

Collective Belief Revision

Theofanis I. Aravanis

TARAVANIS@UPATRAS.GR

Department of Mechanical Engineering

School of Engineering

University of the Peloponnese

Patras 263 34, Greece

Abstract

In this article, we study the dynamics of collective beliefs. As a first step, we formulate David Westlund’s Principle of Collective Change (PCC) —a criterion that characterizes the evolution of collective knowledge— in the realm of belief revision. Thereafter, we establish a number of unsatisfiability results pointing out that the widely-accepted revision operators of Alchourrón, Gärdenfors and Makinson, combined with fundamental types of merging operations —including the ones proposed by Konieczny and Pino Pérez as well as Baral et al.— collide with the PCC. These impossibility results essentially extend in the context of belief revision the negative results established by Westlund for the operations of contraction and expansion. At the opposite of the impossibility results, we also establish a number of satisfiability results, proving that, under certain (rather strict) requirements, the PCC is indeed respected for specific merging operators. Overall, it is argued that the PCC is a rather unsuitable property for characterizing the process of collective change. Last but not least, mainly in response to the unsatisfactory situation related to the PCC, we explore some alternative criteria of collective change, and evaluate their compliance with belief revision and belief merging.

1. Introduction

An intelligent agent, either *individual* or *collective*, should be capable of gathering information about the world, and changing her state of belief in response to new evidential information.¹ As a consequence, the agent should be able to implement *belief change*, an operation which is heavily studied in the realm of Artificial Intelligence (Gärdenfors, 1988; Fermé & Hansson, 2018).

Two fundamental types of belief change are *belief revision* (or simply revision) and *belief merging* (or simply merging). Belief revision is the process by which a rational agent *modifies* her beliefs in the light of new information (Gärdenfors, 1988), whereas, belief merging is the process by which a *community* of agents *aggregates* their beliefs, so that the *collective knowledge* of the community be formulated (Konieczny & Pérez, 2011). To illustrate the importance of these two types of belief change, let us present the subsequent real-world

1. A collective agent could, indicatively, be an organization, a company, a scientific committee, or any multi-agent system. Yet, other, more radical, definitions of collective agency have been proposed. For instance, Minsky (1988) in his book portrays the mind of an individual agent as a “society” of tiny components (tiny agents), which are in fact the fundamental thinking-entities; in that sense, even a single individual can be thought of as a collective entity. We shall formally define in Section 5 the notion of collective agency adopted herein.

situation related to space exploration, which at the same time highlights the motivation for this work.

“Very recently, NASA’s James Webb Space Telescope discovered six gigantic galaxies that emerged shortly after the Big Bang. According to experts, this is an astounding and unexpected discovery, as the uncovered galaxies have been formed at a speed that upends the current understanding of the early universe (Labbé et al., 2023).”

The above scenario describes a situation where a new scientific discovery *contradicts* the current collective knowledge of the scientific community. In response to this contradiction, each expert of the community ought to (individually) *revise* her beliefs. After the revision is complete, the collective knowledge of the scientific community will have changed. The issue that will concern us in this article is whether, and under which conditions, such transitions, from an initial to a revised state of collective knowledge, satisfy certain criteria of collective belief change.

While there is a substantial body of literature on the knowledge-dynamics of *individuals* (see (Fermé & Hansson, 2018) or (Peppas, 2008) for two surveys on the topic), the manner in which the knowledge of *collective agents* evolves in response to new evidential information —as illustrated in the aforementioned space-related scenario— has not been studied as much. A considerable portion of works on collective knowledge focuses on its *static* aspects (for example on the nature, representation and attitude of collective beliefs (Gilbert, 1987, 2004; Maynard-Zhang & Lehmann, 2003; Wray, 2001)), whereas, some representative studies examining collective belief change are the works of Kfir-Dahav and Tennenholtz (1996), Malheiro, Jennings, and Oliveira (1994), Liu and Williams (2001), Paranamana, Wang, and Shafto (2022), Liu, Seligman, and Girard (2014), Tojo (2013), and Westlund (2010). From these studies, the ones of Kfir-Dahav and Tennenholtz (1996), Malheiro et al. (1994), and Liu and Williams (2001) present basic formal frameworks that model different perspectives of multi-agent belief revision. In a recent work by Paranamana et al. (2022), a theoretical framework that applies tools from Markov chain theory to analyze the evolution of societal beliefs is proposed. Liu et al. (2014), on the other hand, model the way in which the beliefs of the members of a community are influenced by their social relations. In a different vein, Tojo (2013) studies collective belief revision by means of a numerical framework of linear algebra.

In our study, we shall focus on the work of David Westlund (2010), who formally addresses the problem of collective belief change, by means of well-established tools developed by Alchourrón, Gärdenfors, and Makinson (1985) for modelling the dynamics of knowledge. In that work, Westlund explores the cases of *contraction* and *expansion*, which are change-operations similar in spirit to belief revision, but at the same time retain notable distinctions — as a matter of fact, contraction can be defined in terms of revision, whereas, expansion is a special (degenerative) case of revision (Alchourrón et al., 1985). Confined to these change-operations, Westlund assumes that every individual of a community is rational, and, moreover, that the synthesis of individuals’ beliefs is employed by means of a merging operator. On these premises, he shows that interesting types of merging operators *violate* what he calls the *Principle of Collective Change* (abbreviated as PCC), a criterion that characterizes the evolution of collective knowledge (Westlund, 2010, Observations 6–10).

Following the route of Westlund, we study in this article the evolution of collective knowledge in the realm of belief revision; namely, in the principal case where the new information *contradicts* the belief corpus of each individual of a community. Accordingly, our contributions can be summarized in the following points:

- By formulating Westlund’s PCC in the realm of belief revision (Westlund, 2010, p. 217), we firstly establish a number of *unsatisfiability results* indicating that the rationality of individuals of a community —formalized by the *AGM model* of Alchourrón et al. (1985) for belief revision— and the aggregation of the beliefs of these individuals by means of *fundamental* merging operators —including the ones of Konieczny and Pérez (2002) and Baral et al. (1992)— do *not* necessarily imply a revision of the community’s collective knowledge adhering to the PCC. The obtained impossibilities essentially *extend* in the realm of belief revision Westlund’s negative results.
- At the opposite of the impossibility results, we also establish a number of *satisfiability results* proving that, if the individual agents of a community revise their beliefs employing operators that implement *uniform belief revision* (Areces & Becher, 2001; Aravanis, 2020), then the PCC is indeed *respected* by certain types of merging operators. These satisfiability results are of interest in their own right, but also offer a notable *proof-of-concept* for uniform belief revision.
- Despite the obtained satisfiability results, which, as we explain, rely on *strict* requirements, it is argued that the PCC is a rather *unsuitable* property for characterizing collective belief revision.
- Lastly, mainly in response to the unsatisfactory circumstance concerning the PCC, we present two *alternative* criteria that characterize the revision of collective knowledge, which are based on Konieczny and Pino Pérez’s main approach for belief merging (Konieczny & Pérez, 2002). The compliance of the proposed alternative criteria of collective change with belief revision and belief merging is evaluated, and some controversial outcomes are identified, as both satisfiability and unsatisfiability results are obtained.

The remainder of this article is structured as follows. The next section introduces the formal prerequisites required for our exposition. Sections 3 and 4 present an overview of the processes of belief revision and belief merging, respectively. Thereafter, Section 5 introduces Westlund’s PCC, expressed in the realm of belief revision. In Sections 6 and 7, the obtained unsatisfiability and satisfiability results concerning the PCC are presented, respectively. Section 8 explores the alluded alternative criteria of collective change, whereas, Section 9 concludes the article with a discussion on the ramifications of the derived results and some avenues of future research.

2. Formal Preliminaries

Throughout this article, we shall work with a propositional language \mathcal{L} , built over a non-empty, *finite* set \mathcal{P} of atoms (propositional variables), using the standard Boolean connectives \wedge (conjunction), \vee (disjunction), \rightarrow (implication), \leftrightarrow (equivalence), \neg (negation), and

governed by *classical propositional logic*. The classical consequence relation is denoted by \models . The lower-case English letters a, b and c are used to denote atoms of \mathcal{P} .

For a set of sentences Γ of \mathcal{L} , $Cn(\Gamma)$ denotes the set of all logical consequences of Γ ; i.e., $Cn(\Gamma) = \{\varphi \in \mathcal{L} : \Gamma \models \varphi\}$. An agent's corpus of beliefs shall be represented by a *theory*, also referred to as a *belief set*. A theory K is a deductively closed set of sentences of \mathcal{L} ; i.e., $K = Cn(K)$. The set of all theories is denoted by \mathbb{T} . For a belief set K and a sentence φ of \mathcal{L} , $K + \varphi$ denotes the set (theory) resulting from the *expansion* of K by φ , defined as $K + \varphi = Cn(K \cup \{\varphi\})$.

A *literal* is an atom $a \in \mathcal{P}$ or its negation. A *possible world* (or simply *world*) r is a consistent set of literals, such that, for any atom $a \in \mathcal{P}$, either $a \in r$ or $\neg a \in r$. The set of all possible worlds is denoted by \mathbb{M} . For a sentence (set of sentences) φ of \mathcal{L} , $[[\varphi]]$ is the set of all possible worlds at which φ is true.

Lastly, a *preorder* over a set of possible worlds M is any reflexive and transitive binary relation in M . A preorder \preceq is called *total* iff, for any $r, r' \in M$, $r \preceq r'$ or $r' \preceq r$. The strict part of \preceq is denoted by \prec ; i.e., $r \prec r'$ iff $r \preceq r'$ and $r' \not\preceq r$. The indifference part of \preceq is denoted by \approx ; i.e., $r \approx r'$ iff $r \preceq r'$ and $r' \preceq r$. We shall call a total preorder \preceq over M *indistinguishable* iff, for any $r, r' \in M$, $r \approx r'$. Lastly, $min(M, \preceq)$ denotes the set of all \preceq -minimal possible worlds of M ; i.e., $min(M, \preceq) = \{r \in M : \text{for all } r' \in M, \text{ if } r' \preceq r, \text{ then } r \preceq r'\}$.

3. Belief Revision

The process of belief revision has been formally captured by the seminal work of Alchourrón et al. (1985). In this section, we present the axiomatic (postulational) characterization of belief revision as developed by the AGM trio, as well as its semantic characterization in terms of a special kind of total preorders over possible worlds, as subsequently developed by Katsuno and Mendelzon (1991).

3.1 Axiomatic Characterization

A *revision function* is a binary function $*$ that maps a belief set K and a sentence φ (also referred to as *epistemic input*) to a belief set $K * \varphi$, representing the result of revising K by φ . We shall say that a revision function $*$ is an *AGM revision function* iff it satisfies the following widely-accepted rationality postulates of Alchourrón et al. (1985), known as *AGM revision postulates*.

- (K * 1) $K * \varphi$ is a theory of \mathbb{T} .
- (K * 2) $\varphi \in K * \varphi$.
- (K * 3) $K * \varphi \subseteq K + \varphi$.
- (K * 4) If $\neg\varphi \notin K$, then $K + \varphi \subseteq K * \varphi$.
- (K * 5) If φ is consistent, then $K * \varphi$ is consistent.
- (K * 6) If $Cn(\{\varphi\}) = Cn(\{\psi\})$, then $K * \varphi = K * \psi$.
- (K * 7) $K * (\varphi \wedge \psi) \subseteq (K * \varphi) + \psi$.
- (K * 8) If $\neg\psi \notin K * \varphi$, then $(K * \varphi) + \psi \subseteq K * (\varphi \wedge \psi)$.

The AGM revision postulates circumscribe the territory of all *rational* strategies by means of which a belief set of an agent can be revised. Therefore, if an agent employs an AGM revision function in order to respond to new evidential information, it is considered that the agent behaves in a rational manner. The rationale behind the AGM revision postulates is concretely discussed by Gärdenfors (1988, Section 3.3) and Peppas (2008, Section 8.3.1). We solely mention here that their guiding principle is the *economy of information*, according to which the new information φ is consistently incorporated into the belief set K , changing the latter as little as possible. Note lastly that, in the special case where φ is *consistent* with K , it follows from postulates $(K * 3)$ & $(K * 4)$ that revision reduces to *expansion*; thus, $K * \varphi = K + \varphi$. As earlier stated however, herein we consider only the principal case of genuine revision; namely, the case where the epistemic input φ *contradicts* the initial belief set K .

3.2 Semantic Characterization

Katsuno and Mendelzon (1991) showed that the revision functions satisfying the AGM revision postulates are precisely those that are induced by means of *faithful preorders*, a special kind of total preorders over all possible worlds.

Definition 1 (Faithful Preorder). *A total preorder \preceq_K over \mathbb{M} is faithful to a belief set K iff, for any possible worlds $r, r' \in \mathbb{M}$, the following two conditions hold:*

- (i) *If $r, r' \in \llbracket K \rrbracket$, then $r \approx_K r'$.*
- (ii) *If $r \in \llbracket K \rrbracket$ and $r' \notin \llbracket K \rrbracket$, then $r \prec_K r'$.*

Intuitively, a faithful preorder \preceq_K over \mathbb{M} encodes the comparative plausibility of the possible worlds of \mathbb{M} , relative to the belief set K . Specifically, the assertion $r \preceq_K r'$ states that the world r is at least as plausible as the world r' , with respect to K .

Definition 2 (Faithful Assignment, (Katsuno & Mendelzon, 1991)). *A faithful assignment is a function that maps each belief set K of \mathbb{T} to a total preorder \preceq_K over \mathbb{M} , that is faithful to K .*

Based on the notion of faithful assignment, Katsuno and Mendelzon obtained the following representation theorem, which characterizes the class of AGM revision functions in terms of faithful preorders.

Theorem 3 (Katsuno & Mendelzon, 1991). *A revision function $*$ satisfies postulates $(K * 1)$ – $(K * 8)$ iff there exists a faithful assignment that maps each belief set K to a total preorder \preceq_K over \mathbb{M} , such that, for any $\varphi \in \mathcal{L}$:*

$$(R) \quad \llbracket K * \varphi \rrbracket = \min(\llbracket \varphi \rrbracket, \preceq_K).$$

Hence, according to condition (R), the revised belief set $K * \varphi$ is specified as the theory corresponding to the most plausible φ -worlds, relative to the initial belief set K .

Since revision by self-contradictory epistemic input constitutes a limiting uninteresting case, we consider in the course of this article only the principal case of *consistent* new information.

4. Belief Merging

Belief merging is the process of aggregating the (perhaps conflicting) beliefs of individuals. In formal terms, belief merging studies the following question: Given a collection of belief sets K_1, \dots, K_n , each one representing the beliefs of an agent of a community, which is a belief set that can be considered to represent a plausible aggregation of K_1, \dots, K_n ? This question is particularly important in the context of many Artificial Intelligence fields, and thus, several merging operators that implement the aforementioned aggregation-process have been proposed in the literature. In this section, we review several interesting proposals of such operators, some of which were already discussed by Westlund (2010). Before presenting the alluded proposals, let us fix the required notation and terminology.

A *belief profile* (or simply *profile*) E is a non-empty multi-set of belief sets; thus, E could contain multiple instances of a belief set. The set of all profiles is denoted by \mathbb{E} . The multi-set union is denoted by \sqcup . For a profile $E \in \mathbb{E}$, we shall write $\llbracket E \rrbracket$ instead of $\llbracket \bigcup E \rrbracket$. We denote by E^n the profile in which E appears n times; that is, $E^n = \underbrace{E \sqcup \dots \sqcup E}_n$. Two profiles E_1, E_2 are equivalent, denoted by $E_1 \equiv E_2$, iff there is a bijective function f from E_1 to E_2 , such that, for any $K \in E_1$, $f(K) = K$.

Definition 4 (Merging Function). *A merging function \mathcal{M} is a binary function that maps a profile E of \mathbb{E} and a sentence μ of \mathcal{L} (representing integrity constraints) to a belief set $\mathcal{M}_\mu(E)$ of \mathbb{T} ; i.e., $\mathcal{M} : \mathbb{E} \times \mathcal{L} \mapsto \mathbb{T}$.*

Hence, a merging function \mathcal{M} generates a belief set $\mathcal{M}_\mu(E)$ which represents the aggregation of all the belief sets of the profile E that, at the same time, respects the *integrity constraints* encoded into the logical sentence μ . The integrity constraints encoded into μ , essentially, may be imposed by physical constraints, unquestionable knowledge and so forth. For a profile $\{\{K_1, \dots, K_n\}\}$ of \mathbb{E} , we shall write $\mathcal{M}_\mu(K_1, \dots, K_n)$ instead of $\mathcal{M}_\mu(\{\{K_1, \dots, K_n\}\})$. Furthermore, in case the sentence μ is a tautology, the subscript μ shall be dropped.²

In what follows, with an exception to be made in Section 8, we assume that the merging-process implemented by merging functions does *not* adhere to specific integrity constraints. In such circumstance, the sentence μ is a *tautology* and, thus, we shall treat (albeit abusively) a merging function as a *unary* function that aggregates multiple belief sets into a single belief set.

4.1 Union-When-Possible Merging Functions

Against this background, we introduce firstly the so-called *union-when-possible* merging functions, which were also considered by Westlund (2010).

Definition 5 (Union-When-Possible Merging Function). *A merging function \mathcal{M} is a union-when-possible merging function iff, for any profile $E \in \mathbb{E}$, $\mathcal{M}(E) = Cn(\bigcup E)$, when the union $\bigcup E$ is consistent.*

2. Notice that, for sets we use the traditional curly brackets “{” and “}”, whereas, for multi-sets we utilize double curly brackets “{” and “}”.

A union-when-possible merging function produces a belief set resulting from the *set-theoretic union* of all the beliefs of the members of a community, provided that there is *no* inconsistency among these beliefs. Thus, if the two belief sets $K_1 = Cn(\{a\})$ and $K_2 = Cn(\{a \rightarrow b\})$ ought to be aggregated by means of a union-when-possible merging function \mathcal{M} , we would have that $\mathcal{M}(K_1, K_2) = Cn(K_1 \cup K_2) = Cn(\{a, a \rightarrow b\}) = Cn(\{a, b\})$, since the union $K_1 \cup K_2$ is consistent.

Observe that, in case the union $\bigcup E$ is inconsistent, the aggregation implemented by a union-when-possible merging function is *not constrained* at all. Based on this observation, Konieczny and Pérez (2002) identified a *proper* sub-class of the whole class of union-when-possible merging functions, which contains merging operators whose behaviour is constrained by certain rationality-principles. These merging operators are called *Integrity Constraints quasi-merging functions*, are abbreviated as *IC quasi-merging functions*, and are introduced subsequently.

4.1.1 IC QUASI-MERGING FUNCTIONS

In the spirit of the AGM trio, Konieczny and Pérez (2002) characterized the class of IC quasi-merging functions in terms of a collection of *rationality-postulates*, which essentially encode reasonable logical properties of an aggregation-operation.

Konieczny and Pino Pérez consider merging functions whose aggregation-operation adheres to specific *integrity constraints*, which are in turn represented by a logical formula μ . In particular, they define an *Integrity Constraints quasi-merging function*, or *IC quasi-merging function* for short, to be any merging function \mathcal{M} that satisfies the following postulates (IC0)–(IC8).³

- (IC0) $\mathcal{M}_\mu(E)$ is a theory of \mathbb{T} containing μ .
- (IC1) If μ is consistent, then $\mathcal{M}_\mu(E)$ is consistent.
- (IC2) If $(\bigcup E) \cup \{\mu\}$ is consistent, then $\mathcal{M}_\mu(E) = (\bigcup E) + \mu$.
- (IC3) If $E_1 \equiv E_2$ and $Cn(\{\mu_1\}) = Cn(\{\mu_2\})$, then $\mathcal{M}_{\mu_1}(E_1) = \mathcal{M}_{\mu_2}(E_2)$.
- (IC4) If $\mu \in K_1$ and $\mu \in K_2$, then $\mathcal{M}_\mu(\{\{K_1, K_2\}\}) \cup K_1$ is consistent iff $\mathcal{M}_\mu(\{\{K_1, K_2\}\}) \cup K_2$ is consistent.
- (IC5) $\mathcal{M}_\mu(E_1 \sqcup E_2) \subseteq Cn(\mathcal{M}_\mu(E_1) \cup \mathcal{M}_\mu(E_2))$.
- (IC6) If $\mathcal{M}_\mu(E_1) \cup \mathcal{M}_\mu(E_2)$ is consistent, then $Cn(\mathcal{M}_\mu(E_1) \cap \mathcal{M}_\mu(E_2)) \subseteq \mathcal{M}_\mu(E_1 \sqcup E_2)$.
- (IC7) $\mathcal{M}_{\mu_1 \wedge \mu_2}(E) \subseteq \mathcal{M}_{\mu_1}(E) + \mu_2$.
- (IC8) If $\mathcal{M}_{\mu_1}(E) \cup \{\mu_2\}$ is consistent, then $\mathcal{M}_{\mu_1}(E) + \mu_2 \subseteq \mathcal{M}_{\mu_1 \wedge \mu_2}(E)$.

3. We have rephrased Konieczny and Pino Pérez's postulates for IC quasi-merging functions in terms of the notation adopted in this paper.

For a detailed discussion on the above postulates, the interested reader is referred to (Konieczny & Pérez, 2002, 2011); in these works, one can also find the definition of several notable sub-classes of the whole class of IC quasi-merging functions, such as the *IC merging functions*, the *majority merging functions* and the *arbitration functions*.

Konieczny and Pino Pérez *semantically* characterized the class of IC quasi-merging functions in terms of a special kind of total preorders over possible worlds. Their semantic characterization is in the spirit of Katsuno and Mendelzon’s faithful-preorders characterization for belief revision (cf. Subsection 3.2). For presenting this characterization, let us firstly introduce the notion of *quasi-syncretic assignment*.

Definition 6 (Quasi-Syncretic Assignment, (Konieczny & Pérez, 2002)). *A quasi-syncretic assignment is a function that maps each profile E to a total preorder \sqsubseteq_E over \mathbb{M} , such that, for any profiles $E_1, E_2 \in \mathbb{E}$, any belief sets $K_1, K_2 \in \mathbb{T}$, and any possible worlds $r, r' \in \mathbb{M}$, the following conditions hold.⁴*

- (i) *If $r, r' \in \llbracket E \rrbracket$, then $r \simeq_E r'$.*
- (ii) *If $r \in \llbracket E \rrbracket$ and $r' \notin \llbracket E \rrbracket$, then $r \sqsubset_E r'$.*
- (iii) *If $E_1 \equiv E_2$, then $\sqsubseteq_{E_1} = \sqsubseteq_{E_2}$.*
- (iv) *For any $r \in \llbracket K_1 \rrbracket$, there is a world $r' \in \llbracket K_2 \rrbracket$, such that $r' \sqsubseteq_{\{K_1\} \sqcup \{K_2\}} r$.*
- (v) *If $r \sqsubseteq_{E_1} r'$ and $r \sqsubseteq_{E_2} r'$, then $r \sqsubseteq_{E_1 \sqcup E_2} r'$.*
- (vi) *If $r \sqsubset_{E_1} r'$ and $r \sqsubset_{E_2} r'$, then $r \sqsubset_{E_1 \sqcup E_2} r'$.*

As in the case of a faithful preorder, a total preorder \sqsubseteq_E over \mathbb{M} , as defined in Definition 6, encodes the comparative plausibility of the possible worlds of \mathbb{M} , relative to the profile E . Similarly thus, the assertion $r \sqsubseteq_E r'$ states that the world r is at least as plausible as the world r' , with respect to E . Based on the notion of quasi-syncretic assignment, Konieczny and Pino Pérez obtained the following representation theorem.

Theorem 7 (Konieczny & Pérez, 2002). *A merging function \mathcal{M} satisfies postulates (IC0)–(IC8) iff there exists a quasi-syncretic assignment that maps each profile E to a total preorder \sqsubseteq_E over \mathbb{M} , such that, for any $\mu \in \mathcal{L}$:*

$$(M) \quad \llbracket \mathcal{M}_\mu(E) \rrbracket = \min(\llbracket \mu \rrbracket, \sqsubseteq_E).$$

According to condition (M), the resulting merged belief set $\mathcal{M}_\mu(E)$ is defined in terms of the most plausible (with respect to E) possible worlds satisfying μ , a scheme that strongly resembles the way that belief sets are revised according to condition (R) (cf. Section 3).

As earlier mentioned, we assume herein (with an exception to be made in Section 8) that the merging-process implemented by merging functions does *not* adhere to integrity constraints. On that premise, the sentence μ is a tautology and, thus, the merged belief set $\mathcal{M}_\mu(E)$, which can be denoted by $\mathcal{M}(E)$, is always *consistent*, in view of postulate (IC1). Furthermore, assuming that the sentence μ is a tautology, the following remark follows immediately from postulate (IC2).

4. \sqsubset_E and \simeq_E denote the strict and indifference part of the preorder \sqsubseteq_E , respectively.

Remark 8. *The class of IC quasi-merging functions is a sub-class of the class of union-when-possible merging functions.*

4.1.2 FORMULA-BASED MERGING FUNCTIONS

The class of union-when-possible merging functions, beyond the IC quasi-merging functions, *subsumes* other central types of merging operators as well, such as the *formula-based* merging functions (Baral et al., 1992; Benferhat, Dubois, Lang, Prade, Saffiotti, & Smets, 1998; Konieczny, 2000). Contrary to IC quasi-merging functions which are *model-based* operators defined in terms of possible worlds (models), formula-based merging functions are sensitive to the *formulae* of the belief sets to be aggregated, as they select some *maximal consistent* subsets of the union of these belief sets. In particular, consider the following merging operator.⁵

$$\mathcal{M}^{\subseteq}(E) = \bigcap \left\{ M \subseteq Cn(\bigcup E) : M \neq \mathcal{L}, \text{ and if } M \subset M' \subseteq Cn(\bigcup E), \text{ then } M' = \mathcal{L} \right\}$$

The operator \mathcal{M}^{\subseteq} is a formula-based merging function that yields the intersection of *every inclusion-maximal consistent* subset of the belief set $Cn(\bigcup E)$. A merging operator analogous to \mathcal{M}^{\subseteq} can also be defined when maximality is specified in terms of *cardinality*.⁶

$$\mathcal{M}^{card}(E) = \bigcap \left\{ M \subseteq Cn(\bigcup E) : M \neq \mathcal{L}, \text{ and if } |M| < |M'| \text{ and } M' \subseteq Cn(\bigcup E), \text{ then } M' = \mathcal{L} \right\}$$

Both merging operators \mathcal{M}^{\subseteq} and \mathcal{M}^{card} are defined by intersecting *all* the consistent subsets of the belief set $Cn(\bigcup E)$ that are maximal, either with respect to set inclusion or with respect to cardinality. One may argue that this scheme is not satisfactory from a merging point of view, as it does not take into account the distribution of the information among the belief sets to be aggregated. In response to this weakness, Konieczny (2000) proposed the use of particular *selection mechanisms* that pick up for intersection only the maximal consistent subsets of $Cn(\bigcup E)$ that best fit a certain merging-criterion. These selection mechanisms, which are inspired by the so-called *selection functions* of Alchourrón et al. (1985) by means of which partial-meet contraction is defined, allowed the definition of a wide *class* of formula-based merging functions with better behaviour and, thus, better logical properties (cf. (Konieczny, 2000) for details).

Now, it is easy to verify that, for any profile $E \in \mathbb{E}$, if the union $\bigcup E$ is *consistent*, then *every* formula-based merging function of the aforementioned class of merging operators yields the belief set $Cn(\bigcup E)$. This observation implies the following remark.

Remark 9. *The class of formula-based merging functions is a sub-class of the class of union-when-possible merging functions.*

We note lastly that *none* of the formula-based merging function considered herein respects postulate (IC3) (Konieczny, 2000). It follows then that *no* formula-based merging function that we study is an IC quasi-merging function, a fact which brings us to the following conclusion.

5. Contrary to the original works on formula-based merging functions (Baral et al., 1992; Benferhat et al., 1998), herein we assume that such operators do *not* adhere to integrity constraints.

6. For a set V , $|V|$ denotes the cardinality of V .

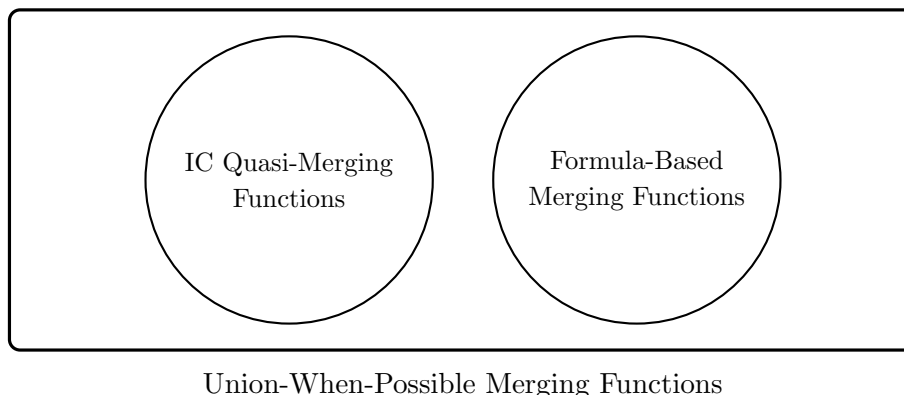


Figure 1: The class of union-when-possible merging functions relative to the classes of IC quasi-merging functions and formula-based merging functions (cf. Remarks 8, 9 and 10).

Remark 10. *The class of formula-based merging functions is disjoint with the class of IC quasi-merging functions.*

Figure 1 depicts the class of union-when-possible merging functions relative to the classes of IC quasi-merging functions and formula-based merging functions, in view of Remarks 8, 9 and 10.

4.2 Other Principal Merging Operators

In the remainder of this section, we introduce three other principal types of merging operators. We begin with the *consensus* merging function, which is a *unique* merging operator also considered by Westlund (2010).

Definition 11 (Consensus Merging Function). *A merging function \mathcal{M} is the consensus merging function iff, for any profile $E \in \mathbb{E}$, $\mathcal{M}(E) = \bigcap E$.*

The consensus merging function produces the consensus among *all* the members of a community. Thus, if the two belief sets $K_1 = Cn(\{a, b\})$ and $K_2 = Cn(\{a, \neg b\})$ ought to be aggregated by means of the consensus merging function \mathcal{M} , we would have that $\mathcal{M}(K_1, K_2) = K_1 \cap K_2 = Cn(\{a\})$.⁷

We shall now elaborate on the second type of merging operators, which we refer to as *radical* merging functions.

Definition 12 (Radical Merging Function). *A merging function \mathcal{M} is a radical merging function iff, for any profile $E \in \mathbb{E}$ and any non-tautological sentence $\varphi \in \mathcal{L}$ such that $\varphi \in \bigcap E$, $\varphi \notin \mathcal{M}(E)$.*

7. It is noteworthy that the consensus merging function respects postulate (Disj) of Everaere, Konieczny, and Marquis (2010, p. 828).

Intuitively, a merging function \mathcal{M} is radical whenever, for every profile E , any non-tautological belief that is *common* to all belief sets of E is *not* contained in the aggregated belief set $\mathcal{M}(E)$. Due to their “reactive” operation, radical merging functions may be considered unreasonable operators for combining information; yet, there are certain scenarios in which their outcome can be perfectly suitable. Consider, for instance, a profile E which is specified by the two belief sets $K_1 = Cn(\{a, \neg b\})$ and $K_2 = Cn(\{a \rightarrow b, \neg b\})$. Then, there is a radical merging function \mathcal{M} that yields the belief set $\mathcal{M}(E) = Cn(\{b\})$, which does not contain non-tautological sentences that are common to both K_1 and K_2 .⁸ Arguably, the belief set $\mathcal{M}(E)$ represents a plausible synthesis of the belief sets K_1 and K_2 .

Finally, let us introduce the *singular* merging functions, which represent the last type of merging operators introduced herein.

Definition 13 (Singular Merging Function). *A merging function \mathcal{M} is a singular merging function iff, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$, there is an $i \in \{1, \dots, n\}$ such that $\mathcal{M}(K_1, \dots, K_n) = K_i$.*

A singular merging function yields a *single unaltered* belief set of the profile it receives as input. Singular merging functions could find application in situations where it is preferable to choose one among all the possible states of belief, rather than to melt all the available states of belief in a single belief set. Suppose, for example, that three physicians propose three different medical treatments for the condition of a patient. It seems plausible that the patient should receive one of the proposed treatments, and not a mixture of them.

A special kind of singular merging function is a *dictator* merging function, also considered by Westlund (2010).

Definition 14 (Dictator Merging Function). *Let n be an arbitrary, but fixed, positive integer. A (singular) merging function \mathcal{M} is a dictator merging function iff there is an $i \in \{1, \dots, n\}$ such that, for any n -elements profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$, $\mathcal{M}(K_1, \dots, K_n) = K_i$.*

Given an arbitrary, but fixed, positive integer n , a dictator merging function always chooses the belief set corresponding to a specific *fixed* index $i \in \{1, \dots, n\}$, irrespectively of the (n -elements) profile $\{\{K_1, \dots, K_n\}\}$ that the merging function receives as input.

5. The Principle of Collective Change

As stated in the Introduction section, Westlund introduced the *Principle of Collective Change* (PCC) in the realm of contraction and expansion (Westlund, 2010, p. 217). This section is devoted to the formulation of the PCC in the realm of *genuine belief revision*; namely, in the *principal* case where the new information *contradicts* the belief set of each individual of a collective entity.

Evidently, a *community* of individuals could be an organization, a company, a scientific committee, or any multi-agent system. A formal definition of the notion of a community that consists of rational agents is presented in Definition 15.

8. Informally, one can think of the merged belief set $\mathcal{M}(E)$ as being generated as follows: The interaction of the belief a of the belief set K_1 with the belief $a \rightarrow b$ of the belief set K_2 produced the belief b of $\mathcal{M}(E)$, and, at the same time, cancelled the belief $\neg b$ that is common to K_1 and K_2 .




Figure 2: A community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ satisfies the PCC iff there exists an AGM revision function $*$ such that, for any profile $\{K_1, \dots, K_n\}$ and any sentence φ that contradicts each belief set of $\{K_1, \dots, K_n\}$, $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * _1 \varphi, \dots, K_n * _n \varphi)$.

Definition 15 (Community of Rational Individuals). *A community of n rational agents is a tuple $\langle \mathcal{M}, *_1, \dots, *_n \rangle$, where n is an arbitrary positive integer, \mathcal{M} is a merging function, and $*_1, \dots, *_n$ are the AGM revision functions that the agents $1, \dots, n$ employ, respectively, to implement revision.*

Each AGM revision function $*_i$, where $i \in \{1, \dots, n\}$, of a community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ is employed by the agent i in order to rationally respond to new epistemic input, whereas, the merging function \mathcal{M} is used by the whole community in order to aggregate the beliefs of its agents, and formulate its collective knowledge. Based on this definition of collective agency, a belief profile $\{K_1, \dots, K_n\}$ essentially accommodates the atomic states of belief of the n members of the community, and the belief set $\mathcal{M}(K_1, \dots, K_n)$ represents the state of collective belief of the community.

On that premise, we shall say that a community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ of n rational agents *satisfies* Westlund’s PCC iff there exists an AGM revision function $*$ such that, for any profile $\{K_1, \dots, K_n\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, the following condition (CC) holds:

$$(CC) \quad \mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * _1 \varphi, \dots, K_n * _n \varphi).$$

Obviously, the community *violates* Westlund’s PCC iff, for some profile $\{K_1, \dots, K_n\} \in \mathbb{E}$ and some sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, there is *no* AGM revision function $*$ such that $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * _1 \varphi, \dots, K_n * _n \varphi)$.

The PCC —whose essence is graphically represented in Figure 2— constitutes a criterion that characterizes the revision of collective knowledge. This criterion concerns the case of *genuine* revision (that does not reduce to expansion), and requires that, if each rational agent of a community revises her beliefs by an epistemic input that *contradicts* her initial belief set, then the community as a whole ought to revise its collective beliefs

by the *same* epistemic input as well. More precisely, the satisfaction of PCC by a community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ boils down to the existence of an AGM revision function $*$, such that, for any profile $\{\{K_1, \dots, K_n\}\}$ and any sentence φ that contradicts each belief set of $\{\{K_1, \dots, K_n\}\}$, the $*$ -revision of the belief set $\mathcal{M}(K_1, \dots, K_n)$ by φ coincides with the belief set $\mathcal{M}(K_1 *_1 \varphi, \dots, K_n *_n \varphi)$. The AGM revision function $*$ (if it exists) can be seen as the revision operator that the community employs as a collective entity, in order to modify its collective knowledge.⁹

6. Unsatisfiability Results

Having introduced the PCC in the realm of genuine belief revision, we establish in this section the alluded unsatisfiability results concerning this property. We begin with the case of union-when-possible merging functions and their specializations IC quasi-merging functions and formula-based merging functions (Subsection 6.1), we then turn to the consensus merging function (Subsection 6.2), and lastly we investigate the case of radical merging functions (Subsection 6.3).

6.1 Union-When-Possible Merging Functions

Our first unsatisfiability result is Theorem 16, which concerns the case of *union-when-possible* merging functions, and proves that there exist communities of rational individuals that, for *any* union-when-possible merging function, *violate* the PCC.

Theorem 16. *There exist AGM revision functions $*_1, \dots, *_n$, such that, for any union-when-possible merging function \mathcal{M} , the community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ violates the PCC.*

Proof. Assume the non-trivial case where \mathcal{L} contains at least two atoms; hence, there are at least four possible worlds in \mathbb{M} . Let w_1, w_2, w_3, w_4 be four distinct worlds of \mathbb{M} , and let K_1, K_2 be two theories such that $\llbracket K_1 \rrbracket = \{w_1\}$ and $\llbracket K_2 \rrbracket = \{w_1, w_2\}$. Moreover, let φ be a sentence of \mathcal{L} such that $\llbracket \varphi \rrbracket = \{w_3, w_4\}$; clearly, φ contradicts both K_1 and K_2 , as $\llbracket \varphi \rrbracket \cap \llbracket K_1 \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket K_2 \rrbracket = \emptyset$.

Consider an AGM revision function $*_1$ that assigns (via (R)) at the belief set K_1 a faithful preorder $\preceq_{K_1}^{*1}$, such that $w_3 \approx_{K_1}^{*1} w_4$. Moreover, consider an AGM revision function $*_2$ that assigns (via (R)) at the belief set K_1 a faithful preorder $\preceq_{K_1}^{*2}$, such that $w_3 \prec_{K_1}^{*2} w_4$, and at the belief set K_2 a faithful preorder $\preceq_{K_2}^{*2}$, such that $w_4 \prec_{K_2}^{*2} w_3$. Then, it follows from condition (R) that $\llbracket K_1 *_1 \varphi \rrbracket = \{w_3, w_4\}$, $\llbracket K_1 *_2 \varphi \rrbracket = \{w_3\}$ and $\llbracket K_2 *_2 \varphi \rrbracket = \{w_4\}$ (the belief sets mentioned above in this proof are graphically represented in Figure 3).

Now, let \mathcal{M} be an arbitrary union-when-possible merging function. Since $\llbracket K_1 \rrbracket \cap \llbracket K_2 \rrbracket = \{w_1\}$, we derive that $\mathcal{M}(K_1, K_1) = \mathcal{M}(K_1, K_2) = K_1$. Furthermore, since $\llbracket K_1 *_1 \varphi \rrbracket \cap \llbracket K_1 *_2 \varphi \rrbracket = \{w_3\}$ and $\llbracket K_1 *_1 \varphi \rrbracket \cap \llbracket K_2 *_2 \varphi \rrbracket = \{w_4\}$, we have that

9. In the context of the space-related scenario stated in the Introduction section, the AGM revision functions $*_1, \dots, *_n$ are employed by the n experts of the scientific community for *individually* revising their current atomic states of knowledge, which are represented by the belief sets K_1, \dots, K_n . The merging function \mathcal{M} is employed by the scientific community so that its *collective* knowledge be generated. Then, assuming that φ represents the new scientific discovery, the belief sets $\mathcal{M}(K_1, \dots, K_n)$ and $\mathcal{M}(K_1 *_1 \varphi, \dots, K_n *_n \varphi)$ correspond, respectively, to the initial and revised states of collective knowledge of the scientific community.




Figure 3: Graphical representation of the belief sets K_1 and K_2 and their revisions by φ , via the AGM revision functions $*_1$ and $*_2$, described in the proof of Theorem 16. The marked areas represent the sets of all possible worlds satisfying the corresponding belief sets.

$\llbracket \mathcal{M}(K_1 *_1 \varphi, K_1 *_2 \varphi) \rrbracket = \llbracket \text{Cn}((K_1 *_1 \varphi) \cup (K_1 *_2 \varphi)) \rrbracket = \llbracket K_1 *_1 \varphi \rrbracket \cap \llbracket K_1 *_2 \varphi \rrbracket = \{w_3\}$
 and $\llbracket \mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi) \rrbracket = \llbracket \text{Cn}((K_1 *_1 \varphi) \cup (K_2 *_2 \varphi)) \rrbracket = \llbracket K_1 *_1 \varphi \rrbracket \cap \llbracket K_2 *_2 \varphi \rrbracket = \{w_4\}$,
 respectively. Obviously then, it follows that $\mathcal{M}(K_1 *_1 \varphi, K_1 *_2 \varphi) \neq \mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi)$.

Next, suppose towards contradiction that the community $\langle \mathcal{M}, *_1, *_2 \rangle$ of two rational agents satisfies the PCC. Therefore, there exists an AGM revision function $*$ such that $\mathcal{M}(K_1, K_1) * \varphi = \mathcal{M}(K_1 *_1 \varphi, K_1 *_2 \varphi)$ and $\mathcal{M}(K_1, K_2) * \varphi = \mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi)$. Given that $\mathcal{M}(K_1 *_1 \varphi, K_1 *_2 \varphi) \neq \mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi)$, we derive that $\mathcal{M}(K_1, K_1) * \varphi \neq \mathcal{M}(K_1, K_2) * \varphi$. This, however, contradicts the fact that $\mathcal{M}(K_1, K_1) = \mathcal{M}(K_1, K_2)$. Consequently, for any union-when-possible merging function \mathcal{M} , the community $\langle \mathcal{M}, *_1, *_2 \rangle$ violates the PCC. ■

Given that union-when-possible merging functions subsume IC quasi-merging functions and formula-based merging functions (cf. Remarks 8 and 9), the following important corollary is an immediate consequence of Theorem 16.

Corollary 17. *There exist AGM revision functions $*_1, \dots, *_n$, such that, for any IC quasi-merging function or any formula-based merging function \mathcal{M} , the community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ violates the PCC.*

Corollary 17 points out that there exist communities of rational agents (who behave rationally at the individual level as each agent employs an AGM revision function), such that, for *any* IC quasi-merging function or formula-based merging function, do *not* adhere to the PCC.

6.2 The Consensus Merging Function

Let us now turn to the case of the *consensus* merging function. It turns out that an unsatisfiability result can be obtained in that circumstance too, indicating that there exist communities of rational individuals, equipped with the consensus merging function, that do *not* respect the PCC.

Theorem 18. *Let \mathcal{M} be the consensus merging function. There exist AGM revision functions $*_1, \dots, *_n$, such that the community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ violates the PCC.*

Proof. Assume the non-trivial case where \mathcal{L} contains at least two atoms; hence, there are at least four possible worlds in \mathbb{M} . Let w_1, w_2, w_3, w_4 be four distinct worlds of \mathbb{M} , let K_1, K_2 be two theories such that $\llbracket K_1 \rrbracket = \{w_1\}$ and $\llbracket K_2 \rrbracket = \{w_2\}$, and let φ be a sentence of \mathcal{L} such that $\llbracket \varphi \rrbracket = \{w_3, w_4\}$. Clearly, φ contradicts both K_1 and K_2 , as $\llbracket \varphi \rrbracket \cap \llbracket K_1 \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket K_2 \rrbracket = \emptyset$.

Consider an AGM revision function $*_1$ that assigns (via (R)) at the belief sets K_1 and K_2 two faithful preorders $\preceq_{K_1}^{*1}$ and $\preceq_{K_2}^{*1}$, respectively, such that $w_3 \prec_{K_1}^{*1} w_4$ and $w_4 \prec_{K_2}^{*1} w_3$. Moreover, consider an AGM revision function $*_2$ that assigns (via (R)) at the belief sets K_1 and K_2 two faithful preorders $\preceq_{K_1}^{*2}$ and $\preceq_{K_2}^{*2}$, respectively, such that $w_4 \prec_{K_1}^{*2} w_3$ and $w_3 \prec_{K_2}^{*2} w_4$. Then, it follows from condition (R) that $\llbracket K_1 *_1 \varphi \rrbracket = \{w_3\}$ and $\llbracket K_2 *_1 \varphi \rrbracket = \{w_4\}$, and moreover, $\llbracket K_1 *_2 \varphi \rrbracket = \{w_4\}$ and $\llbracket K_2 *_2 \varphi \rrbracket = \{w_3\}$.

Now, since \mathcal{M} is the consensus merging function, it follows that $\mathcal{M}(K_1, K_2) = \mathcal{M}(K_2, K_1) = K_1 \cap K_2$. Furthermore, it is true that $\mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi) = (K_1 *_1 \varphi) \cap (K_2 *_2 \varphi)$ and $\mathcal{M}(K_2 *_1 \varphi, K_1 *_2 \varphi) = (K_2 *_1 \varphi) \cap (K_1 *_2 \varphi)$, a fact which entails, respectively, that $\llbracket \mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi) \rrbracket = \llbracket K_1 *_1 \varphi \rrbracket \cup \llbracket K_2 *_2 \varphi \rrbracket = \{w_3\}$ and $\llbracket \mathcal{M}(K_2 *_1 \varphi, K_1 *_2 \varphi) \rrbracket = \llbracket K_2 *_1 \varphi \rrbracket \cup \llbracket K_1 *_2 \varphi \rrbracket = \{w_4\}$. Obviously then, it follows that $\mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi) \neq \mathcal{M}(K_2 *_1 \varphi, K_1 *_2 \varphi)$.

Next, suppose towards contradiction that the community $\langle \mathcal{M}, *_1, *_2 \rangle$ of two rational agents satisfies the PCC. Therefore, there exists an AGM revision function $*$ such that $\mathcal{M}(K_1, K_2) * \varphi = \mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi)$ and $\mathcal{M}(K_2, K_1) * \varphi = \mathcal{M}(K_2 *_1 \varphi, K_1 *_2 \varphi)$. Given that $\mathcal{M}(K_1 *_1 \varphi, K_2 *_2 \varphi) \neq \mathcal{M}(K_2 *_1 \varphi, K_1 *_2 \varphi)$, we derive that $\mathcal{M}(K_1, K_2) * \varphi \neq \mathcal{M}(K_2, K_1) * \varphi$. This, however, contradicts the fact that $\mathcal{M}(K_1, K_2) = \mathcal{M}(K_2, K_1)$. Consequently, for the consensus merging function \mathcal{M} , the community $\langle \mathcal{M}, *_1, *_2 \rangle$ violates the PCC. \blacksquare

It should be noted at this point that Westlund, confined in the realm of contraction and expansion, also established impossibility results for union-when-possible merging functions and the consensus merging function. In particular, he showed that *no* union-when-possible or consensus merging function respects the PCC for the operations of contraction and expansion (Westlund, 2010, Observations 6 and 10). Theorems 16 and 18 and Corollary 17, essentially, come to *extend* in the realm of belief revision the aforementioned Westlund's impossibility results.

6.3 Radical Merging Functions

We close this section with the case of *radical* merging functions, which also lead to communities not adhering to the PCC. This is shown in Theorem 19, which proves that *every* community of rational individuals equipped with a radical merging function *violates* the PCC.

Theorem 19. *Every community of rational agents equipped with a radical merging function violates the PCC.*

Proof. Let $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ be a community of n rational agents, such that \mathcal{M} is a radical merging function. Let $\{\{K_1, \dots, K_n\}\}$ be a profile of \mathbb{E} , and let φ be a non-tautological

sentence of \mathcal{L} such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$. From postulate $(K*2)$, it follows that $\varphi \in K_i *_i \varphi$, for all $i \in \{1, \dots, n\}$. Therefore, given that \mathcal{M} is a radical merging function, we derive that $\varphi \notin \mathcal{M}(K_1 *_1 \varphi, \dots, K_n *_n \varphi)$. Postulate $(K*2)$ also entails that, for *any* AGM revision function $*'$, $\varphi \in \mathcal{M}(K_1, \dots, K_n) *' \varphi$. Hence, there is *no* AGM revision function $*$ such that $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 *_1 \varphi, \dots, K_n *_n \varphi)$. Consequently, the community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ violates the PCC. ■

7. Satisfiability Results

The impossibility results established in the previous section are, certainly, not a good indication for the suitability of the PCC, and some notable implications of them shall be discussed later in this article. Still, in the present section, it is shown that it is possible to formulate a collection of satisfiability results too concerning the PCC. The alluded satisfiability results relate to the cases of *union-when-possible*, *consensus* and *singular* merging functions, and shall be established with the aid of a special kind of belief revision, called *uniform belief revision*. The process of uniform belief revision has been developed by Areces and Becher (2001), and was subsequently investigated by Aravanis (2020).¹⁰ We proceed subsequently to the formal definition of this kind of belief revision.

7.1 Uniform Belief Revision

Let \preceq be an arbitrary, but *fixed*, total preorder over the set of all possible worlds \mathbb{M} . Intuitively, the total preorder \preceq —which is *belief-set-independent*—encodes the (state-independent) comparative plausibility of possible worlds, which in turn reflects the *dynamics* of a domain, imposed by a body of pre-established background knowledge, such as objective laws of physics, legal or moral codes. In view of condition (US) presented below, the total preorder \preceq can generate a whole *family* $\{\preceq_K\}_{K \in \mathbb{T}}$ of faithful preorders over worlds, that is, one total preorder associated with *each* belief set of \mathbb{T} .

$$(US) \quad \text{For any possible worlds } r, r' \notin \llbracket K \rrbracket, \quad r \preceq_K r' \quad \text{iff} \quad r \preceq r'.$$

According to condition (US), for each belief set $K \in \mathbb{T}$, the faithful preorder \preceq_K is set *identical* to the total preorder \preceq , on their sub-domain $(\mathbb{M} - \llbracket K \rrbracket) \times (\mathbb{M} - \llbracket K \rrbracket)$. In this way, the total preorder \preceq that encodes the domain's dynamics generates a family $\{\preceq_K\}_{K \in \mathbb{T}}$ of faithful preorders, and thus, \preceq generates the *revision-strategy* for every belief set of the language.

Now, the family $\{\preceq_K\}_{K \in \mathbb{T}}$ of faithful preorders induces, through condition (R), an AGM revision function. An AGM revision function so constructed is called *uniform revision operator*, or *UR operator* for short. As proved by Areces and Becher (2001), the proper

10. Areces and Becher named the revision operators that implement uniform belief revision “iterable AGM revision functions”, as the authors focus on the favourable properties of these operators related to iterated revision (Peppas, 2014). In this article, following Aravanis (2020), we adopt the term “uniform revision operators” which reflects the way the operators act on belief sets.

sub-class of AGM revision functions satisfying the subsequent postulate (UR) is precisely the class of UR operators.¹¹

(UR) For any sentence $\neg\varphi \in K_1 \cap K_2$, $K_1 * \varphi = K_2 * \varphi$.

Hence, given a UR operator $*$, any two belief sets of the language, which are both contradicted by an epistemic input φ , are $*$ -revised by φ in a *uniform manner*. As we show in the remainder of this section, this property of uniform belief revision will help us to establish our satisfiability results concerning the PCC.

7.2 Union-When-Possible Merging Functions

Against this background, we establish our first satisfiability result related to the PCC, which is Theorem 20 that refers to *union-when-possible* merging functions, and points out that the *uniformity* of the revision-process (as implemented by a UR operator) guarantees the fulfilment of the PCC.

Theorem 20. *Let $*$ be a UR operator. Moreover, let \mathcal{M} be a union-when-possible merging function such that, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, $\varphi \in \mathcal{M}(K_1, \dots, K_n)$. Then, the community $\langle \mathcal{M}, *, \dots, * \rangle$ of n rational agents satisfies the PCC.*

Proof. To prove that the community $\langle \mathcal{M}, *, \dots, * \rangle$ of n rational agents satisfies the PCC, it suffices to show that, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$.

To that end, let $\{\{K_1, \dots, K_n\}\}$ be an arbitrary profile of \mathbb{E} , and let φ be a sentence of \mathcal{L} such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$. Firstly, by the consideration (hypothesis) of \mathcal{M} , we have that $\neg\varphi \in \mathcal{M}(K_1, \dots, K_n)$. Furthermore, observe that, since $*$ is a UR operator, it follows from postulate (UR) that $K_j * \varphi = K_k * \varphi = \mathcal{M}(K_1, \dots, K_n) * \varphi = H$, for all $j, k \in \{1, \dots, n\}$. From our assumption that the epistemic input φ is consistent (cf. Section 3), postulate ($K * 5$) entails that $K_j * \varphi$ is consistent, for all $j \in \{1, \dots, n\}$. Then, the fact that $K_j * \varphi = K_k * \varphi$, for all $j, k \in \{1, \dots, n\}$, entails that the union $\bigcup_{i=1}^n (K_i * \varphi)$ is consistent as well. Since \mathcal{M} is a union-when-possible merging function, we derive then that $\mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi) = \bigcup_{i=1}^n (K_i * \varphi) = \underbrace{H \cup \dots \cup H}_n = H$. Combining the above, we obtain that $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$, as desired. \blacksquare

Theorem 20 assumes a union-when-possible merging function \mathcal{M} such that the belief set $\mathcal{M}(K_1, \dots, K_n)$ resulting from the \mathcal{M} -aggregation of any belief sets K_1, \dots, K_n is *contradicted* by the epistemic input φ .¹² It proves, then, that a community equipped with *any*

11. A concrete example that outlines the construction of a UR operator is Example 4.2 by Aravanis (2020, p. 1363). Furthermore, an interesting *generalization* of UR operators has been recently formulated by Aravanis and Peppas (2022).

12. This means that the revision of $\mathcal{M}(K_1, \dots, K_n)$ by φ is genuine, in the sense that revision does *not* reduce to expansion.

such union-when-possible merging function (thus, with any such IC quasi-merging function or formula-based merging function) satisfies the PCC, provided that all its members (individually) revise their beliefs in a uniform manner, by means of a UR operator.

7.3 The Consensus Merging Function

The second satisfiability result concerning the PCC that we establish is Theorem 21, stated below. Similarly to Theorem 20, Theorem 21 proves that any community equipped with the consensus merging function respects the PCC, provided that its agents uniformly revise their beliefs.

Theorem 21. *Let $*$ be a UR operator, and let \mathcal{M} be the consensus merging function. Then, the community $\langle \mathcal{M}, *, \dots, * \rangle$ of n rational agents satisfies the PCC.*

Proof. To prove that the community $\langle \mathcal{M}, *, \dots, * \rangle$ of n rational agents satisfies the PCC, it suffices to show that, for any profile $\{\{K_1, \dots, K_n\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$, or equivalently, $\left(\bigcap_{i=1}^n K_i\right) * \varphi = \bigcap_{i=1}^n (K_i * \varphi)$.

To that end, let $\{\{K_1, \dots, K_n\} \in \mathbb{E}$, and let φ be a sentence of \mathcal{L} such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$. Firstly, notice that $\neg\varphi \in \left(\bigcap_{i=1}^n K_i\right)$. Furthermore, observe that, since $*$ is a UR operator, it follows from postulate (UR) that $K_j * \varphi = K_k * \varphi = \left(\bigcap_{i=1}^n K_i\right) * \varphi = H$, for all $j, k \in \{1, \dots, n\}$. Therefore, we derive that $\bigcap_{i=1}^n (K_i * \varphi) = \underbrace{H \cap \dots \cap H}_n = H$. Combining the above, we obtain that $\left(\bigcap_{i=1}^n K_i\right) * \varphi = \bigcap_{i=1}^n (K_i * \varphi)$, as desired. ■

In contrast to Theorems 16 and 18 of the previous section which are impossibility results concerning union-when-possible merging functions and the consensus merging function, Theorems 20 and 21 indicate that, under certain conditions, the satisfaction of the PCC is indeed *feasible* for the aforementioned types of merging operators.

It is noteworthy that Westlund (2010) did *not* establish satisfiability results for union-when-possible merging functions and the consensus merging function — although, as stated in the previous section, he did establish impossibility results for these types of merging operators.

7.4 Singular Merging Functions

Westlund (2010, Observation 5) has already shown that, for any profile of \mathbb{E} , every dictator merging function respects the PCC, for the operations of contraction and expansion. Theorem 22 shows that this is also the case for the operation of revision, by proving that every community equipped with a dictator merging function respects the PCC.

Theorem 22. *Every community of rational agents equipped with a dictator merging function satisfies the PCC.*

Proof. Let $\langle \mathcal{M}, *_{1}, \dots, *_{n} \rangle$ be a community of n rational agents, such that \mathcal{M} is a dictator merging function, and let $\{\{K_1, \dots, K_n\}\}$ be a profile of \mathbb{E} . Since \mathcal{M} is a dictator merging function and n is fixed, it follows that there is an $i \in \{1, \dots, n\}$ such that $\mathcal{M}(K_1, \dots, K_n) = K_i$. Furthermore, for any sentence $\varphi \in \mathcal{L}$, it is also true that $\mathcal{M}(K_1 *_{1} \varphi, \dots, K_n *_{n} \varphi) = K_i *_{i} \varphi$. Combining the above, we derive that there exists an AGM revision function, and that is $*_{i}$, such that, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_j$, for all $j \in \{1, \dots, n\}$, $\mathcal{M}(K_1, \dots, K_n) *_{i} \varphi = \mathcal{M}(K_1 *_{1} \varphi, \dots, K_n *_{n} \varphi)$. Therefore, the community $\langle \mathcal{M}, *_{1}, \dots, *_{n} \rangle$ satisfies the PCC, as desired. ■

It turns out that a result analogous to Theorem 22 can be obtained for communities equipped with the more general *singular* merging functions, provided that the individual agents of such a community employ a UR operator for implementing revision. This is shown in Theorem 23, stated below.

Theorem 23. *Let $*$ be a UR operator, and let \mathcal{M} be a singular merging function. Then, the community $\langle \mathcal{M}, *, \dots, * \rangle$ of n rational agents satisfies the PCC.*

Proof. To prove that the community $\langle \mathcal{M}, *, \dots, * \rangle$ of n rational agents satisfies the PCC, it suffices to show that, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$.

To that end, let $\{\{K_1, \dots, K_n\}\}$ be an arbitrary profile of \mathbb{E} , and let φ be a sentence of \mathcal{L} such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$. Since \mathcal{M} is a singular merging function, it follows that there is a $j \in \{1, \dots, n\}$ such that $\mathcal{M}(K_1, \dots, K_n) = K_j$. Furthermore, it is also true that there is a $k \in \{1, \dots, n\}$ such that $\mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi) = K_k * \varphi$.¹³ Given that $*$ is a UR operator, we derive from postulate (UR) that $K_j * \varphi = K_k * \varphi$. Combining the above, we have that $\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$, as desired. ■

7.5 Comments on the Satisfiability Results

Let us now summarize our main satisfiability results. Both Theorems 21 and 23 show that the consensus merging function or *any* singular merging function, combined with *any* UR operator, can induce a community adhering to the dictates of the PCC. Essentially, these two results indicate that a community equipped with any of the aforementioned types of merging operators fulfil the PCC, if both the community itself and all its members revise their beliefs with the exact *same* strategy. This can in turn be achieved when both the community and its members perceive the dynamics of the domain in the same way, as encoded into a *single* total preorder \preceq over possible worlds.¹⁴ Theorem 20 shows that the previous observations apply to the case of union-when-possible merging functions as well, when some weak constraints are imposed on their operation.

Against the background of the established satisfiability results, the following concrete example illustrates representative scenarios of collective belief revision that respect the PCC.

13. The indices j and k are in general distinct.

14. The uniform perception of the dynamics of the domain could be possible in an homogeneous community (probably, for example, within a small scientific committee); yet, in diverge communities such a requirement is undoubtedly strong.

Example 24 (Satisfiability of the PCC). Let $\mathcal{P} = \{a, b, c\}$, let K_1, K_2, K_3 be three belief sets such that $K_1 = \text{Cn}(\{a, b, c\})$, $K_2 = \text{Cn}(\{\neg a, b, c\})$ and $K_3 = \text{Cn}(\{a, \neg b, c\})$, and let φ be a sentence of \mathcal{L} such that $\varphi = \neg c$. Notice that φ contradicts each one of K_1, K_2 and K_3 . Moreover, let $*$ be a UR operator induced, through conditions (US) and (R), from the following total preorder \preceq over \mathbb{M} .¹⁵

$$\begin{array}{ccccccc} abc & & & & ab\bar{c} & & \\ \bar{a}bc & \prec & \bar{a}b\bar{c} & \prec & \bar{a}\bar{b}\bar{c} & \prec & \bar{a}bc \\ & & & & a\bar{b}c & & \\ & & & & \bar{a}bc & & \end{array}$$

Then, it follows from conditions (US) and (R) that $K_1 * \varphi = K_2 * \varphi = K_3 * \varphi = \text{Cn}(\{\neg a, b, \neg c\})$.

Lastly, let $\langle \mathcal{M}, *, *, * \rangle$ be a community (e.g., a scientific committee) of 3 rational agents, where \mathcal{M} can be either a union-when-possible merging function, or a consensus merging function, or a singular merging function.

Assume, firstly, that \mathcal{M} is a union-when-possible merging function such that, for any profile $\{K_1, \dots, K_n\} \in \mathbb{E}$ and any sentence $\psi \in \mathcal{L}$ such that $\psi \in K_i$, for all $i \in \{1, \dots, n\}$, $\psi \in \mathcal{M}(K_1, \dots, K_n)$. In view of the fact that $K_1 \cup K_2 \cup K_3$ is inconsistent, suppose that \mathcal{M} is such that $\mathcal{M}(K_1, K_2, K_3) = \text{Cn}(\{(a \wedge b \wedge c) \vee (\neg a \wedge b \wedge c) \vee (a \wedge \neg b \wedge c)\})$. Expectedly, $\neg\varphi \in \mathcal{M}(K_1, K_2, K_3)$, and therefore, due to (US) and (R), we derive that $\mathcal{M}(K_1, K_2, K_3) * \varphi = \text{Cn}(\{\neg a, b, \neg c\})$. Since \mathcal{M} is a union-when-possible merging function, it also follows that $\mathcal{M}(K_1 * \varphi, K_2 * \varphi, K_3 * \varphi) = (K_1 * \varphi) \cup (K_2 * \varphi) \cup (K_3 * \varphi) = \text{Cn}(\{\neg a, b, \neg c\})$. Observe, then, that $\mathcal{M}(K_1, K_2, K_3) * \varphi = \mathcal{M}(K_1 * \varphi, K_2 * \varphi, K_3 * \varphi)$, as it was expected due to Theorem 20.

Thereafter, assume that \mathcal{M} is a consensus merging function. Then, we have that $\mathcal{M}(K_1, K_2, K_3) = K_1 \cap K_2 \cap K_3 = \text{Cn}(\{(a \wedge b \wedge c) \vee (\neg a \wedge b \wedge c) \vee (a \wedge \neg b \wedge c)\})$. Clearly, $\neg\varphi \in \mathcal{M}(K_1, K_2, K_3)$, and therefore, due to (US) and (R), we derive that $\mathcal{M}(K_1, K_2, K_3) * \varphi = \text{Cn}(\{\neg a, b, \neg c\})$. Since \mathcal{M} is a consensus merging function, it also follows that $\mathcal{M}(K_1 * \varphi, K_2 * \varphi, K_3 * \varphi) = (K_1 * \varphi) \cap (K_2 * \varphi) \cap (K_3 * \varphi) = \text{Cn}(\{\neg a, b, \neg c\})$. Observe, again, that $\mathcal{M}(K_1, K_2, K_3) * \varphi = \mathcal{M}(K_1 * \varphi, K_2 * \varphi, K_3 * \varphi)$, as it was expected due to Theorem 21.

Lastly, assume that \mathcal{M} is a singular merging function, and suppose that it is such that $\mathcal{M}(K_1, K_2, K_3) = K_2 = \text{Cn}(\{\neg a, b, c\})$. As $\neg\varphi \in \mathcal{M}(K_1, K_2, K_3)$, we derive from conditions (US) and (R) that $\mathcal{M}(K_1, K_2, K_3) * \varphi = \text{Cn}(\{\neg a, b, \neg c\})$. Since \mathcal{M} is a singular merging function and $K_1 * \varphi = K_2 * \varphi = K_3 * \varphi = \text{Cn}(\{\neg a, b, \neg c\})$, it also follows that $\mathcal{M}(K_1 * \varphi, K_2 * \varphi, K_3 * \varphi) = \text{Cn}(\{\neg a, b, \neg c\})$. Once again, $\mathcal{M}(K_1, K_2, K_3) * \varphi = \mathcal{M}(K_1 * \varphi, K_2 * \varphi, K_3 * \varphi)$, as it was expected due to Theorem 23.

Notwithstanding that the aforementioned satisfiability results show that, under certain (rather strict) conditions, the satisfaction of the PCC is indeed obtainable, the impossibilities of Section 6 are certainly not a good indication for the suitability of that property.

15. \prec denotes the strict part of \preceq . Furthermore, for the sake of readability, possible worlds are represented as sequences (rather than sets) of literals, whereas, the negation of an atom a is represented as \bar{a} , instead of $\neg a$.

Indeed, it can be argued that the PCC is *too restrictive*, as it demands that, if all individual agents revise by an epistemic input, then the collective agent should revise by the same sentence as well. Yet, this strong requirement for collective change may be completely undesirable, since, in a sense, it restricts the independence of agents. Let us illustrate the restrictive requirements of the PCC through the subsequent concrete example.

Example 25 (Restrictive Requirements of the PCC). *Let $\mathcal{P} = \{a, b, c\}$, let K_1, K_2, K_3 be three belief sets such that $K_1 = Cn(\{\neg a, \neg b, \neg c\})$, $K_2 = Cn(\{\neg a, \neg b, c\})$ and $K_3 = Cn(\{\neg a, b, \neg c\})$, and let φ be a sentence of \mathcal{L} such that $\llbracket \varphi \rrbracket = \{\{a, \neg b, \neg c\}, \{a, \neg b, c\}, \{a, b, \neg c\}, \{a, b, c\}\}$. Notice that φ contradicts each one of K_1, K_2 and K_3 . Moreover, let $\langle \mathcal{M}, *_1, \dots, *_6 \rangle$ be a community of 6 rational agents (e.g., an organization's board of directors), where \mathcal{M} is a Hamming-based majority merging function.¹⁶ Lastly, consider the two profiles $E_1 = \{\{K_1, K_1, K_1, K_1, K_1, K_1\}\}$ and $E_2 = \{\{K_1, K_1, K_2, K_2, K_3, K_3\}\}$ of \mathbb{E} .*

On that premises, Table 1 presents the sum of the Hamming distances between each one of the possible worlds $\{\neg a, \neg b, \neg c\}$, $\{\neg a, \neg b, c\}$ and $\{\neg a, b, \neg c\}$ (corresponding to K_1, K_2, K_3 , respectively) and each possible world of \mathbb{M} .¹⁷ As any Hamming-based majority merging function is defined by means of the possible worlds of \mathbb{M} to which the minimum sum of the Hamming distances corresponds, it follows that $\mathcal{M}(E_1) = \mathcal{M}(E_2) = Cn(\{\neg a, \neg b, \neg c\})$ (cf. last line of Table 1).

Now, suppose, for simplicity, that all 6 individual agents utilize a single AGM revision function $$, such that it assigns at K_1, K_2, K_3 (via condition (R)) the following faithful preorders $\preceq_{K_1}, \preceq_{K_2}, \preceq_{K_3}$ over \mathbb{M} , respectively:*

$$\begin{aligned} \llbracket K_1 \rrbracket &= \{\neg a, \neg b, \neg c\} \prec_{K_1} \{a, \neg b, \neg c\} \prec_{K_1} \text{All other possible worlds of } \mathbb{M} \\ \llbracket K_2 \rrbracket &= \{\neg a, \neg b, c\} \prec_{K_2} \{a, b, c\} \prec_{K_2} \text{All other possible worlds of } \mathbb{M} \\ \llbracket K_3 \rrbracket &= \{\neg a, b, \neg c\} \prec_{K_3} \{a, b, c\} \prec_{K_3} \text{All other possible worlds of } \mathbb{M} \end{aligned}$$

*Then, it follows from condition (R) that $\llbracket K_1 * \varphi \rrbracket = \{\{a, \neg b, \neg c\}\}$ and $\llbracket K_2 * \varphi \rrbracket = \llbracket K_3 * \varphi \rrbracket = \{\{a, b, c\}\}$. By calculating the required Hamming distances (in an analogous way as it was done in Table 1), we can derive that*

$$\begin{aligned} \mathcal{M}(K_1 * \varphi, K_1 * \varphi, K_1 * \varphi, K_1 * \varphi, K_1 * \varphi, K_1 * \varphi) &= Cn(\{a, \neg b, \neg c\}) \text{ and} \\ \mathcal{M}(K_1 * \varphi, K_1 * \varphi, K_2 * \varphi, K_2 * \varphi, K_3 * \varphi, K_3 * \varphi) &= Cn(\{\neg a, \neg b, \neg c\}). \end{aligned}$$

Observe, however, that the aforementioned entirely sensible scenario is prohibited by the PCC, as condition (CC) implies that

$$\mathcal{M}(K_1 * \varphi, K_1 * \varphi, K_1 * \varphi, K_1 * \varphi, K_1 * \varphi, K_1 * \varphi) = \mathcal{M}(K_1 * \varphi, K_1 * \varphi, K_2 * \varphi, K_2 * \varphi, K_3 * \varphi, K_3 * \varphi).$$

Obviously yet, the latter equality is an unreasonable requirement, which imposes unnecessary/unnatural restrictions to the process of collective belief revision.

16. A majority merging function \mathcal{M} is an IC quasi-merging function that satisfies the following majority postulate: $\mathcal{M}_\mu(E_2) \subseteq \exists n \mathcal{M}_\mu(E_1 \sqcup E_2^n)$. A concrete example of the operation of Hamming-based majority merging functions is Example 4.4 of Konieczny and Pérez (2002, p. 786).

17. Recall that the Hamming distance between two possible worlds is the number of atoms on which the two possible worlds disagree.

Possible World	K_1 $\{\neg a, \neg b, \neg c\}$	K_2 $\{\neg a, \neg b, c\}$	K_3 $\{\neg a, b, \neg c\}$	Sum of Hamming Distances
$\{a, b, c\}$	3	2	2	7
$\{\neg a, b, c\}$	2	1	1	4
$\{a, \neg b, c\}$	2	1	3	6
$\{a, b, \neg c\}$	2	3	1	6
$\{\neg a, \neg b, c\}$	1	0	2	3
$\{a, \neg b, \neg c\}$	1	2	2	5
$\{\neg a, b, \neg c\}$	1	2	0	3
$\{\neg a, \neg b, \neg c\}$	0	1	1	2

Table 1: The sum of the Hamming distances between each one of the possible worlds $\{\neg a, \neg b, \neg c\}$, $\{\neg a, \neg b, c\}$ and $\{\neg a, b, \neg c\}$ (corresponding to K_1 , K_2 , K_3 , respectively) and each possible world of \mathbb{M} . Observe that the minimum sum of the Hamming distances corresponds to the possible world $\{\neg a, \neg b, \neg c\}$ (last line of the table).

8. Alternative Criteria of Collective Change

As we showed in Section 6, the compliance of the PCC with widely-accepted tools of belief revision and belief merging is facing serious hurdles. In response to this unsatisfactory situation, we propose in this section two *alternative* criteria that characterize the evolution of collective knowledge, and evaluate their compliance with the aforementioned tools of belief change. We, certainly, do not claim that the proposed criteria are the only route forward; our sole aim here is to provide potential properties that the process of collective belief revision should/could respect.

The proposed properties of collective change are encoded into the subsequent conditions (CP1) & (CP2), which are *not* captured by (i.e., do not follow from) the PCC. Both conditions (CP1) & (CP2) are properties that a community $\langle \mathcal{M}, *_1, \dots, *_n \rangle$ of n rational agents can respect, and are expressed in terms of *IC quasi-merging functions* that can adhere to integrity constraints.

$$\text{(CP1)} \quad \mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}(K_1, \dots, K_n) * \varphi.$$

$$\text{(CP2)} \quad \text{If } \neg\varphi \in K_i \text{ for all } i \in \{1, \dots, n\}, \text{ then } \mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}(K_1 *_1 \varphi, \dots, K_n *_n \varphi).$$

Condition (CP1) asserts that the \mathcal{M}_φ -aggregation of K_1, \dots, K_n should coincide with the \mathcal{M} -aggregation of K_1, \dots, K_n revised by the epistemic input φ , through an *AGM revision function* $*$. Condition (CP2) states that, for any sentence φ that contradicts each one

of the belief sets K_1, \dots, K_n , the \mathcal{M}_φ -aggregation of K_1, \dots, K_n should be identical to the \mathcal{M} -aggregation of the revised belief sets $K_1 * \varphi, \dots, K_n * \varphi$. It is easy to verify that, if both conditions (CP1) & (CP2) hold, then, for any sentence φ that contradicts each one of K_1, \dots, K_n , there exists an AGM revision function $*$ such that

$$\mathcal{M}(K_1, \dots, K_n) * \varphi = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi).$$

The previous equality is precisely condition (CC) of Section 5 that encodes the PCC.

Against this background, an *unsatisfiability* result can be formulated proving that there is an IC quasi-merging function \mathcal{M} , such that, for *any* AGM revision function $*$, condition (CP1) is *violated*.

Theorem 26. *There exists an IC quasi-merging function \mathcal{M} , for which there is no AGM revision function $*$ such that, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in \mathcal{M}(K_1, \dots, K_n)$, $\mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}(K_1, \dots, K_n) * \varphi$.*

Proof. Let \mathcal{M} be an IC quasi-merging function such that, for a sentence $\varphi \in \mathcal{L}$ and two profiles $\{\{K_1, \dots, K_n\}\}, \{\{K'_1, \dots, K'_n\}\} \in \mathbb{E}$, $\mathcal{M}(K_1, \dots, K_n) = \mathcal{M}(K'_1, \dots, K'_n)$ and $\mathcal{M}_\varphi(K_1, \dots, K_n) \neq \mathcal{M}_\varphi(K'_1, \dots, K'_n)$. Next, let $*$ be an arbitrary AGM revision function, and suppose towards contradiction that $\mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}(K_1, \dots, K_n) * \varphi$ and $\mathcal{M}_\varphi(K'_1, \dots, K'_n) = \mathcal{M}(K'_1, \dots, K'_n) * \varphi$. Since $\mathcal{M}(K_1, \dots, K_n) = \mathcal{M}(K'_1, \dots, K'_n)$, we derive that $\mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}_\varphi(K'_1, \dots, K'_n)$, which is a contradiction. ■

Despite the impossibility of Theorem 26, it is noteworthy that condition (CP1) is indeed satisfied for *any singleton* profile of \mathbb{E} . This observation follows from the following two remarks: Firstly, from the fact that, for any IC quasi-merging function \mathcal{M} and any belief set K , $\mathcal{M}(K) = K$. Secondly, from a result of Konieczny and Pérez (2002), who showed that IC quasi-merging functions induce AGM revision functions, since, as they prove, for any IC quasi-merging function \mathcal{M} , the revision function $*$, defined as $K * \varphi = \mathcal{M}_\varphi(K)$, is an AGM revision function.¹⁸

Let us now turn to the investigation of condition (CP2). For that case, we present Theorem 27 which is a *satisfiability* result showing that, when the individual agents of a community employ a *specific UR operator* for implementing revision (namely, a UR operator induced from an *indistinguishable* total preorder \preceq over possible worlds), there exists an IC quasi-merging function \mathcal{M} such that condition (CP2) is *satisfied*.¹⁹

Theorem 27. *Let \preceq be an indistinguishable total preorder over \mathbb{M} , and let $*$ be the UR operator generated from \preceq , by means of conditions (US) and (R). Then, there exists an IC quasi-merging function \mathcal{M} such that, for any profile $\{\{K_1, \dots, K_n\}\} \in \mathbb{E}$ and any sentence $\varphi \in \mathcal{L}$ such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$, $\mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$.*

Proof. For proving the theorem, firstly, we shall semantically construct an appropriate IC quasi-merging function \mathcal{M} . To that end, for any profile $E \in \mathbb{E}$, we define a total preorder \sqsubseteq_E over \mathbb{M} , such that:

18. To be precise, this result of Konieczny and Pérez (2002), which is their Theorem 5.3, concerns IC merging functions; however, it can be easily verified by the proof of this theorem that the result also holds for the more general IC-quasi merging functions.

19. Recall that a total preorder \preceq over \mathbb{M} is indistinguishable iff, for any $r, r' \in \mathbb{M}$, $r \approx r'$ (cf. Section 2).

- For all $r, r' \in \llbracket E \rrbracket$, $r \simeq_E r'$.
- If $r \in \llbracket E \rrbracket$ and $r' \notin \llbracket E \rrbracket$, then $r \sqsubseteq_E r'$.
- For all $r, r' \notin \llbracket E \rrbracket$, $r \sqsubseteq_E r'$ iff $r \preceq r'$.

The total preorder \sqsubseteq_E over \mathbb{M} sets all worlds in $\llbracket E \rrbracket$ equally plausible and strictly more plausible (with respect to E) than any world outside $\llbracket E \rrbracket$. Furthermore, \sqsubseteq_E is defined to be identical to the total preorder \preceq , on their sub-domain $(\mathbb{M} - \llbracket E \rrbracket) \times (\mathbb{M} - \llbracket E \rrbracket)$.

Next, we will show that the total preorder \sqsubseteq_E satisfies all conditions (i)–(vi) of Definition 6 (cf. Section 4). The fact that \sqsubseteq_E satisfies conditions (i), (ii) and (iii) is obvious from the definition of \sqsubseteq_E .

For condition (iv), let K_1, K_2 be two belief sets, and let r be any world in $\llbracket K_1 \rrbracket$. If $\llbracket K_1 \rrbracket \cap \llbracket K_2 \rrbracket \neq \emptyset$, then conditions (i) and (ii) entail that there is a world r' in $\llbracket \{K_1\} \sqcup \{K_2\} \rrbracket = \llbracket K_1 \rrbracket \cap \llbracket K_2 \rrbracket$, and thus in $\llbracket K_2 \rrbracket$, such that $r' \sqsubseteq_{\{K_1\} \sqcup \{K_2\}} r$. If $\llbracket K_1 \rrbracket \cap \llbracket K_2 \rrbracket = \emptyset$, then obviously, for any worlds $z, z' \in \mathbb{M}$, $z, z' \notin \llbracket \{K_1\} \sqcup \{K_2\} \rrbracket = \llbracket K_1 \rrbracket \cap \llbracket K_2 \rrbracket$, and thus, $z \sqsubseteq_{\{K_1\} \sqcup \{K_2\}} z'$ iff $z \preceq z'$. Since the total preorder \preceq is indistinguishable, it follows that, for any worlds $z, z' \in \mathbb{M}$, $z \simeq_{\{K_1\} \sqcup \{K_2\}} z'$. Hence, we derive that there is a world $r' \in \llbracket K_2 \rrbracket$, such that $r' \sqsubseteq_{\{K_1\} \sqcup \{K_2\}} r$, as desired.

For condition (v), let E_1, E_2 be two profiles of \mathbb{E} and let r, r' be two worlds of \mathbb{M} , such that $r \sqsubseteq_{E_1} r'$ and $r \sqsubseteq_{E_2} r'$. If $r \in \llbracket E_1 \sqcup E_2 \rrbracket$, then conditions (i) and (ii) entail that $r \sqsubseteq_{E_1 \sqcup E_2} r'$, as desired. Assume, therefore, that $r \notin \llbracket E_1 \sqcup E_2 \rrbracket$. Hence, $r \notin \llbracket E_1 \rrbracket \cap \llbracket E_2 \rrbracket$, which means that either $r \notin \llbracket E_1 \rrbracket$ or $r \notin \llbracket E_2 \rrbracket$, or both. If $r \notin \llbracket E_1 \rrbracket$, we derive from $r \sqsubseteq_{E_1} r'$ that $r' \notin \llbracket E_1 \rrbracket$ as well, and thus $r' \notin \llbracket E_1 \rrbracket \cap \llbracket E_2 \rrbracket$. Therefore, it follows from $r \sqsubseteq_{E_1} r'$ and condition (US) that $r \preceq r'$, and again, from $r \preceq r'$ and the same condition that $r \sqsubseteq_{E_1 \sqcup E_2} r'$. If $r \notin \llbracket E_2 \rrbracket$, we get with a symmetric line of reasoning that $r \sqsubseteq_{E_1 \sqcup E_2} r'$. Hence, in any case, it is true that $r \sqsubseteq_{E_1 \sqcup E_2} r'$, as desired.

For condition (vi), let E_1, E_2 be two profiles of \mathbb{E} and let r, r' be two worlds of \mathbb{M} , such that $r \sqsubseteq_{E_1} r'$ and $r \sqsubseteq_{E_2} r'$. Given that the total preorder \preceq is indistinguishable, $r \sqsubseteq_{E_1} r'$ entails that $r \in \llbracket E_1 \rrbracket$ and $r' \notin \llbracket E_1 \rrbracket$, and similarly, $r \sqsubseteq_{E_2} r'$ entails that $r \in \llbracket E_2 \rrbracket$ and $r' \notin \llbracket E_2 \rrbracket$. Hence, $r \in \llbracket E_1 \rrbracket \cap \llbracket E_2 \rrbracket$ and $r' \notin \llbracket E_1 \rrbracket \cap \llbracket E_2 \rrbracket$. Since $\llbracket E_1 \rrbracket \cap \llbracket E_2 \rrbracket = \llbracket E_1 \sqcup E_2 \rrbracket$, condition (ii) implies then that $r \sqsubseteq_{E_1 \sqcup E_2} r'$, as desired.

Now, let $E = \{\{K_1, \dots, K_n\}\}$ be a profile of \mathbb{E} , and let φ be a sentence of \mathcal{L} such that $\neg\varphi \in K_i$, for all $i \in \{1, \dots, n\}$. Since the total preorder \sqsubseteq_E over \mathbb{M} satisfies all conditions (i)–(vi) of Definition 6, it induces, via condition (M) of Theorem 7, an IC quasi-merging function \mathcal{M} such that $\llbracket \mathcal{M}_\varphi(K_1, \dots, K_n) \rrbracket = \min(\llbracket \varphi \rrbracket, \sqsubseteq_E)$. Given that $\llbracket K_i \rrbracket \cap \llbracket \varphi \rrbracket = \emptyset$, for all $i \in \{1, \dots, n\}$, it follows that $\left(\bigcap_{i=1}^n \llbracket K_i \rrbracket \right) \cap \llbracket \varphi \rrbracket = \emptyset$, which means that $\llbracket E \rrbracket \cap \llbracket \varphi \rrbracket = \emptyset$. Then, the construction of the total preorder \sqsubseteq_E entails that $\llbracket \mathcal{M}_\varphi(K_1, \dots, K_n) \rrbracket = \min(\llbracket \varphi \rrbracket, \preceq)$.

Furthermore, since $*$ is a UR operator, it follows from postulate (UR) that $K_i * \varphi = K_j * \varphi = H$, for all $i, j \in \{1, \dots, n\}$.²⁰ Therefore, we have from postulate (IC2) that $\mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi) = \bigcup_{i=1}^n (K_i * \varphi) = \underbrace{H \cup \dots \cup H}_n = H$. Moreover, from condi-

20. The UR operator $*$ is an AGM revision function that implements a type of revision named by Alchourrón et al. (1985) and Gärdenfors (1988) *full-meet* revision.

tion (US), it is true that $\llbracket H \rrbracket = \min(\llbracket \varphi \rrbracket, \preceq)$. Combining the above, we derive that $\mathcal{M}_\varphi(K_1, \dots, K_n) = \mathcal{M}(K_1 * \varphi, \dots, K_n * \varphi)$, as desired. ■

Although Theorem 27 shows that condition (CP2) can be satisfied for an IC quasi-merging function and an AGM revision function, it can be argued that its assumptions are rather *strong*; not only it demands that the alluded AGM revision function is a UR operator—a fact implying a uniform perception of the domain’s dynamics by the community—, but also it requires this UR operator be induced (via (US) and (R)) by means of a *specific* type of total preorder \preceq over possible worlds—i.e., an indistinguishable total preorder over \mathbb{M} —that sets *all* possible worlds in a *single* equivalence class.

Summing up, in the present section, we took a first step towards the formulation of (plausible) properties to which collective change should/could adhere. Given that the knowledge-dynamics of collective agents are equally vital as the knowledge-dynamics of individual agents, further investigation towards this direction constitutes an imperative task.

9. Conclusion

In this article, we studied the dynamics of collective agents. By formulating David Westlund’s PCC in the realm of genuine belief revision, we firstly established a number of *unsatisfiability results* concerning the combination of AGM revision functions with *principal* types of merging operators (namely, union-when-possible, consensus and radical merging functions). In particular, we showed that there exist communities of rational agents (who behave rationally at the individual level as each agent employs an AGM revision function), such that, for *any* union-when-possible or consensus merging function, violate the PCC (Theorems 16 and 18). This outcome concerns also the IC quasi-merging functions and the formula-based merging functions, as these are union-when-possible merging functions (Corollary 17). The case of radical merging functions is even worse, as the use of *any* radical merging function to aggregate the beliefs of rational agents always leads to the violation of the PCC (Theorem 19). All the aforementioned impossibility results come to *extend* in the realm of belief revision the negative results established by Westlund (2010, Observations 6–10) for the operations of contraction and expansion.

At the opposite of the impossibility results, we established a number of *satisfiability results* too. Specifically, we showed that the satisfaction of the PCC is indeed *feasible* for union-when-possible, consensus and singular merging functions, if the agents of a community revise their beliefs employing operators that implement *uniform belief revision* (Theorems 20, 21 and 23). These satisfiability results—although demand the rather strong requirement of an homogeneous perception of the dynamics of the domain—are of interest in their own right, but also offer a proof-of-concept for uniform belief revision.

Notwithstanding that the aforementioned satisfiability results show that, under certain (rather strict) requirements, the satisfaction of the PCC is indeed obtainable, the derived impossibilities are certainly not a good indication for the suitability of that property. Indeed, as we demonstrated in the representative Example 25, the PCC is *too restrictive*, as it prohibits perfectly sensible scenarios of collective belief revision.

Mainly in response to the unsatisfactory situation regarding the suitability of the PCC, two *alternative* criteria that characterize the revision of collective knowledge, based on Konieczny and Pino Pérez’s IC quasi-merging functions, were sought. The evaluation of the compliance of these alternative criteria of collective change with belief revision and belief merging showed controversial indications, as both satisfiability and unsatisfiability results were obtained.

Future work shall be devoted to the investigation of (plausible) properties to which the evolution of collective knowledge should/could adhere (as earlier stated), as well as to the study of the interrelation between the results established herein and Westlund’s results, by means of the well-known Levi and Harper Identities.

Acknowledgments

The author expresses his gratitude to Pavlos Peppas for the fruitful discussions on the topic of this study, as well as to the anonymous reviewers for their constructive input on previous versions of this article.

References

- Alchourrón, C., Gärdenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, *50*(2), 510–530.
- Aravanis, T. (2020). On uniform belief revision. *Journal of Logic and Computation*, *30*, 1357–1376.
- Aravanis, T., & Peppas, P. (2022). Theory-relational belief revision. *Annals of Mathematics and Artificial Intelligence*, *90*, 573–594.
- Areces, C., & Becher, V. (2001). Iterable AGM functions. In Williams, M.-A., & Rott, H. (Eds.), *Frontiers in Belief Revision*, Vol. 22 of *Applied Logic Series*, pp. 165–196. Springer, Dordrecht.
- Baral, C., Kraus, S., Minker, J., & Subrahmanian, V. S. (1992). Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, *8*, 45–71.
- Benferhat, S., Dubois, D., Lang, J., Prade, H., Saffiotti, A., & Smets, P. (1998). A general approach for inconsistency handling and merging information in prioritized knowledge bases. In *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR 1998)*, pp. 466–477.
- Everaere, P., Konieczny, S., & Marquis, P. (2010). Disjunctive merging: Quota and Gmin merging operators. *Artificial Intelligence*, *174*, 824–849.
- Fermé, E., & Hansson, S. O. (2018). *Belief Change: Introduction and Overview*. Springer International Publishing.
- Gärdenfors, P. (1988). *Knowledge in Flux – Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Massachusetts.
- Gilbert, M. (1987). Modelling collective belief. *Synthese*, *73*, 185–204.

- Gilbert, M. (2004). Collective epistemology. *Episteme*, 1, 95–107.
- Katsuno, H., & Mendelzon, A. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3), 263–294.
- Kfir-Dahav, N. E., & Tennenholtz, M. (1996). Multi-agent belief revision. In *Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 1996)*, pp. 175–194.
- Konieczny, S. (2000). On the difference between merging knowledge bases and combining them. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR 2000)*, pp. 135–144.
- Konieczny, S., & Pérez, R. P. (2002). Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5), 773–808.
- Konieczny, S., & Pérez, R. P. (2011). Logic based merging. *Journal of Philosophical Logic*, 40, 239–270.
- Labbé, I., van Dokkum, P., Nelson, E., Bezanson, R., Suess, K. A., Leja, J., Brammer, G., Whitaker, K., Mathews, E., Stefanon, M., & Wang, B. (2023). A population of red candidate massive galaxies ~ 600 Myr after the Big Bang. *Nature*, 616, 266–269.
- Liu, F., Seligman, J., & Girard, P. (2014). Logical dynamics of belief change in the community. *Synthese*, 191, 2403–2431.
- Liu, W., & Williams, M.-A. (2001). A framework for multi-agent belief revision. *Studia Logica*, 67, 291–312.
- Malheiro, B., Jennings, N. R., & Oliveira, E. (1994). Belief revision in multi-agent systems. In *Proceedings of the 11th European Conference on Artificial Intelligence (ECAI 1994)*, pp. 294–298.
- Maynard-Zhang, P., & Lehmann, D. (2003). Representing and aggregating conflicting beliefs. *Journal of Artificial Intelligence Research*, 19, 155–203.
- Minsky, M. (1988). *The Society of Mind*. Simon & Schuster.
- Paranamana, P., Wang, P., & Shafto, P. (2022). Evolution of beliefs in social networks. *Collective Intelligence*, 1, 1–14.
- Peppas, P. (2008). Belief revision. In van Harmelen, F., Lifschitz, V., & Porter, B. (Eds.), *Handbook of Knowledge Representation*, pp. 317–359. Elsevier Science.
- Peppas, P. (2014). A panorama of iterated revision. In Hansson, S. O. (Ed.), *David Makinson on Classical Methods for Non-Classical Problems*, pp. 71–94. Springer Netherlands.
- Tojo, S. (2013). Collective belief revision in linear algebra. In *2013 Federated Conference on Computer Science and Information Systems*, pp. 175–178.
- Westlund, D. (2010). Rational belief changes for collective agents. In Olsson, E. J., & Enqvist, S. (Eds.), *Belief Revision meets Philosophy of Science*, Vol. 21 of *Logic, Epistemology, and the Unity of Science*, pp. 213–224. Springer, Dordrecht.
- Wray, K. B. (2001). Collective belief and acceptance. *Synthese*, 129, 313–333.