over worlds. This is achieved by introducing the concept of a *CB faithful assignment*, which associates each formula $\varphi \in \mathcal{L}_*$ with what we call a *context structure space*.

The core component of a context structure space is a tuple, referred to as a *context structure*, denoted by $\mathcal{C} = (Cont, Cred, \preceq)$, where $\emptyset \neq Cred \subseteq Cont \subseteq \Omega$, and $\preceq$ is a total preorder on $Cred$. A context structure $\mathcal{C} = (Cont, Cred, \preceq)$ is said to be *absolute* if $Cont = \Omega$. In what follows, when a context structure is annotated with a superscript integer i, it implicitly denotes $\mathcal{C}^i = (Cont^i, Cred^i, \preceq^i)$.

A *context structure space* is a tuple $\mathcal{S} = \langle \mathcal{C}^0, \ldots, \mathcal{C}^k \rangle$, $k \geq 0$, where for each $i \in \{0, \ldots, k\}$, $\mathcal{C}^i$ is a context structure, $\mathcal{C}^0$ is absolute, and for each $i, j \in \{1, \ldots, k\}$, if $i \neq j$ then $Cred^i_\varphi \cap Cont^j_\varphi = \emptyset$. In what follows, we often associate a formula $\varphi \in \mathcal{L}_*$ with a context structure space, denoted $\mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$, where for each $i \in \{0, \ldots, k\}$, $\mathcal{C}^i_\varphi = (Cont^i_\varphi, Cred^i_\varphi, \preceq^i_\varphi)$.

Definition 7 (CB Faithful Assignment)**.** *A* CB faithful assignment *(CB assignment for short) Ψ is a mapping associating each formula $\varphi \in \mathcal{L}_*$ with a context structure space $\mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ where $k \geq 0$ and such that the following conditions are satisfied for all $\varphi, \psi \in \mathcal{L}_*$:*

1. $[\varphi] = \min(Cred^0_\varphi, \preceq^0_\varphi)$
2. *If $\varphi \equiv \psi$, then $\mathcal{S}_\varphi = \mathcal{S}_\psi$*

We intend to show that CB operators can be represented in terms of CB assignments. To start, let us provide some intuition about the elements involved in a CB assignment. When a formula φ is associated with a context structure space $\mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$, the set $Cred^0_\varphi$ within the absolute context structure $\mathcal{C}^0_\varphi$ is a set of credible worlds that determine absolute credibility: precisely, the formulas whose models overlap with $Cred^0_\varphi$ are absolutely credible. Thus, $Cred^0_\varphi$ plays a role analogous to the set of credible worlds C_φ in a CL faithful assignment (Definition 4). The other context structures, $\mathcal{C}^i_\varphi$ for $i \in \{1, \ldots, k\}$, are central to the CB framework. In each such $\mathcal{C}^i_\varphi = (Cont^i_\varphi, Cred^i_\varphi, \preceq^i_\varphi)$ ($i \geq 1$), $Cred^i_\varphi$ represents the *set of credible worlds* within that context structure, and $Cont^i_\varphi$ is a *contextual set of worlds*. Each context structure characterizes credible formulas: a formula is credible with respect to a given context structure $\mathcal{C}^i_\varphi$ if all of its models lie in $Cont^i_\varphi$, and at least one model is in $Cred^i_\varphi$. Let us make this notion precise:

Definition 8. *Given a context structure $\mathcal{C} = (Cont, Cred, \preceq)$, we say that a formula $\alpha \in \mathcal{L}$ is* $\mathcal{C}$-credible *if $[\alpha] \cap Cred \neq \emptyset$ and $[\alpha] \subseteq Cont$.*

Let us observe that:

Lemma 1. *Let $\mathcal{S} = \langle \mathcal{C}^0, \ldots, \mathcal{C}^k \rangle$ be a context structure space and $\alpha \in \mathcal{L}$. If α is $\mathcal{C}^i$-credible for some $i \in \{1, \ldots, k\}$, then for all $j \in \{1, \ldots, k\}$, if $i \neq j$, then $[\alpha] \cap Cred^j = \emptyset$ and $[\alpha] \not\subseteq Cont^j$.*

Lemma 1 implies that for each formula $\alpha \in \mathcal{L}$, α is $\mathcal{C}^i$-credible for at most one non-absolute context structure $\mathcal{C}^i \in \mathcal{S}$. When there exists $i \in \{1, \ldots, k\}$ such that α is $\mathcal{C}^i$-credible, in what follows we simply denote such index i by i_α, and we say that α is $\mathcal{C}^{i_\alpha}$-credible.

Let us now define a belief change operator from a CB assignment:

Definition 9. *Let Ψ be a CB assignment $\varphi \mapsto \mathcal{S}_\varphi$ with $\mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$, $k \geq 0$, and let $\circ$ be a belief change operator. We say that* $\circ$ is represented by Ψ *if for all $\varphi \in \mathcal{L}_*$, $\alpha \in \mathcal{L}$,*

$$[\varphi \circ \alpha] = \begin{cases} \min([\alpha], \preceq^0_\varphi) & \text{if } \alpha \text{ is } \mathcal{C}^0_\varphi\text{- credible} \\ \min([\alpha], \preceq^{i_\alpha}_\varphi) & \text{if } \alpha \text{ is not } \mathcal{C}^0_\varphi\text{- credible} \\ & \text{and } \alpha \text{ is } \mathcal{C}^{i_\alpha}_\varphi\text{- credible} \\ [\varphi] & \text{in the remaining case} \end{cases}$$

We can prove that:

Proposition 6. *If a belief change operator $\circ$ is represented by a CB assignment, then $\circ$ is a CB operator.*

The proof takes advantage of the following lemma:

Lemma 2. *If a belief change operator $\circ$ is represented by a CB assignment $\varphi \mapsto \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$, $k \geq 0$, then for each $\varphi \in \mathcal{L}_*$ and each $\alpha \in \mathcal{L}$, if α is φ-contextually credible, then α is not $\mathcal{C}^0_\varphi$-credible and α is $\mathcal{C}^{i_\alpha}_\varphi$-credible.*

This lemma shows that there is exactly one (non-absolute) context structure that supports a given contextually credible formula.

We now construct an assignment $\Psi_\circ$ from a given belief change operator $\circ$. We will show that if $\circ$ is a CB operator, then this assignment is a CB assignment and represents $\circ$ as in Definition 9. The construction is done as follows:

Definition 10 (Assignment Corresponding to $\circ$)**.** *Given a belief change operator $\circ$, the* assignment corresponding to $\circ$, *denoted by $\Psi_\circ$, is an assignment associating each formula $\varphi \in \mathcal{L}_*$ with a set $\mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$, $k \geq 0$, where for each $i \in \{0, \ldots, k\}$, $\mathcal{C}^i_\varphi = (Cont^i_\varphi, Cred^i_\varphi, \preceq^i_\varphi)$, $Cred^i_\varphi, Cont^i_\varphi \subseteq \Omega$ and $\preceq^i_\varphi$ is a binary relation on $Cred^i_\varphi$, all of which are defined from $\circ$ as follows, for each $\varphi \in \mathcal{L}_*$:*

- $Cont^0_\varphi = \Omega$
- $Cred^0_\varphi = \{\omega \in \Omega \mid \forall \omega' \in \Omega \quad \varphi \circ \phi_{\omega,\omega'} \models \phi_{\omega,\omega'}\}$
- $\{Cont^i_\varphi \mid 1 \leq i \leq k\} = \max(\mathcal{E}, \subseteq)$
- $\forall i \in \{1, \ldots, k\}$, $Cred^i_\varphi = \{\omega \in Cont^i_\varphi \mid \varphi \circ \phi_\omega \models \phi_\omega\}$
- $\forall i \in \{0, \ldots, k\}$ $\forall \omega, \omega' \in Cred^i_\varphi \quad \omega \preceq^i_\varphi \omega'$ *iff* $\omega \in [\varphi \circ \phi_{\omega,\omega'}]$

where $\max(\mathcal{E}, \subseteq) = \{W \in \mathcal{E} \mid \forall W' \in \mathcal{E} \ \ W \not\subset W'\}$, *and* $\mathcal{E} = \{W \subseteq (\Omega \setminus Cred^0_\varphi) \mid \exists \omega \in W : \forall \omega' \in W \ \varphi \circ \phi_{\omega,\omega'} \models \phi_{\omega,\omega'}\}$.

For the sake of illustration, take again the CB operator $\bullet_{\mu,\sigma}$ introduced in the previous section, that is defined by two formulas $\mu, \sigma \in \mathcal{L}$. It is not difficult to verify that the following assignment, denoted by $\Psi_{\bullet_{\mu,\sigma}}$, corresponds to $\bullet_{\mu,\sigma}$:

Definition 11 (Assignment $\Psi_{\bullet_{\mu,\sigma}}$)**.** *For each $\mu, \sigma \in \mathcal{L}$, let $\Psi_{\bullet_{\mu,\sigma}} : \varphi \mapsto \mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ be assignment defined for each $\varphi \in \mathcal{L}_*$ as:*

- *if $\neg\sigma \models \varphi$ or $\mu \wedge \sigma \models \varphi$, then $\mathcal{S}_\varphi = (\mathcal{C}^0_\varphi)$, with $Cred^0_\varphi = [\varphi \vee (\mu \wedge \sigma)]$ and $\forall \omega, \omega' \in Cred^0_\varphi$, $\omega \preceq^0_\varphi \omega'$ iff $\omega \in [\varphi]$ or $\omega' \in [\neg\varphi]$*
- *otherwise, $\mathcal{S}_\varphi = (\mathcal{C}^0_\varphi, \mathcal{C}^1_\varphi)$, with $Cred^0_\varphi = [\varphi]$, $Cont^1_\varphi = [\neg\varphi \wedge \sigma]$, $Cred^1_\varphi = [\neg\varphi \wedge \mu \wedge \sigma]$, and $\forall i \in \{0, 1\}$, $\preceq^i_\varphi = Cred^i_\varphi \times Cred^i_\varphi$.*

Indeed, accordingly:

Proposition 7. *For each μ, σ, the assignment $\Psi_{\bullet_{\mu,\sigma}}$ corresponds to $\bullet_{\mu,\sigma}$.*

We can now show that:

Proposition 8. *If $\circ$ is a CB operator, then its corresponding assignment $\Psi_\circ$ is a CB assignment.*

The proof goes through a set of notable observations:

Lemma 3. *Let $\circ$ be a CB operator. Then the assignment $\varphi \mapsto \mathcal{S}_\varphi$ that represents $\circ$ satisfies the following conditions, for each $\varphi \in \mathcal{L}_*$, all $i, j \in \{0, \ldots, k\}$, all $\omega, \omega' \in \Omega$ and each $W \subseteq \Omega$:*

1. $[\varphi] \subseteq Cred^0_\varphi$
2. $\emptyset \neq Cred^i_\varphi \subseteq Cont^i_\varphi$
3. *if $W \cap Cred^i_\varphi \neq \emptyset$ and $W \subseteq Cont^i_\varphi$, then $\varphi \circ \phi_W \models \phi_W$*
4. *if $j \geq 1$ and $i \neq j$, then $Cred^i_\varphi \cap Cont^j_\varphi = \emptyset$*

Lemma 3.1 says that all models of φ are in the set $Cred^0_\varphi$, that is the set defining absolute credibility. Lemma 3.2 states that each contextual set of worlds contains the set of credible worlds within the same context structure, and that each such set of credible worlds is not empty. Lemma 3.3 is key: a formula is credible if all its models are in the contextual set $Cont^i_\varphi$ and at least one is in the set of credible worlds $Cred^i_\varphi$. Finally, Lemma 3.4, together with Lemma 3.2, implies that each world can belong to the set of credible worlds of at most one context structure, ensuring that the sets of credible worlds do not overlap.

The last step of our representation result consists in showing that the construction of a CB assignment $\Psi_\circ$ wrt. a CB operator is faithful to the representation of $\circ$ by $\Psi_\circ$ in Definition 9:

Proposition 9. *If a belief change operator $\circ$ is a CB operator, then it is represented by $\Psi_\circ$.*

The proof takes advantage of the following results:

Lemma 4. *Let $\circ$ be a CB operator and $\Psi_\circ$ be its corresponding CB assignment. Then for all $\varphi \in \mathcal{L}_*$, $\alpha \in \mathcal{L}$,*

1. *$\forall i \in \{0, \ldots, k\}$, if α is $\mathcal{C}^i_\varphi$-credible, then $[\varphi \circ \alpha] \subseteq Cred^i_\varphi$*
2. *if $\varphi \circ \alpha \models \alpha$, then $\exists i \in \{0, \ldots, k\}$ s.t. α is $\mathcal{C}^i_\varphi$-credible*

Lemma 4.1 states that the CB revision by a credible formula, whether absolutely or contextually credible, results in an output formula whose models all lie in the set of credible worlds of a single context structure. Lemma 4.2 says that any credible formula must be $\mathcal{C}^i_\varphi$-credible with respect to at least one context structure (absolute or not).

Propositions 7 and 9 together mean that one can recover the definition of $\bullet_{\mu,\sigma}$ through Definition 9 using its representation by the assignment $\Psi_{\bullet_{\mu,\sigma}}$ (cf. Definition 11).

Propositions 6, 8 and 9 together provide us with a representation result, stated in the following theorem:

Theorem 4. *A belief change operator is a CB operator if and only if it is represented by a CB assignment.*

Before concluding this section, we illustrate how the CB assignment $\Psi_\circ$ representing a CB operator $\circ$ can be restricted to recover the representations of CL and KM operators. We also show how to capture CB operators that satisfy the original trichotomy condition (P6) satisfied by CL operators:

Proposition 10. *Let $\circ$ be a CB operator and $\Psi_\circ : \varphi \mapsto \mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ be its corresponding CB assignment. Then:*

1. *$\circ$ is a CL operator iff for each $\varphi \in \mathcal{L}_*$, $\mathcal{S}_\varphi = (\mathcal{C}^0_\varphi)$*
2. *$\circ$ is a KM operator iff for each $\varphi \in \mathcal{L}_*$, $Cred^0_\varphi = \Omega$*
3. *$\circ$ is satisfies (P6) iff for each $\varphi \in \mathcal{L}_*$, $k \leq 1$*

Propositions 10.1 and 10.2 provide alternative representations of KM and CL operators in terms of CB assignments. Proposition 10.1 shows that a CB assignment represents a CL operator precisely when it comprises only a single (absolute) context structure, which is equivalent to a structure $(C_\varphi, \preceq_\varphi)$ used in a CL faithful assignment (cf. Definition 4). Likewise, Proposition 10.2 says that KM operators are represented by CB assignments with a single (absolute) context structure where all worlds are credible, mirroring the notion of faithful assignments as in Definition 2. Finally, Proposition 10.3 offers a noteworthy insight: CB operators satisfying the original trichotomy condition (P6) of CL revision are exactly those whose corresponding CB assignment comprises at most one non-absolute context structure.

5 Two Subclasses of CB Operators

This section introduces additional constraints on CB revision, focusing on how contextual credibility and non-credibility are further governed in their interaction with absolutely non-credible formulas. We define two subclasses of CB operators, each distinguished by an extra postulate.

The first of these is the Context-Exclusive postulate (CE), so named because it requires the exclusion of any absolutely non-credible formula from being consistent with a contextually credible one. Specifically, it requires that contextually credible formulas exclude all worlds that are never deemed credible, i.e., worlds that do not lie in any of the sets of credible worlds from the context structures in the corresponding CB assignment. This postulate enforces a stricter notion of contextual credibility, ensuring that successful changes never result from revision based on inputs that involve worlds that are fundamentally non-credible.

Definition 12 (Context-Exclusive Revision Operator)**.** *A CB revision operator $\circ$ is called a* Context-Exclusive (CE) operator *if it satisfies the following condition:*

(CE) *If $\varphi \circ \alpha \models \alpha$, $\varphi \circ (\alpha \vee \beta) \not\models \alpha \vee \beta$ and $\alpha \wedge \gamma \not\models \bot$, then $\varphi \circ (\alpha \wedge \gamma) \models \alpha \wedge \gamma$*

This postulate can be restated using the terms from Definition 5: if a formula α is contextually credible and its conjunction with another formula γ is consistent, then this conjunction must also be credible. In other words, (CE) requires that a formula is contextually credible only if all its models represent situations the agent considers credible. As a result, if a formula is consistent but not credible, making it credible again by weakening can only make it absolutely credible. It can be verified that (CE) strengthens (CC) as (if $\varphi \circ \alpha \models \alpha$, $\varphi \circ \beta \models \beta$ and $\alpha \wedge \beta \not\models \bot$, then $\varphi \circ (\alpha \vee \beta) \models \alpha \vee \beta$); while (CC) requires in its precondition the conjunction $\alpha \wedge \gamma$ to be credible, (CE) only requires it to be consistent.

Example 1 (continued)**.** *Recall that p is non-credible for Alice, while $p \wedge q$ is credible. Now, under (CE), $p \wedge q$ is "*strongly*" credible: every possible situation satisfying it would be accepted by Alice. In this view, absolutely non-credible formulas act in a similar way to contexts (but non-credible ones) so that any overlap with them turns an otherwise contextually credible statement into one that is ambiguous and thus lacks sufficient contextual support. Thus, if a statement is contextually credible, it is credible in this "strong" sense, and adding further information, for instance that Bob also had dinner with the prime minister ($p \wedge q \wedge s$), will also be credible, provided it is consistent.*

This matches applications where being compatible with an absolutely non-credible situation is taken to mean there is insufficient contextual support.

Note that (P5) trivially implies (CE), since (P5) rules out contextually credible inputs and thus ensures the precondition of (CE) always holds. As a result, CE operators still properly generalize the CL framework:

Proposition 11. *Every CL revision operator is a CE revision operator.*

The CE representation result in terms of CB assignments can be found through the following proposition:

Proposition 12. *Let $\circ$ be a CB revision operator and $\Psi_\circ : \varphi \mapsto \mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ be its corresponding CB assignment. Then $\circ$ is a CE revision operator if and only if for each $\varphi \in \mathcal{L}_*$ and for all $i \in \{1, \ldots, k\}$,*

$$Cont^i_\varphi = Cred^i_\varphi$$

For a CE operator $\circ$, each contextual set $Cont^i_\varphi$ coincides with its corresponding set of credible worlds $Cred^i_\varphi$ for $i \geq 1$. Thus, a contextually credible formula α is $\mathcal{C}^{i_\alpha}_\varphi$-credible ($i_\alpha \geq 1$) when all its models lie within $Cred^{i_\alpha}_\varphi$. As a result, any contextually credible formula β consistent with α satisfies $[\alpha \vee \beta] \subseteq Cred^{i_\alpha}_\varphi$, making $\alpha \vee \beta$ also $\mathcal{C}^{i_\alpha}_\varphi$-credible.

The following result shows that the generalization is strict, as some instances of the CB operators $\bullet_{\mu,\sigma}$ introduced earlier lead to CE operators that are not CL ones. It also shows that the class of CE operators is a strict subclass of CB operators:

Proposition 13. *For each $\mu, \sigma \in \mathcal{L}$, the operator $\bullet_{\mu,\sigma}$ is a CE operator if and only if σ is valid or $\mu \wedge \sigma \models \bot$ or $\sigma \models \mu$.*

Proposition 13 directly shows that CE operators form a strict subclass of CB operators. Together with Proposition 5.1, which characterizes CL operators with a stricter condition, this implies that CE operators strictly generalize CL operators.

The second postulate we introduce is the Context-Tolerant postulate (CT), so named because it allows contextually credible formulas to be consistent with any absolutely non-credible formula. This principle captures the intuition that the contextual credibility of a formula should not be lost just because it is consistent with situations that are never considered credible. Rather, a formula is not contextually credible only because it remains ambiguous between two credible situations arising from different contexts.

Definition 13 (Context-Tolerant Revision Operator)**.** *A CB operator $\circ$ is called a* Context-Tolerant (CT) operator *if it satisfies the following condition:*

(CT) *If $\varphi \circ \alpha \models \alpha$ and for all γ, (if $\gamma \models \beta$ then $\varphi \circ \gamma \not\models \gamma$), then $\varphi \circ (\alpha \vee \beta) \models \alpha \vee \beta$*

In other words, (CT) says that extending a contextually credible formula by disjoining it with an absolutely non-credible formula does not affect its credibility, as non-credible worlds are irrelevant to contextual credibility.

Example 1 (continued)**.** *Suppose again that $p \wedge q$ is contextually credible for Alice. Now, consider the statement: "King Charles III visited Bob's institute yesterday" (t), which Alice considers absolutely non-credible, meaning no situation involving t is credible to her. The statement $p \wedge (q \vee t)$ is a weaker form of $p \wedge q$, by including t as an alternative. However, it still retains enough context for Alice to recognize it as credible.*

Recall that under (CE), absolutely non-credible formulas act like other contexts, in the sense that any weakening of a contextually credible statement that overlaps with them

makes it non-credible. Under (CT), such overlaps do not affect contextual credibility; only credible alternative contexts create ambiguity.

Just as (P5) implies (CE) by vacuity, establishing that all CE operators are CL operators, we can similarly observe that (P5) implies (CT), which shows that:

Proposition 14. *Every CL operator is a CT operator.*

An interesting point is that the full CL framework coincides with the class of CT operators that satisfy (P6). This follows from Proposition 10.3, which shows that CB operators with (P6) correspond to at most one non-absolute context structure, and from the next lemma, which rules out for CT operators with (P6) the case of exactly one non-absolute context structure:

Lemma 5. *Let $\circ$ be a CT operator and $\Psi_\circ : \varphi \mapsto \mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ be its corresponding CB assignment. Then for each $\varphi \in \mathcal{L}_*$, $\langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ is such that $k \neq 1$.*

Proposition 15. *A belief change operator is CL operator if and only if it is a CT operator satisfying (P6).*

The CT representation result can be found through the following proposition:

Proposition 16. *Let $\circ$ be a CB operator and $\Psi_\circ : \varphi \mapsto \mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ be its corresponding CB assignment. Then $\circ$ is a CT operator if and only if for each $\varphi \in \mathcal{L}_*$ and for all $i \in \{1, \ldots, k\}$,*

$$Cont^i_\varphi = \Omega \setminus \bigcup \{Cred^j_\varphi \mid j \in \{0, \ldots, k\}, j \neq i\}$$

Proposition 16 says that for a CT operator $\circ$, any world outside all sets of credible worlds $Cred^j_\varphi$ belongs to every contextual set $Cont^i_\varphi$. It follows that a formula α is contextually non-credible precisely when its models intersect at least two distinct sets of credible worlds $Cred^i_\varphi$ and $Cred^j_\varphi$ with $i \neq j$ and $i, j \geq 1$. Equivalently, there exists a formula β such that $\alpha \wedge \beta$ is $\mathcal{C}^{i_{\alpha\wedge\beta}}_\varphi$-credible ($[\alpha \wedge \beta] \cap Cred^{i_{\alpha\wedge\beta}}_\varphi \neq \emptyset$ and $[\alpha \wedge \beta] \subseteq Cont^{i_{\alpha\wedge\beta}}_\varphi$), and every world $\omega \in [\neg\alpha \wedge \beta]$ lies in some set(s) of credible worlds $Cred^j_\varphi$ with $j \neq i_{\alpha\wedge\beta}$.

From Definition 11, Lemma 5 and Proposition 10.1, any CB operator $\bullet_{\mu,\sigma}$ satisfying (CT) is necessarily a CL operator. In combination with Proposition 14, we obtain:

Proposition 17. *For each $\mu, \sigma \in \mathcal{L}$, the operator $\bullet_{\mu,\sigma}$ is a CT operator if and only if it is a CL operator.*

This means that the class of CB operators $\bullet_{\mu,\sigma}$ cannot yield a CT operator that is not also a CL operator. Indeed, any operator $\bullet_{\mu,\sigma}$ corresponds to a CB assignment that has at most one non-absolute context structure determined by σ (cf. Definition 11), and Lemma 5 combined with Proposition 10.1 show that at least two such structures are required for a CT operator to fall outside the CL class.

Thus, we introduce a second class of concrete operators denoted by $\bullet^{CT}_{\mu,\sigma}$, which are also defined by two given formulas $\mu, \sigma \in \mathcal{L}$ as follows, for each $\varphi \in \mathcal{L}_*$ and each $\alpha \in \mathcal{L}$:

$$\varphi \bullet^{CT}_{\mu,\sigma} \alpha = \begin{cases} \varphi \wedge \alpha & \text{if } \varphi \wedge \alpha \not\models \bot \\ \mu \wedge \alpha & \text{if } \varphi \wedge \alpha \models \bot, \mu \wedge \alpha \not\models \bot, \\ & \text{and } (\alpha \models \sigma \text{ or } \alpha \models \neg\sigma) \\ \varphi & \text{in the remaining case} \end{cases}$$

As intended, we get that:

Proposition 18. *For each $\mu, \sigma \in \mathcal{L}$, the operator $\bullet^{CT}_{\mu,\sigma}$ is a CT operator.*

For an operator $\bullet^{CT}_{\mu,\sigma}$, the formula μ defines the set of credible worlds (for worlds not already models of the beliefs φ), similar to an operator $\bullet_{\mu,\sigma}$. However, unlike $\bullet_{\mu,\sigma}$ operators, here the formula σ defines two distinct contextual sets from non-absolute context structures: both $[\sigma]$ and $[\neg\sigma]$ serve as contextual sets. Thus, an input α is contextually credible only if it entails either σ or $\neg\sigma$.

Using this new class of operators, we can now give a concrete example of a CT operator that is not a CL operator. Consider the case where $\mu = \sigma = p \wedge q$, which matches the example in Section 3 illustrating a CB operator that is not a CL operator. We verify that $\neg p \bullet_{\mu,\sigma} p \equiv \neg p$ and $\neg p \bullet_{\mu,\sigma} (p \wedge q) \equiv p \wedge q$, showing that $\bullet_{\mu,\sigma}$ does not satisfy (P5). Moreover, this example highlights that (P5) is not satisfied in cases where μ is consistent and σ is consistent but not valid, which are precisely the interesting scenarios:

Proposition 19. *For each $\mu, \sigma \in \mathcal{L}$, the operator $\bullet^{CT}_{\mu,\sigma}$ is a CL operator if and only if $\sigma \models \bot$ or σ is valid or $\mu \models \bot$*

The final connection we examine concerns CB operators that satisfy both postulates (CE) and (CT). In the rest of this section, we refer to them simply as CET operators. Interestingly, as a direct consequence of the characterizations of CE and CT operators via CB assignments (cf. Prop. 12 and Prop. 16), we obtain the following:

Corollary 1. *Let $\circ$ be a CB revision operator and $\Psi_\circ : \varphi \mapsto \mathcal{S}_\varphi = \langle \mathcal{C}^0_\varphi, \ldots, \mathcal{C}^k_\varphi \rangle$ be its corresponding CB assignment. Then $\circ$ is a CET operator if and only if for each $\varphi \in \mathcal{L}_*$, $\bigcup \{Cred^i_\varphi \mid i \in \{0, \ldots, k\}\} = \Omega$*

Corollary 1 states that for a CET operator $\circ$, each world belongs to a credible set $Cred^i_\varphi$ in the corresponding CB assignment. This implies that every possible world is "credible," meaning any non-credible formula can become credible by strengthening it with suitable contextual support:

Corollary 2. *A CB operator $\circ$ is a CET operator if and only if it satisfies:*

(CET) *If α is a complete formula, then $\varphi \circ \alpha \models \alpha$*

This means that for CET operators, every formula is either credible or contextually non-credible; in other words, CET operators rule out absolutely non-credible formulas.

Another notable consequence of our representation results is the relationship between CET operators and the trichotomy postulate (P6). Indeed, CET operators satisfying (P6) fully capture the KM framework:

Proposition 20. *A belief change operator is a KM operator if and only if it is a CET operator that satisfies (P6).*

Lastly, CL operators form a strict subclass of CET operators. In particular, we can construct a CT operator $\bullet^{CT}_{\mu,\sigma}$ that is also a CE operator but not a CL operator. For example, the case $\mu = \sigma = p \wedge q$, already discussed earlier, satisfies (CT) and (CE) but does not satisfy (P5). More precisely, we have the following characterization:

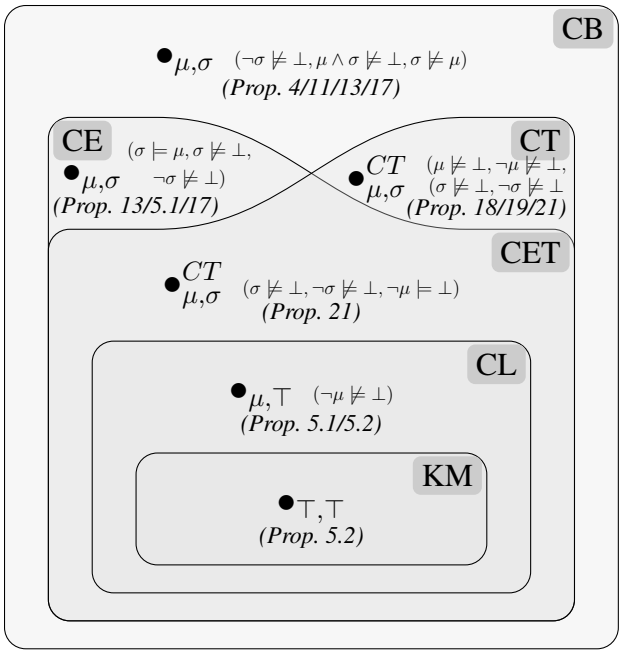

Figure 1: Dependencies between CB, CE, CT, CET, CL and KM operators. We have that KM $\subseteq$ CL (Th. 3), CL $\subseteq$ CE, (Prop. 11), CL $\subseteq$ CT (Prop. 14), and CE $\subseteq$ CB and CT $\subseteq$ CB by definition. Each operator $\bullet_{\mu,\sigma}$, $\bullet^{CT}_{\mu,\sigma}$ in the figure is representative of the class in which it is depicted, showing that each inclusion is strict.

Proposition 21. *For each $\mu, \sigma \in \mathcal{L}$, the CT operator $\bullet^{CT}_{\mu,\sigma}$ is a CE operator and not a CL operator if and only if $\sigma \not\models \bot$, $\neg\sigma \not\models \bot$, and μ is valid.*

We have now established the complete relationships between CB, CE, CT, CET, CL, and KM operators. Furthermore, we have shown how adding the postulate (P6) places additional constraints on some of these classes.

Figure 1 summarizes key results by illustrating the dependencies among these categories and showing that all six classes CB, CE, CT, CET, CL, and KM are distinct. For each class C in this collection, the figure presents a concrete example of an operator (either $\bullet_{\mu,\sigma}$ or $\bullet^{CT}_{\mu,\sigma}$) that belongs to C but not to any other class that does not contain C.

6 Summary and Discussion

The success condition in belief revision, formalized by the postulate (R1) in the KM framework, has long been debated. A common relaxation, relative success (P1), restricts changes to cases where the new input belongs to a set of credible formulas, leading to the credibility-limited (CL) revision framework. In this paper, we introduced a strict generalization of CL revision, termed context-based (CB) revision, by removing the success monotonicity postulate (P5) and weakening the trichotomy postulate (P6) to Relative Trichotomy (RT). Success monotonicity demands that all consequences of a credible formula are themselves credible. In contrast, CB revision allows a formula to be non-credible not because it lacks any credible world, but because it lacks sufficient contextual support. This means the formula can gain credibility if it is strengthened with enough context.

To capture this intuition, we proposed two new postulates, Contextual Monotonicity (CM) and Contextual Cohesion (CC), which serve as weakened versions of (P5) and describe expected behaviors of CB revision. We then provided a representation theorem for CB revision operators in terms of CB assignments, involving context structures made up of sets of credible worlds, contextual sets of worlds, and total preorders over credible worlds. Additionally, we introduced postulates (CE) and (CT) to characterize specific CB scenarios focusing on their interaction with non-credible worlds. Finally, we examined how the various resulting classes of operators relate to one another by leveraging their respective representation results through CB assignments.

Before concluding, we briefly relate our approach to two relevant works that touch on aspects of weakening success conditions and incorporating explanations in belief revision.

The work of Garapa, Fermé, and Reis (2020) examines belief base revision and systematically explores subclasses of operators obtained by weakening individual postulates. One such postulate is strict improvement, the counterpart of success monotonicity (P5) at the belief set level. Their results show that operators can exist without satisfying this condition. However, their analysis does not introduce new postulates tailored to specific applications, so it does not explicitly address contextual information.

A different but related line of research investigates the role of explanations in belief revision (Falappa, Kern-Isberner, and Simari 2002). In their framework, an explanation consists of an explanans E (supporting beliefs) and an explanandum O (the conclusion), working over belief bases and requiring E to be given separately and entail O. In CB revision, the added context can be viewed as an explanans supporting the input formula. However, since our framework assumes full rational inference with syntax independence (P4), there is no need to separate explanans and explanandum explicitly: the strengthened input already incorporates its context. While Falappa, Kern-Isberner, and Simari (2002) explicitly link E and O via explanations, in CB revision the context supports credibility without necessarily acting as a formal explanation.

Our work opens several research avenues. The two concrete classes of operators we introduced, $\bullet_{\mu,\sigma}$ and $\bullet^{CT}_{\mu,\sigma}$, are simple and could be extended to exploit the full expressivity of total preorderings, for example by incorporating distance-based revision operators such as Dalal's (Dalal 1988) to define a broader class of CB operators. Another interesting direction is exploring how contexts evolve for iterative use, inspired by Darwiche and Pearl's epistemic states (Darwiche and Pearl 1997), as done in the CL framework (Booth et al. 2012). Further generalizations of CB revision are also possible, particularly in relaxing the vacuity postulate (P2) when contextual credibility is extended to belief expansion. For example, Alice's initial beliefs about Bob's lunch may not exclude the possibility of him dining with the prime minister, yet she may still reject this due to insufficient context. Weakening (P2) to apply relative success in expansion would allow for more flexibility, similar to our modification of (P6). This potential extension is left for future research.

Acknowledgements

This work was supported by JSPS KAKENHI Grant Number JP25K00375.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

Booth, R.; Fermé, E.; Konieczny, S.; and Pérez, R. P. 2012. Credibility-limited revision operators in propositional logic. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR'12)*, 116–125.

Dalal, M. 1988. Investigations into a theory of knowledge base revision: preliminary report. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI'88)*, 475–479.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89(1-2):1–29.

Falappa, M. A.; Kern-Isberner, G.; and Simari, G. R. 2002. Explanations, belief revision and defeasible reasoning. *Artificial Intelligence* 141(1/2):1–28.

Fermé, E. L., and Hansson, S. O. 1999. Selective revision. *Studia Logica* 63(3):331–342.

Fermé, E., and Hansson, S. O. 2011. AGM 25 years: Twenty-five years of research in belief change. *Journal of Philosophical Logic* 40(2):295–331.

Garapa, M.; Fermé, E.; and Reis, M. D. L. 2018. Studies in credibility-limited base revision. In *Proceedings of the 16th International Conference on the Principles of Knowledge Representation and Reasoning (KR'18)*, 240–247.

Garapa, M.; Fermé, E.; and Reis, M. D. L. 2020. Credibility-limited base revision: New classes and their characterizations. *Journal of Artificial Intelligence Research* 69:1023–1075.

Garapa, M.; Fermé, E.; and Reis, M. D. L. 2025. Maps between different kinds of two-level credibility-limited revision operators. *ACM Transactions on Computational Logic* 26(3):16:1–16:33.

Gärdenfors, P. 1988. *Knowledge in flux*. MIT Press.

Hansson, S. O.; Fermé, E. L.; Cantwell, J.; and Falappa, M. A. 2001. Credibility limited revision. *The Journal of Symbolic Logic* 66(4):1581–1596.

Hansson, S. O. 1997a. Semi-revision. *Journal of Applied Non-Classical Logics* 7(2).

Hansson, S. O. 1997b. What's new isn't always best. *Theoria* 63(1-2):1–13.

Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.

Schwind, N.; Konieczny, S.; and Marquis, P. 2018. On belief promotion. In *Proceedings of the 16th International Conference on the Principles of Knowledge Representation and Reasoning (KR'18)*, 297–307.