

On Consensus in Belief Merging

Nicolas Schwind

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

Pierre Marquis

CRIL-CNRS, Université d'Artois,
Institut Universitaire de France
Lens, France
marquis@cril.fr

Abstract

We define a consensus postulate in the propositional belief merging setting. In a nutshell, this postulate imposes the merged base to be consistent with the pieces of information provided by each agent involved in the merging process. The interplay of this new postulate with the IC postulates for belief merging is studied, and an incompatibility result is proved. The maximal sets of IC postulates which are consistent with the consensus postulate are exhibited. When satisfying some of the remaining IC postulates, consensus operators are shown to suffer from a weak inferential power. We then introduce two families of consensus operators having a better inferential power by setting aside some of these postulates.

Introduction

In this paper, we are interested in defining consensus operators for aggregating pieces of propositional information. We consider a set of n communicating agents, each of them wanting to refine her own propositional beliefs φ_i by merging them with the beliefs of the other agents of the group. Reaching this goal requires first to merge all the belief bases, then to determine how to refine each φ_i with the result of the merging process. To achieve the first step, many belief merging (BM) operators can be exploited, see e.g., (Lin 1996; Revesz 1997; Liberatore and Schaerf 1998; Konieczny and Pino Pérez 2002a; Konieczny and Pino Pérez 2011). The second step calls for revision policies, as considered in (Schwind et al. 2015; 2016).

Here the focus is laid on agents which are *reluctant to change*: each agent is ready to accept the merged base, provided that it allows her to refine her prior beliefs φ_i but does not question them. Thus, the revision policy which is adopted by each such agent consists in expanding her belief base φ_i by the merged base if the conjunction is consistent, and to keep φ_i unchanged otherwise. In order to avoid the latter case, we are interested in defining *consensus operators*, i.e., merging operators such that the merged base C that is generated satisfies the *consensus condition*: C is consistent with every input base φ_i that is consistent with the integrity constraint μ . In such a case, C is said to be a consensus for \mathcal{K} under μ .

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Since consensus operators are BM operators, it is important to situate them within the family of BM operators. This family received much attention in AI and some rationality postulates (the so-called IC postulates) associated with representation theorems have been pointed out (Konieczny and Pino Pérez 2002a). As BM operators, consensus operators are thus expected to satisfy as many IC postulates as possible. Unfortunately, using fully rational merging operators does not lead to compute merged bases which are consensus in the general case.

Example 1 Consider a set of three US citizens, Alice, Bob, Charlie traveling together to Paris, and planning a trip to Normandy. They all know that they have to take a train at station Gare St Lazare but have different beliefs about the location and availability of the station from CDG Airport. Thus, Alice (agent 1) believes that Gare St Lazare is located at the north of Paris midtown, and at the west. Bob (agent 2) has been told that Gare St Lazare is located at the south of Paris midtown, and that it is not reachable from CDG airport in less than 30mn, and Charlie (agent 3) believes that it is located at the north of Paris midtown, and is reachable from CDG airport in less than 30mn. Representing the information using three propositional symbols (a : Gare St Lazare is located at the south of Paris midtown, b : Gare St Lazare is located at the east of Paris midtown, and c : Gare St Lazare is reachable from CDG airport in less than 30mn), the belief bases of Alice, Bob, and Charlie are respectively $\varphi_1 = \{\neg a \wedge \neg b\}$, $\varphi_2 = \{a \wedge \neg c\}$, and $\varphi_3 = \{\neg a \wedge c\}$. Here there are no integrity constraint ($\mu = \top$). Suppose one takes advantage of the distance-based merging operator $\Delta^{d_H, \Sigma}$ to merge the belief bases of the profile $\mathcal{K} = \{\varphi_1, \varphi_2, \varphi_3\}$ under μ . $\Delta^{d_H, \Sigma}$ is based on the Hamming distance and uses sum as the aggregation function. It is known to satisfy all the IC postulates (Konieczny and Pino Pérez 2002a), thus it appears *prima facie* as a good candidate for this job. However, $\Delta_{\mu}^{d_H, \Sigma}(\{\varphi_1, \varphi_2, \varphi_3\})$ is equivalent to the base $\{\neg a \wedge \neg b\}$,¹ which is not a consensus for $\{\varphi_1, \varphi_2, \varphi_3\}$ under μ since it conflicts with Bob's beliefs.

In order to satisfy the consensus condition for this example, a logically weaker merged base must be computed. Thus, merging operators having quite a low inferential power look at a first glance as suitable consensus operators.

¹Details of the computations are reported in Table 1.

Among them is the basic merging operator Δ^b (Konieczny and Pérez 1999): $\Delta_\mu^b(\{\varphi_1, \varphi_2, \varphi_3\})$ is equivalent to the disjunction of the three bases $\varphi_1, \varphi_2, \varphi_3$, hence to $\{(\neg a \vee \neg c) \wedge (a \vee \neg b \vee c)\}$, so that it is consistent with every of them. However, this consensus is not convenient since it does not convey any new piece of information to any of the three agents (expanding φ_1, φ_2 , or φ_3 with it does not change anything: there is no belief refinement at all). One consequence is that, for instance, it does not entail $\neg b$ which is not questioned by any of the three agents. Contrastingly, merged bases equivalent to $\{(\neg a \vee \neg c) \wedge \neg b\}$ or to $\{(\neg a \vee \neg c) \wedge (a \vee c) \wedge \neg b\}$ would be much better consensuses, since the first one would lead Bob and Charlie to improve their beliefs and the second one would lead each of the three agents to improve her beliefs. This example shows that one must take care of the inferential power of the consensus operators which are used.

The very objective of this paper is to determine the extent to which the consensus condition is compatible with the IC postulates, to point out some consensus operators, and to delineate borders of the trade-off between the rationality conditions and the inferential power offered by such operators. The main contributions are as follows. First of all, a consensus postulate is formally defined. The interplay of this new postulate with the “standard” IC postulates for belief merging is studied. We show that the consensus postulate is incompatible with the conjunction of the IC postulates **(IC2)** and **(IC6)**, while compatible with each of them taken separately. Since **(IC2)** is more central to BM than **(IC6)**, the focus is laid on consensus operators satisfying **(IC2)**. We also investigate how our consensus postulate interacts with the majority, arbitration and disjunction postulates. Then we show that the consensus operators satisfying **(IC0)** and **(IC8)** have a weak inferential power. Thus setting aside **(IC8)** is necessary for getting consensus operators satisfying **(IC0)** and **(IC2)**, and which are not too weak from the inferential standpoint. On this ground, we introduce two families of consensus operators. The first one is composed of the distance-based operators relying on a Pareto aggregation. Such operators satisfy all IC postulates but **(IC6)** and **(IC8)**. The second one gathers consensus operators ∇ induced by merging operators Δ , and can be viewed as a generalization of the family of arbitration operators from (Liberatore and Schaerf 1998). When the underlying BM operator Δ is an IC one, the corresponding consensus operator ∇ is ensured to satisfy all IC postulates but **(IC5)**, **(IC6)** and **(IC8)**. We show how the expected consensuses for Example 1 as given above can be computed using some of the introduced consensus operators, and provide comparative results on the behaviour of these operators from the inferential standpoint. For space reasons, most proofs are omitted. An extended version of the paper containing all the proofs is available from www.cril.fr/~marquis/aaai18-extended.pdf.

A Glimpse at Propositional Merging

Let $\mathcal{L}_{\mathcal{P}}$ be a propositional language built up from a finite set of propositional variables \mathcal{P} and the usual connectives. \perp (resp. \top) is the Boolean constant always false (resp. true).

An interpretation is a mapping from \mathcal{P} to $\{0, 1\}$, denoted by a bit vector whenever a strict total order on \mathcal{P} is specified. The set of all interpretations is denoted \mathcal{W} . $[\varphi]$ denotes the set of models of the formula φ , i.e., $[\varphi] = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. \models denotes logical entailment and \equiv logical equivalence, i.e., $\varphi \models \psi$ iff $[\varphi] \subseteq [\psi]$ and $\varphi \equiv \psi$ iff $[\varphi] = [\psi]$.

An integrity constraint μ is a propositional formula. A belief base φ denotes the set of beliefs of an agent, it is a finite and consistent set of propositional formulae, interpreted conjunctively, so that φ is identified with the conjunction of its elements. A profile $\mathcal{K} = \{\varphi_1, \dots, \varphi_n\}$ is a finite, multi-set of belief bases. $\bigwedge \mathcal{K}$ (resp. $\bigvee \mathcal{K}$) denotes the conjunction (resp. the disjunction) of all bases from \mathcal{K} . \sqcup refers to the union of multi-sets. Given a profile \mathcal{K} and a formula μ , $MC(\mathcal{K}, \mu)$ denotes the set of maximal subsets of bases from \mathcal{K} which are jointly consistent with μ , that is, $MC(\mathcal{K}, \mu) = \{\mathcal{K}' \subseteq \mathcal{K} \mid \bigwedge \mathcal{K}' \wedge \mu \not\models \perp, \forall \mathcal{K}'' \subseteq \mathcal{K}, \mathcal{K}' \subsetneq \mathcal{K}'' \Rightarrow \bigwedge \mathcal{K}'' \wedge \mu \models \perp\}$. Lastly, $\bigcap MC(\mathcal{K}, \mu)$ denotes the set $\bigcap_{\mathcal{K}' \in MC(\mathcal{K}, \mu)} \mathcal{K}'$.

A BM operator Δ is a mapping associating with every integrity constraint μ and every profile $\mathcal{K} = \{\varphi_1, \dots, \varphi_n\}$ with $n \geq 1$, a belief base $\Delta_\mu(\mathcal{K})$ called the merged base. A set of “standard” properties denoted **(IC0)**-**(IC8)** (called IC postulates) expected for BM operators have been pointed out (Konieczny and Pino Pérez 2002a). Operators satisfying them are called *IC merging operators*.

(IC0) $\Delta_\mu(\mathcal{K}) \models \mu$;

(IC1) If μ is consistent, then $\Delta_\mu(\mathcal{K})$ is consistent;

(IC2) If $\bigwedge \mathcal{K} \wedge \mu$ is consistent, then $\Delta_\mu(\mathcal{K}) \equiv \bigwedge \mathcal{K} \wedge \mu$;

(IC3) If $\mathcal{K}_1 \equiv \mathcal{K}_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(\mathcal{K}_1) \equiv \Delta_{\mu_2}(\mathcal{K}_2)$;

(IC4) If $\varphi_1 \models \mu$, $\varphi_2 \models \mu$ and $\Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_1$ is consistent, then $\Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_2$ is consistent;

(IC5) $\Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2) \models \Delta_\mu(\mathcal{K}_1 \sqcup \mathcal{K}_2)$;

(IC6) If $\Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2)$ is consistent, then $\Delta_\mu(\mathcal{K}_1 \sqcup \mathcal{K}_2) \models \Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2)$;

(IC7) $\Delta_{\mu_1}(\mathcal{K}) \wedge \mu_2 \models \Delta_{\mu_1 \wedge \mu_2}(\mathcal{K})$;

(IC8) If $\Delta_{\mu_1}(\mathcal{K}) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(\mathcal{K}) \models \Delta_{\mu_1}(\mathcal{K}) \wedge \mu_2$.

In these postulates, when $\mathcal{K}_1 = \{\varphi_{1,1}, \dots, \varphi_{1,n}\}$ and $\mathcal{K}_2 = \{\varphi_{2,1}, \dots, \varphi_{2,n}\}$, $\mathcal{K}_1 \equiv \mathcal{K}_2$ means that there exists a bijection π from $\{1, \dots, n\}$ to $\{1, \dots, n\}$ such that for each $i \in \{1, \dots, n\}$, we have $\varphi_{1,i} \equiv \varphi_{2,\pi(i)}$. We refer the reader to (Konieczny and Pino Pérez 2002a) for a detailed explanation about the rationale of these postulates.

The following additional postulates will also be considered in this paper (Konieczny and Pino Pérez 2002a; Everaere, Konieczny, and Marquis 2010): **(Maj)**, **(Arb)** and **(Disj)**, that respectively characterize the class of majority, arbitration (Konieczny and Pino Pérez 2002a) and disjunctive (Everaere, Konieczny, and Marquis 2010) operators:

(Maj) $\exists n \geq 1 \Delta_\mu(\mathcal{K}_1 \sqcup \underbrace{\mathcal{K}_2 \sqcup \dots \sqcup \mathcal{K}_2}_n) \models \Delta_\mu(\mathcal{K}_2)$.

$$(\text{Arb}) \text{ If } \begin{cases} \Delta_{\mu_1}(\{\varphi_1\}) \equiv \Delta_{\mu_2}(\{\varphi_2\}) \\ \Delta_{\mu_1 \leftrightarrow \neg \mu_2}(\{\varphi_1, \varphi_2\}) \equiv (\mu_1 \leftrightarrow \neg \mu_2) \\ \mu_1 \not\equiv \mu_2 \\ \mu_2 \not\equiv \mu_1 \end{cases}$$

then $\Delta_{\mu_1 \vee \mu_2}(\{\varphi_1, \varphi_2\}) \equiv \Delta_{\mu_1}(\{\varphi_1\})$.

(Disj) If $\bigvee \mathcal{K} \wedge \mu$ is consistent, then $\Delta_\mu(\mathcal{K}) \models \bigvee K$.

Among IC merging operators are some *distance-based* operators, i.e., operators which are based on the selection of some models of the integrity constraint, the “closest” ones to the given profile. Those operators are characterized by a distance d between interpretations and an aggregation function f (Konieczny, Lang, and Marquis 2004). They associate with every integrity constraint μ and every profile \mathcal{K} a belief base $\Delta_\mu^{d,f}(\mathcal{K})$ which satisfies $[\Delta_\mu^{d,f}(\mathcal{K})] = \min([\mu], \leq_{\mathcal{K}}^{d,f})$, where $\leq_{\mathcal{K}}^{d,f}$ is the total preorder over interpretations induced by \mathcal{K} defined by $\omega \leq_{\mathcal{K}}^{d,f} \omega'$ if and only if $d^f(\omega, \mathcal{K}) \leq d^f(\omega', \mathcal{K})$, where $d^f(\omega, \mathcal{K}) = f_{\varphi_i \in \mathcal{K}}\{d(\omega, \varphi_i)\}$ and $d(\omega, \varphi_i) = \min_{\omega' \models \varphi_i} d(\omega, \omega')$. Usual distances are d_D , the drastic distance and d_H , the Hamming distance. Note that some distance-based operators are not IC merging ones (some conditions must be satisfied by f , see (Konieczny, Lang, and Marquis 2004)) but taking advantage of usual aggregation functions as Σ , GMax and GMin (Everaere, Konieczny, and Marquis 2010) lead to IC merging operators.

Example 1 (continued) Let $\mathcal{L}_{\mathcal{P}}$ be built up from the set of propositional variables $\mathcal{P} = \{a, b, c\}$, $\varphi_1 = \{\neg a \wedge \neg b\}$, $\varphi_2 = \{a \wedge \neg c\}$, $\varphi_3 = \{\neg a \wedge c\}$ and $\mu = \top$. We consider Σ , GMax and GMin operators based on the drastic distance on the one hand, the Hamming distance on the other hand. Table 1 gives for each interpretation $\omega \in [\mu]$ the distances $d(\omega, \varphi_i)$ for $d \in \{d_D, d_H\}$ and $i \in \{1, 2, 3\}$, and shows for each interpretation whether it is selected in $\Delta_\mu^{d,\Sigma}(\mathcal{K})$, $\Delta_\mu^{d,\text{GMax}}(\mathcal{K})$ and $\Delta_\mu^{d,\text{GMin}}(\mathcal{K})$ respectively (interpretations ω are denoted as binary sequences following the ordering $a < b < c$).

- $\Delta_\mu^{d_D, \Sigma}(\mathcal{K}) \equiv \Delta_\mu^{d_D, \text{GMax}}(\mathcal{K}) \equiv \Delta_\mu^{d_D, \text{GMin}}(\mathcal{K}) \equiv \neg a \wedge \neg b \wedge c \equiv \varphi_1 \wedge \varphi_3$.
- $\Delta_\mu^{d_H, \Sigma}(\mathcal{K}) \equiv \neg a \wedge \neg b \equiv \varphi_1$;
- $\Delta_\mu^{d_H, \text{GMax}}(\mathcal{K}) \equiv \neg a \wedge \neg b \wedge \neg c$;
- $\Delta_\mu^{d_H, \text{GMin}}(\mathcal{K}) \equiv \neg a \wedge \neg b \wedge c \equiv \varphi_1 \wedge \varphi_3$.

Consensus and Propositional Merging

A *consensus operator* is a BM operator satisfying the following consensus postulate **(CO)**:

(CO) $\forall \varphi_i \in \mathcal{K}$, if $\varphi_i \wedge \mu$ is consistent, then $\varphi_i \wedge \Delta_\mu(\mathcal{K})$ is consistent.

When Δ is a merging operator satisfying **(CO)**, the merged base $\Delta_\mu(\mathcal{K})$ is said to be a consensus for \mathcal{K} under μ .

Let us investigate how the consensus postulate interacts with the IC postulates. First of all, **(CO)** can be viewed as a stronger version of the equity postulate **(IC4)**:

Proposition 1 Every consensus operator satisfies **(IC4)**.

Then a key issue is to determine whether **(CO)** is compatible or not with the IC postulates. It turns out that the answer to it is negative. More precisely:

Proposition 2 There is no BM operator jointly satisfying **(IC2)**, **(IC6)**, and **(CO)**.

Proof: Let $\mathcal{K} = \{\varphi_1, \varphi_2\}$ such that $[\varphi_1] = \{\omega_1\}$ and $[\varphi_2] = \{\omega_2\}$, and let μ such that $[\mu] = \{\omega_1, \omega_2\}$. Towards a contradiction, assume there exists a BM operator Δ satisfying **(IC2)**, **(IC6)**, and **(CO)**. By **(IC2)**, $\Delta_\mu(\{\varphi_2\}) \equiv \varphi_2$, so $[\Delta_\mu(\{\varphi_2\})] = \{\omega_2\}$. Furthermore, by **(CO)**, $\{\omega_1, \omega_2\} \subseteq [\Delta_\mu(\{\varphi_1, \varphi_2\})]$. Hence, $[\Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \Delta_\mu(\{\varphi_2\})] = \{\omega_2\}$, so $\Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \Delta_\mu(\{\varphi_2\}) \not\models \perp$. Then, by **(IC6)**, $\Delta_\mu(\{\varphi_1, \varphi_2, \varphi_2\}) \models \Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \Delta_\mu(\{\varphi_2\})$. Thus $[\Delta_\mu(\{\varphi_1, \varphi_2, \varphi_2\})] = \{\omega_2\}$. However, by **(CO)** we must also have $\{\omega_1, \omega_2\} \subseteq [\Delta_\mu(\{\varphi_1, \varphi_2, \varphi_2\})]$. Contradiction. ■

The other IC postulates do not jointly conflict with **(CO)**. Indeed, one can find a BM operator Δ satisfying **(CO)** and all the IC postulates but **(IC2)**. The *trivial merging operator* Δ^t given by $\Delta_\mu^t(\mathcal{K}) = \{\mu\}$ does the job.

Proposition 3 Δ^t satisfies **(CO)** and all IC postulates but **(IC2)**.

But Δ^t is not an “interesting” operator as it gives up all information from any input profile, even when all belief bases are jointly consistent with the integrity constraint.

On the other hand, one can define consensus operators satisfying **(IC2)** and “almost” all other IC postulates. More precisely, there exist BM operators Δ satisfying **(CO)** and all the IC postulates but **(IC6)**. This is the case of the *drastic merging operator* Δ^d (Konieczny and Pérez 1999) defined by:

$$\Delta_\mu^d(\mathcal{K}) = \begin{cases} \bigwedge \mathcal{K} \wedge \mu & \text{if consistent,} \\ \mu & \text{otherwise;} \end{cases}$$

and of the *basic merging operator* Δ^b (Konieczny and Pérez 1999) defined by:

$$\Delta_\mu^b(\mathcal{K}) = \begin{cases} \bigwedge \mathcal{K} \wedge \mu & \text{if consistent,} \\ \bigvee \mathcal{K} \wedge \mu & \text{if } \bigwedge \mathcal{K} \wedge \mu \text{ is inconsistent} \\ & \text{and } \bigvee \mathcal{K} \wedge \mu \text{ is consistent,} \\ \mu & \text{otherwise.} \end{cases}$$

Proposition 4 Δ^d and Δ^b satisfy **(CO)** and all IC postulates but **(IC6)**.

The last two propositions can be used to strengthen Proposition 2, showing that **(IC2)**, **(IC6)**, and **(CO)** is a *minimal* set of incompatible postulates (i.e., while there is no merging operator satisfying the three postulates, one can find merging operators satisfying any proper subset of it). Yet **(IC2)** is essential to belief merging. Indeed, a fundamental expectation of BM is to ensure that the beliefs of a group of agents is at least as (logically) strong as the individual beliefs, when there is no conflict between them. That way, synergetic effects are possible, i.e., some logical consequences of the beliefs of the group may not be among the consequences of any isolated agent. This is what **(IC2)** captures. On the other hand, **(IC6)** is sometimes considered as too “strong.”

ω	φ_1	φ_2	φ_3	Σ	GMax	GMin	Par	$\nabla^{d_D, f}$		φ_1	φ_2	φ_3	Σ	GMax	GMin	Par	$\nabla^{d_H, \text{GMin}}$	$\nabla^{d_H, g}$
000	0	1	1						$d = d_D$	0	1	1	•	•		•		•
100	1	0	1				•	•		1	0	2				•		•
010	1	1	1							1	1	1						
001	0	1	0	•	•	•	•	•		0	2	0	•		•	•	•	•
110	1	0	1				•	•		2	0	2						
101	1	1	1							1	1	1						
011	1	1	0							1	2	0						
111	1	1	1							2	1	1						
									$d = d_H$									

Table 1: The BM operators $\Delta^{d, \Sigma}$, $\Delta^{d, \text{GMax}}$, $\Delta^{d, \text{GMin}}$, $\Delta^{d, \text{Par}}$ and $\nabla^{d, f}$ for $d \in \{d_D, d_H\}$, $f \in \{\Sigma, \text{GMax}, \text{GMin}\}$ and $g \in \{\Sigma, \text{GMax}\}$. The left part (resp. right part) depicts the case where $d = d_D$ (resp. $d = d_H$).

Loosely speaking it is the counterpart of a very restrictive Pareto condition, considering the aggregation of sets of criteria, not only individual ones. This explains why many merging operators of interest proposed in the literature fail to satisfy **(IC6)**, e.g., the formula-based operators (Konieczny 2000) and the quota operators (Everaere, Konieczny, and Marquis 2010); and this is also why weaker postulates (like **(IC6b)** introduced in (Everaere, Konieczny, and Marquis 2014)²) have been previously considered. Therefore, as we are interested in consensus operators satisfying **(IC2)**, **(IC6)** must be set aside.

Let us now take look at how consensus operators relate to **(Maj)**, **(Arb)** and **(Disj)** together with the IC postulates (except **(IC6)**). It turns out that similarly to **(IC6)**, **(Maj)** appears as antagonistic with **(CO)** under **(IC2)** (Proposition 5), and that **(IC2)**, **(Maj)**, and **(CO)** is a minimal set of incompatible postulates (Proposition 6):

Proposition 5 *There is no BM operator jointly satisfying **(IC2)**, **(Maj)**, and **(CO)**.*

Proposition 6 Δ^t satisfies **(Maj)**.

On the other hand, **(Arb)** and **(Disj)** are compatible with **(CO)** and all IC postulates (except **(IC6)**):

Proposition 7 Δ^d and Δ^b satisfy **(Arb)** and **(Disj)**.

According to Proposition 4, the drastic and basic operators appear as relatively “well-behaved” in terms of IC properties. In addition (Proposition 7), they satisfy **(Arb)** and **(Disj)**. However, whenever $\bigwedge \mathcal{K} \wedge \mu$ is inconsistent nothing new can be inferred from the merged base $\Delta_\mu^d(\mathcal{K})$ or $\Delta_\mu^b(\mathcal{K})$ that cannot be inferred by at least one the bases from the profile. That is to say, for every base $\varphi_i \in \mathcal{K}$, $\varphi_i \models \Delta_\mu^d(\mathcal{K})$ and $\varphi_i \models \Delta_\mu^b(\mathcal{K})$. A natural question is then whether there exist some consensus operators which satisfy all (or a subset of) IC postulates apart from **(IC6)** and which offer a reasonable compromise from the inferential standpoint. It turns out that the consensus operators satisfying a few IC postulates do not preserve much information from the input profile:

Proposition 8 *Let Δ be a BM operator satisfying **(CO)**, **(IC0)** and **(IC8)**. For any profile \mathcal{K} and each base $\varphi_i \in \mathcal{K}$, if $\varphi_i \wedge \mu \not\models \Delta_\mu(\mathcal{K})$ then (i) $\varphi_i \in \bigcap MC(\mathcal{K}, \mu)$ and (ii) $\forall \varphi_j \in \mathcal{K}$, $\varphi_j \notin \bigcap MC(\mathcal{K}, \mu) \Rightarrow \varphi_j \wedge \mu \models \varphi_i \wedge \mu$.*

²**(IC6b)** requires that if $\Delta_\mu(\varphi_1) \wedge \dots \wedge \Delta_\mu(\varphi_n)$ is consistent, then $\Delta_\mu(\varphi_1, \dots, \varphi_n) \models \Delta_\mu(\varphi_1) \wedge \dots \wedge \Delta_\mu(\varphi_n)$.

Proof: We first prove the following lemma:

Lemma 1 *Let Δ be an operator satisfying **(CO)**, **(IC0)** and **(IC8)**. For each base $\varphi_i \in \mathcal{K}$, if $\varphi_i \wedge \mu$ is consistent then $\varphi_i \wedge \mu \models \Delta_\mu(\mathcal{K})$ or $\Delta_\mu(\mathcal{K}) \models \varphi_i \wedge \mu$.*

Proof: Given a set of interpretations M , let $\phi(M)$ denote any formula satisfying $[\phi(M)] = M$. Assume $\varphi_i \wedge \mu \not\models \Delta_\mu(\mathcal{K})$ and let us prove that $\Delta_\mu(\mathcal{K}) \models \varphi_i \wedge \mu$. Let $\omega \models \varphi_i \wedge \mu \wedge \neg \Delta_\mu(\mathcal{K})$. Assume towards a contradiction that $\Delta_\mu(\mathcal{K}) \not\models \varphi_i \wedge \mu$. Then let $\omega' \models \Delta_\mu(\mathcal{K}) \wedge \neg(\varphi_i \wedge \mu)$. By **(IC0)**, $\omega' \models \mu$ so $\omega' \models \Delta_\mu(\mathcal{K}) \wedge \neg \varphi_i$. Since $\omega \not\models \Delta_\mu(\mathcal{K})$ and $\omega' \models \Delta_\mu(\mathcal{K})$, $\Delta_\mu(\mathcal{K}) \wedge \phi(\{\omega, \omega'\}) \equiv \phi(\{\omega'\})$. So by **(IC8)**, $\Delta_{\phi(\{\omega, \omega'\})}(\mathcal{K}) \models \phi(\{\omega'\})$. On the one hand, $\omega' \not\models \varphi_i$ so $\Delta_{\phi(\{\omega, \omega'\})}(\mathcal{K}) \wedge \varphi_i \models \perp$. On the other hand, $\omega \models \varphi_i$ so $\phi(\{\omega, \omega'\}) \wedge \varphi_i \not\models \perp$, and by **(CO)** $\Delta_{\phi(\{\omega, \omega'\})}(\mathcal{K}) \wedge \varphi_i \not\models \perp$. Contradiction. ■

We now prove the proposition. Let $\varphi_i \in \mathcal{K}$, $\varphi_i \wedge \mu \not\models \Delta_\mu(\mathcal{K})$. We prove (i) $\varphi_i \in \bigcap MC(\mathcal{K}, \mu)$. Assume towards a contradiction that $\varphi_i \notin \bigcap MC(\mathcal{K}, \mu)$. Since $\varphi_i \wedge \mu \not\models \perp$, $MC(\mathcal{K}, \mu) \neq \emptyset$. So let $S \subseteq MC(\mathcal{K}, \mu)$, $\varphi_i \notin S$. By Lemma 1, $\Delta_\mu(\mathcal{K}) \models \varphi_i \wedge \mu$. Yet $\varphi_i \notin S$, so $\bigwedge S \wedge \mu \wedge \varphi_i \models \perp$. So let $\varphi_j \in S$, $\Delta_\mu(\mathcal{K}) \not\models \varphi_j \wedge \mu$. We have $\varphi_j \wedge \mu \not\models \perp$, so by Lemma 1, $\varphi_j \wedge \mu \models \Delta_\mu(\mathcal{K})$. We got that $\Delta_\mu(\mathcal{K}) \models \varphi_i \wedge \mu$ and $\varphi_j \wedge \mu \models \Delta_\mu(\mathcal{K})$, therefore $\varphi_j \wedge \mu \models \varphi_i \wedge \mu$. This contradicts that $\bigwedge S \wedge \mu \wedge \varphi_i \models \perp$. Therefore, $\varphi_i \in \bigcap MC(\mathcal{K}, \mu)$. We now prove (ii). We already got that $\Delta_\mu(\mathcal{K}) \models \varphi_i \wedge \mu$. Let $\varphi_j \notin \bigcap MC(\mathcal{K}, \mu)$. By using the contrapositive of (i) on φ_j , we get $\varphi_j \wedge \mu \models \Delta_\mu(\mathcal{K})$. Hence, $\varphi_j \wedge \mu \models \varphi_i \wedge \mu$. This concludes the proof. ■

Proposition 8 (i) tells us that using a consensus operator Δ satisfying **(IC0)** and **(IC8)**, leads the merged base $\Delta_\mu(\mathcal{K})$ to be entailed by each base from the profile (in conjunction with μ) which does not belong to *all* maximal consistent subsets from $MC(\mathcal{K}, \mu)$; and (ii) each one these bases (in conjunction with μ) necessarily entails each base that belongs to all maximal consistent subsets from $MC(\mathcal{K}, \mu)$. This shows that consensus operators satisfying **(IC0)** and **(IC8)** have a weak inferential power.

The next proposition is a noticeable consequence of these results. We say that a profile \mathcal{K} contains a disagreement over μ whenever $\bigwedge \{\varphi_i \in \mathcal{K} \mid \varphi_i \wedge \mu \not\models \perp\}$ is inconsistent:

Proposition 9 *Let Δ be a BM operator satisfying **(CO)**,*

(IC0) and **(IC8)**. For each profile \mathcal{K} and each formula μ , if \mathcal{K} contains a disagreement over μ , then $\bigvee_{\mathcal{K}' \in MC(\mathcal{K}, \mu)} \bigwedge \mathcal{K}' \wedge \mu \models \Delta_\mu(\mathcal{K})$.

Proposition 9 states that in presence of **(IC0)** and **(IC8)**, the disjunction of all maximal consistent subsets of $MC(\mathcal{K}, \mu)$ entails the merged base when the profile \mathcal{K} contains a disagreement over the integrity constraints μ . Yet profiles containing a disagreement over the integrity constraints are the most “interesting” ones in belief merging, as dealing with jointly inconsistent bases is precisely what makes the merging issue a non-trivial one.

Before closing this section, we show that the indecisiveness of consensus operators gets more critical in presence of additional IC postulates. Let us consider a weakening of **(IC6)** introduced in (Konieczny and Pino Pérez 2002b):

(IC6') If $\Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2)$ is consistent, then $\Delta_\mu(\mathcal{K}_1 \sqcup \mathcal{K}_2) \models \Delta_\mu(\mathcal{K}_1) \vee \Delta_\mu(\mathcal{K}_2)$.

Clearly, a BM operator satisfying **(IC6)** also satisfies **(IC6')**. Merging operators satisfying all IC postulates yet replacing **(IC6)** by **(IC6')** are called *quasi-merging operators* (Konieczny and Pino Pérez 2002b). Among them are the Max operators (Konieczny and Pino Pérez 2002b), i.e., distance-based operators using Max as an aggregation function. Interestingly, although **(IC6)** is not compatible with **(CO)** and **(IC2)**, one can find a quasi-merging operator that is a consensus one. Indeed, the drastic merging operator Δ^d is a quasi-merging operator (cf. Theorem 87 in (Konieczny 1999)) and Proposition 4 shows that it satisfies **(CO)**. However, as we argued before this operator makes poor use of information conveyed by the input profile. Actually, in presence of **(IC2)**, **(IC7)** and **(IC6')** in addition to **(IC0)** and **(IC8)**, the result given in Proposition 8 is strengthened: when dealing with any profile whose belief bases are jointly inconsistent with the integrity constraints, every consensus quasi-merging operator returns a merged base which is entailed by the disjunction of all non-valid bases, once conjoined with the integrity constraints:

Proposition 10 Let Δ be a merging operator satisfying **(CO)**, **(IC0)**, **(IC2)**, **(IC6')**, **(IC7)** and **(IC8)**. If $\bigwedge \mathcal{K} \wedge \mu$ is inconsistent, then for each base $\varphi_i \in \mathcal{K}$ such that φ_i is not valid, we have that $\varphi_i \wedge \mu \models \Delta_\mu(\mathcal{K})$.

Let us summarize: first, by Proposition 2, there is no consensus operator satisfying **(IC2)** and **(IC6)**, so that **(IC6)** must be set aside. Second, by Proposition 8, we know that one cannot find a consensus operator with a reasonable inferential power which satisfies **(IC0)** and **(IC8)** (things getting worse in presence of **(IC7)**, **(IC2)** and **(IC6')**, cf. Proposition 10). Yet **(IC0)** captures a fundamental principle whereas **(IC8)** could be considered as too strong: it is the BM counterpart of the postulate **(R6)** in belief revision (Katsuno and Mendelzon 1991), and revision operators of interest yet not satisfying **(R6)** have been proposed in the literature (Katsuno and Mendelzon 1991; Benferhat, Lagrue, and Papini 2005). Therefore, in addition to **(IC6)**, **(IC8)** appears as the best condition to be relaxed to get more interesting consensus operators. This is what we do in the following.

Pareto Consensus Operators

We now focus on a class of consensus operators, called *Pareto operators*. As distance-based operators, Pareto operators take advantage of a distance d between interpretations which defines a distance between an interpretation ω and a belief base as $d(\omega, \varphi) = \min_{\omega' \models \varphi} d(\omega, \omega')$. Every interpretation ω is then associated with a list of numbers $(\delta_1, \dots, \delta_n)$ where for each $\varphi_i \in \mathcal{K}$, $\delta_i = d(\omega, \varphi_i)$. Pairs of interpretations are then compared using the Pareto criterion on the list of numbers associated with them. The selected models of the merged base are the “closest” ones to the profile in terms of *Pareto optimality*:

Definition 1 (Pareto dominance) Given a profile \mathcal{K} , a distance d between interpretations and two interpretations ω, ω' , ω is said to (weakly) Pareto dominate ω' w.r.t. d and \mathcal{K} , noted $\omega \leq_{\mathcal{K}}^{d, Par} \omega'$, if for each $\varphi_i \in \mathcal{K}$, $d(\omega, \varphi_i) \leq d(\omega', \varphi_i)$.

Pareto operators are then formally defined as follows:

Definition 2 (Pareto operator) Given a distance d between interpretations, the Pareto merging operator based on d , denoted by $\Delta^{d, Par}$, associates with every formula μ and every profile \mathcal{K} a belief base $\Delta_\mu^{d, Par}(\mathcal{K})$ which satisfies $[\Delta_\mu^{d, Par}(\mathcal{K})] = \min([\mu], \leq_{\mathcal{K}}^{d, Par})$.

Proposition 11 For any distance d between interpretations, $\Delta^{d, Par}$ satisfies **(CO)**, **(IC0)**, **(IC1)**, **(IC2)**, **(IC3)**, **(IC4)**, **(IC5)** and **(IC7)**. It does not satisfy **(IC6)**, **(IC8)**, **(Arb)** or **(Disj)** in the general case.

Example 1 (continued) We get that (see Table 1) $\Delta_\mu^{d_D, Par}(\mathcal{K}) \equiv (a \wedge \neg c) \vee (\neg a \wedge \neg b \wedge c) \equiv \varphi_2 \vee (\varphi_1 \wedge \varphi_3)$, and $\Delta_\mu^{d_H, Par}(\mathcal{K}) \equiv (\neg a \vee \neg c) \wedge \neg b$. Both $\Delta_\mu^{d_D, Par}(\mathcal{K})$ and $\Delta_\mu^{d_H, Par}(\mathcal{K})$ are consensuses for \mathcal{K} under μ . Using $\Delta^{d_D, Par}$ leads agents 1 and 3 to refine their beliefs, and using $\Delta^{d_H, Par}$ leads agents 2 and 3 to refine their beliefs.

From this example, one can observe that $\Delta^{d_D, Par}$ and $\Delta^{d_H, Par}$ are incomparable in terms of inferential power in the general case, since here $\Delta_\mu^{d_D, Par}(\mathcal{K}) \not\models \Delta_\mu^{d_H, Par}(\mathcal{K})$ and $\Delta_\mu^{d_H, Par}(\mathcal{K}) \not\models \Delta_\mu^{d_D, Par}(\mathcal{K})$. Table 1 also suggests that all three distance-based merging operators have an inferential power stronger than the Pareto operator based on the same distance, i.e., for $d \in \{d_D, d_H\}$, $\Delta_\mu^{d, \Sigma}(\mathcal{K}) \models \Delta_\mu^{d, Par}(\mathcal{K})$ (and similarly for $\Delta_\mu^{d, GMax}(\mathcal{K})$ and $\Delta_\mu^{d, GMin}(\mathcal{K})$). This is actually always the case for any distance d and any aggregation function f satisfying the (strict monotonicity) property (Everaere, Konieczny, and Marquis 2012), (which is offered by standard aggregation functions as GMax, GMin and Σ):

Definition 3 (strict monotonicity) An aggregation function f satisfies (strict monotonicity) if $x < y \Rightarrow f(x_1, \dots, x, \dots, x_n) < f(x_1, \dots, y, \dots, x_n)$.

Proposition 12 Let d be any distance and f be an aggregation function satisfying (strict monotonicity). For every profile \mathcal{K} and formula μ , we have $\Delta_\mu^{d, f}(\mathcal{K}) \models \Delta_\mu^{d, Par}(\mathcal{K})$.

Example 1 also illustrates that the merged base obtained using Pareto operators (for any of the two distances) is not entailed by the disjunction of all three input bases, as it would be required for consensus operators satisfying **(IC0)**, **(IC2)** and **(IC8)**, since none of these three bases belong to all maximal consistent subsets of $MC(\mathcal{K}, \mu)$ (cf. point (i) of Proposition 8). This shows that while ensuring the consensus condition, Pareto operators have a reasonable inferential power compared to consensus merging operators satisfying **(IC0)** and **(IC8)**. Especially, Corollary 1 below shows that the Pareto operator based on the drastic distance has an inferential power which is higher than *any* consensus operator satisfying **(IC0)** and **(IC8)**, when dealing with profiles \mathcal{K} containing a disagreement over μ :

Proposition 13 *For any profile \mathcal{K} and formula μ , we have that $\Delta_{\mu}^{d_D, Par}(\mathcal{K}) \equiv \bigvee_{\mathcal{K}' \in MC(\mathcal{K}, \mu)} \mathcal{K}' \wedge \mu$.*

Proposition 13 states that $\Delta_{\mu}^{d_D, Par}(\mathcal{K})$ is equivalent to the disjunction of the maximal consistent subsets of $MC(\mathcal{K}, \mu)$. It is worth noting that $\Delta_{\mu}^{d_D, Par}$ corresponds to the combination function $Comb_1$ introduced in (Baral et al. 1992) and studied in (Konieczny 2000). A direct consequence of Proposition 13 and Proposition 9 is:

Corollary 1 *Let Δ be any BM operator satisfying **(CO)**, **(IC0)** and **(IC8)**. For any profile \mathcal{K} and formula μ , if \mathcal{K} contains a disagreement over μ , then $\Delta_{\mu}^{d_D, Par}(\mathcal{K}) \models \Delta_{\mu}(\mathcal{K})$.*

From IC Operators to Consensus Ones

We now explain how some consensus operators can be induced from merging ones. We determine the postulates satisfied by the induced operators, depending on the postulates satisfied by the merging operators they are based on.

Definition 4 (∇ operator) *Let Δ be a BM operator. The BM operator ∇ induced by Δ is defined by $\nabla_{\mu}(\mathcal{K}) \equiv \mu$ if $\bigvee \mathcal{K} \wedge \mu$ is inconsistent, and $\nabla_{\mu}(\mathcal{K}) \equiv \bigvee_{\varphi_i \in \mathcal{K}} \Delta_{\varphi_i \wedge \mu}(\mathcal{K} \setminus \{\varphi_i\})$ in the remaining case.³*

Ensuring conditions on Δ is enough to guarantee conditions on ∇ :

Proposition 14

- For $i \in \{0, 1, 2, 3, 7\}$, if Δ satisfies **(ICi)**, then ∇ satisfies **(ICi)**.
- If Δ satisfies **(IC0)** and **(IC1)**, then ∇ satisfies **(IC4)** and **(CO)**.
- If Δ satisfies **(IC0)**, then ∇ satisfies **(Disj)**.
- If Δ satisfies **(IC2)**, then ∇ satisfies **(Arb)**.

However, we have also that:

Proposition 15 *∇ satisfies neither **(IC5)** nor **(IC8)** in general, even if ∇ is induced by an IC merging operator.*

Let us denote $\nabla^{d, f}$ the distance-based operator ∇ operator induced by the distance-based merging operator $\Delta^{d, f}$.

³If $\mathcal{K} = \{\varphi_1\}$, then we assume that $\nabla_{\mu}(\{\varphi_1\}) \equiv \mu$ if $\varphi_1 \wedge \mu$ is inconsistent, and $\nabla_{\mu}(\{\varphi_1\}) \equiv \varphi_1 \wedge \mu$ in the remaining case.

Example 1 (continued) *Let us detail the computation of $\nabla_{\mu}^{d_D, \Sigma}(\mathcal{K})$. We have that $\nabla_{\mu}^{d_D, \Sigma}(\mathcal{K}) \equiv \Delta_{\varphi_1 \wedge \mu}^{d_D, \Sigma}(\{\varphi_2, \varphi_3\}) \vee \Delta_{\varphi_2 \wedge \mu}^{d_D, \Sigma}(\{\varphi_1, \varphi_3\}) \vee \Delta_{\varphi_3 \wedge \mu}^{d_D, \Sigma}(\{\varphi_1, \varphi_2\}) \equiv (\varphi_1 \wedge \varphi_3) \vee \varphi_2 \vee (\varphi_1 \wedge \varphi_3) \equiv \varphi_2 \vee (\varphi_1 \wedge \varphi_3)$. Then it can be checked that $\nabla_{\mu}^{d_D, \Sigma}(\mathcal{K}) \equiv \nabla_{\mu}^{d_D, GMin}(\mathcal{K}) \equiv \nabla_{\mu}^{d_D, GMax}(\mathcal{K}) \equiv \nabla_{\mu}^{d_D, Par}(\mathcal{K}) \equiv \Delta_{\mu}^{d_D, Par}(\mathcal{K})$; $\nabla_{\mu}^{d_H, \Sigma}(\mathcal{K}) \equiv \nabla_{\mu}^{d_H, GMax}(\mathcal{K}) \equiv \nabla_{\mu}^{d_H, Par}(\mathcal{K}) \equiv \Delta_{\mu}^{d_H, Par}(\mathcal{K})$; and $\nabla_{\mu}^{d_H, GMin}(\mathcal{K}) \equiv (\neg a \vee \neg c) \wedge (a \vee c) \wedge \neg b$. Note that when using $\nabla_{\mu}^{d_H, GMin}$ all agents refine their beliefs and the merged base entails $\neg b$ which is not questioned by any agent.*

Interestingly, the family of ∇ operators can be viewed as a generalization of the arbitration operators (or commutative revision operators) introduced in (Liberatore and Schaerf 1998). Let us recall that the arbitration operator \diamond induced by a belief revision operator \circ is defined by $\varphi_1 \diamond \varphi_2 = (\varphi_1 \circ \varphi_2) \vee (\varphi_2 \circ \varphi_1)$ (Liberatore and Schaerf 1998), and that, for a given BM operator Δ , the induced belief revision operator \circ_{Δ} is defined by $\varphi_1 \circ_{\Delta} \varphi_2 = \Delta_{\varphi_2}(\{\varphi_1\})$ (Konieczny and Pino Pérez 2002a). We have that:

Proposition 16 *Let Δ be a BM operator, \circ_{Δ} the induced belief revision operator, and \diamond_{Δ} the corresponding arbitration operator. We have $\varphi_1 \diamond_{\Delta} \varphi_2 \equiv \nabla_{\top}(\{\varphi_1, \varphi_2\})$.*

Note that the generalization of arbitration operators to profiles with more than two bases, consisting of the ∇ operators, differ from the generalization suggested in (Konieczny and Pino Pérez 2002a) where the merged base for a profile \mathcal{K} is defined as $\Delta_{\bigvee \mathcal{K}}(\mathcal{K})$. Indeed, a merging operator \diamond defined by $\diamond_{\mu}(\mathcal{K}) = \Delta_{\mu \wedge \bigvee \mathcal{K}}(\mathcal{K})$ where Δ is an IC merging operator, does not satisfy **(CO)** in general, even if μ is valid.

Let us finally provide some results about the inferential power of ∇ operators. First, from Example 1 one can observe (for instance) that $\nabla_{\mu}^{d_H, GMin}(\mathcal{K}) \models \Delta_{\mu}^{d_H, Par}(\mathcal{K})$: the information contained in $\Delta_{\mu}^{d_H, Par}(\mathcal{K})$ is also contained in $\nabla_{\mu}^{d_H, GMin}(\mathcal{K})$. This does not happen by accident:

Proposition 17 *Let $\nabla^{d, f}$ be the ∇ operator induced by $\Delta^{d, f}$, where d is any distance and f is an aggregation function satisfying (strict monotonicity). For any profile \mathcal{K} and formula μ such that $\bigvee \mathcal{K} \wedge \mu$ is consistent, we have $\nabla_{\mu}^{d, f}(\mathcal{K}) \models \Delta_{\mu}^{d, Par}(\mathcal{K})$.*

Actually, the inferential power of distance-based operators $\nabla^{d, f}$ may be strictly stronger than $\Delta^{d, Par}$. Indeed, Example 1 shows that $\nabla_{\mu}^{d_H, GMin}(\mathcal{K}) \not\models \Delta_{\mu}^{d_H, Par}(\mathcal{K})$. This may also be the case when considering the drastic distance, as shown by the following example:

Example 2 *Let $\mathcal{P} = \{a, b\}$, $\mathcal{K} = \{\varphi_1, \varphi_2, \varphi_3, \varphi_4\}$, $\varphi_1 = \{\neg a\}$, $\varphi_2 = \{\neg b\}$, $\varphi_3 = \varphi_4 = \{a \vee b\}$, and $\mu = \top$. We get that $\Delta_{\mu}^{d_D, \Sigma}(\mathcal{K}) \equiv \neg a \vee \neg b$, whereas $\nabla_{\mu}^{d_D, \Sigma}(\mathcal{K}) \equiv a \Leftrightarrow \neg b$. Hence, $\nabla_{\mu}^{d_D, \Sigma}(\mathcal{K}) \not\models \Delta_{\mu}^{d_D, \Sigma}(\mathcal{K})$.*

Thus, while Pareto operators exhibit a better behaviour than ∇ operators in terms of ‘‘core’’ IC properties (Pareto operators satisfy **(IC5)** and ∇ operators do not), distance-based $\nabla^{d, f}$ operators have an inferential power stronger than Pareto ones (under conditions on f which are often met).

Related Work

Notions of consensus have been considered in several AI settings, e.g., in multi-agent systems (Ephrati and Rosenschein 1996), where a consensus is derived from voting agents and results in maximizing some social utility, or within the Dempster-Shafer theory (Jøsang 2002); however, none of these works is about belief merging. The other way around, (Creignou et al. 2016) introduce refined BM operators, suited to some specific fragments of propositional logic (e.g., the Horn fragment). Among other things, the authors show that no refinement of a distance-based merging operator relying on Σ or GMax satisfies **(IC6)** when the Horn fragment is targeted. This is similar to what happens in our setting when **(CO)** must be satisfied in presence of **(IC2)**. That said, (Creignou et al. 2016) is not about consensus.

(Konieczny, Marquis, and Schwind 2011) point out independence postulates corresponding to some rationalization principles (the fact of “repairing” the input bases for discarding the models of them which do not satisfy the integrity constraints). Similarly to our work, the authors study how those postulates interact with the IC postulates. That said, the motivations of the two works are dissimilar. Indeed, reaching a consensus is useful when the agents involved are reluctant to change, while making a rationalization step consists in modifying the input bases. Thus, the results obtained in (Konieczny, Marquis, and Schwind 2011) are different from ours. Notably, the property of independence by rationalization by expansion or revision is compatible with the IC postulates, while consensus is not.

(Benferhat, Lagrue, and Rossit 2014) analyze sum-based merging operators suited to profiles \mathcal{K} of ranked belief bases (such bases K are multisets of pairs (φ, k) where k is the rank of the formula φ , a positive integer reflecting its plausibility). In this setting, every interpretation ω has a rank for K , equal to 0 when ω satisfies every formula in K and to the largest rank of a formula of K that ω falsifies otherwise. The rank of ω for \mathcal{K} is the sum of its ranks for each base from \mathcal{K} . The models of the beliefs corresponding to a profile \mathcal{K} w.r.t. the operator Δ_Σ are the interpretations having the smallest rank for \mathcal{K} . The authors then introduce a merging operator \blacktriangle_Σ suited to the case when no commensurability assumption between the ranks considered in the bases of \mathcal{K} can be made. A compatible scale is defined as any mapping updating the ranks of the formulae in every base φ of \mathcal{K} while preserving the relative ordering of the formulae in φ . As to \blacktriangle_Σ , ω_1 is at least as preferred as ω_2 for \mathcal{K} if the rank of ω_1 is not strictly greater than the rank of ω_2 for any profile obtained from \mathcal{K} by applying a compatible scale to it. At a first glance, some concepts look quite similar to our own work. To be more precise, the authors define a consensus condition:

(CSS) $\forall \varphi_i \in \mathcal{K}$, if $\varphi_i \models \mu$, then $\varphi_i \wedge \Delta_\mu(\mathcal{K})$ is consistent,

and show that, unlike Δ_Σ , \blacktriangle_Σ satisfies it. They also show that \blacktriangle_Σ can be characterized by a Pareto ordering over propositional interpretations, and does not satisfy any of the postulates **(IC6)** and **(IC8)**, just as our Pareto-based consensus operators. However, a closer inspection shows that our results differ from those reported in (Benferhat, Lagrue, and Rossit 2014) on many aspects. This can be explained

by the fact that the two merging settings under consideration — ranked belief bases merging vs. “flat” belief bases distance-based merging — are quite different. In the former setting, the same formula can be associated with two distinct ranks in two distinct bases, so that the rank of an interpretation ω for a base φ does not solely depend on ω and on the formulae occurring in φ . Contrastingly, in our setting, no ranks are considered and the distance of an interpretation ω to φ depends only on ω and the models of φ . Thus, \blacktriangle_Σ does not belong to our family of Pareto consensus operators. Furthermore, while they look similar, the **(CSS)** postulate is actually strictly weaker than **(CO)**: the antecedent part of **(CSS)**, $\varphi_i \models \mu$, is strictly stronger than the antecedent part of **(CO)**, $\varphi_i \wedge \mu$ is consistent. Finally, most results in (Benferhat, Lagrue, and Rossit 2014) are about the issue of determining the postulates satisfied by Δ_Σ and \blacktriangle_Σ , while our main results are about the compatibility of **(CO)** with the IC postulates (thus, one does not focus on specific operators). In particular, our incompatibility result (cf. Proposition 2) is not of the same nature as the “impossibility” result provided in (Benferhat, Lagrue, and Rossit 2014) (cf. Proposition 12) which concerns Δ_Σ .

(Gauwin, Konieczny, and Marquis 2007) investigate how merging operators can be used to reach a consensus through a *conciliation* operator, i.e., a mapping from the set of belief profiles to the set of belief profiles. Conciliation operators are defined iteratively: at each step, every agent *revises* her belief according to the merged base, from which a new belief profile is derived. Then, this merge-then-revise process is repeated until a consensus is eventually reached. This approach is different from ours in several aspects. First, conciliation operators are induced by merging and revision operators, while our notion of consensus operator does not rely on a revision operator. Second, several iterations are typically required in a conciliation process to get a consensus while our consensus operators are required to achieve the goal in a single step. Third, in (Gauwin, Konieczny, and Marquis 2007) a consensus is reached when $\bigwedge \mathcal{K} \wedge \mu$ is consistent, whereas our consensus condition only requires each base consistent with μ to be consistent with the merged base (which does not impose to reach a common agreement).

(Everaere, Konieczny, and Marquis 2008) introduced:

(temperance) $\forall \varphi_i \in \mathcal{K}$, $\varphi_i \wedge \Delta_\top(\mathcal{K})$ is consistent.

which is close to **(CO)**. Actually, **(temperance)** corresponds to the specific case of **(CO)** when the integrity constraint μ is valid. The authors prove a result similar to Proposition 2, showing that **(temperance)** is incompatible with the conjunction of **(IC2)** and **(IC6)**. They also introduce a merging operator $\Delta_\mu^{\text{diff}, \oplus}$ based on a notion of closeness $\text{diff}(\omega, \varphi)$, measured as the set of all minimal subsets of variables (w.r.t. \subseteq) which have to be flipped in ω for making it a model of φ . The aggregation function used \oplus is set-product. Among other things, the authors show that $\Delta_\mu^{\text{diff}, \oplus}$ satisfies **(temperance)**. However, $\Delta_\mu^{\text{diff}, \oplus}$ does not satisfy **(CO)**: when $\varphi_1 = \{-a\}$, $\varphi_2 = \{a \wedge b\}$, and $\mu = (a \vee \neg b) \wedge (\neg a \vee b)$, we have that $\Delta_\mu^{\text{diff}, \oplus}(\{\varphi_1, \varphi_2\}) \equiv a \wedge b$; and thus $\varphi_1 \wedge \mu$ is consistent, while $\varphi_1 \wedge \Delta_\mu^{\text{diff}, \oplus}(\{\varphi_1, \varphi_2\})$ is inconsistent. As a consequence, $\Delta_\mu^{\text{diff}, \oplus}$ is neither a Pareto consensus operator

nor a ∇ operator induced by an IC merging operator.

Lastly, our notion of consensus C for \mathcal{K} under μ departs from the one considered in (Grégoire, Konieczny, and Lagniez 2016), in which it is required that $C \subseteq \bigcup_{i=1}^n \varphi_i \cup \{\mu\}$ and $\{\mu\} \subseteq C$. An important difference concerns the syntax-sensitivity dimension. In (Grégoire, Konieczny, and Lagniez 2016), replacing any $\varphi_i \in \mathcal{K}$ by an equivalent base may easily lead to a different consensus. This is not the case in our approach as soon as the merging operator satisfies **(IC3)** (and this is ensured by the merging operators we focus on). Clearly enough, when $\mu \equiv \top$, our notion of consensus generalizes the one considered in (Grégoire, Konieczny, and Lagniez 2016): a consensus should be a (conjunctively interpreted) subset of the union of all bases from the profile, while in our approach it could be any propositional formula. For this reason, using our approach, more information-preserving consensus can be established in many cases. For instance, if $\mathcal{K} = \{\{a \wedge b\}, \{\neg a \wedge b\}\}$ (and $\mu \equiv \top$), the unique consensus according to (Grégoire, Konieczny, and Lagniez 2016) is the empty set, while in our approach $\{b\}$ is an acceptable consensus. Lastly, while the main contribution of (Grégoire, Konieczny, and Lagniez 2016) is about computational aspects of consensus generation, our work can be viewed instead as an axiomatic study of the consensus condition in propositional merging.

Conclusion

In this paper, we have stated a consensus condition in propositional BM and proved that it is incompatible with the standard IC postulates. We have then introduced two classes of consensus operators, satisfying a large subset of IC postulates and exhibiting a reasonable behaviour from the inferential standpoint.

A perspective for further research consists in identifying the computational complexity of the inference problem for the Pareto and ∇ operators. The question whether a representation theorem can be found in order to characterize all operators that satisfy **(CO)** plus all standard postulates except **(IC6)** and **(IC8)**, will also be investigated.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number JP17K12746.

References

- Baral, C.; Kraus, S.; Minker, J.; and Subrahmanian, V. S. 1992. Combining knowledge bases consisting of first-order theories. *Computational Intelligence* 8(1):45–71.
- Benferhat, S.; Lagrue, S.; and Papini, O. 2005. Revision of partially ordered information: Axiomatization, semantics and iteration. In *IJCAI'05*, 376–381.
- Benferhat, S.; Lagrue, S.; and Rossit, J. 2014. Sum-based weighted belief base merging: From commensurable to incommensurable framework. *International Journal of Approximate Reasoning* 55(9):2083–2108.
- Creignou, N.; Papini, O.; Rümmele, S.; and Woltran, S. 2016. Belief merging within fragments of propositional logic. *ACM Trans. Comput. Log.* 17(3):20:1–20:28.
- Ephrati, E., and Rosenschein, J. S. 1996. Deriving consensus in multiagent systems. *Artificial Intelligence* 87(1-2):21–74.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2008. A diff-based merging operator. In *NMR'08*, 19–25.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2010. Disjunctive merging: Quota and gmin merging operators. *Artificial Intelligence* 174(12-13):824–849.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2012. Compositional belief merging. In *KR'12*, 603–607.
- Everaere, P.; Konieczny, S.; and Marquis, P. 2014. On egalitarian belief merging. In *KR'14*, 121–130.
- Gauwin, O.; Konieczny, S.; and Marquis, P. 2007. Conciliation through iterated belief merging. *Journal of Logic and Computation* 17(5):909–937.
- Grégoire, E.; Konieczny, S.; and Lagniez, J.-M. 2016. On consensus extraction. In *IJCAI'16*, 1095–1101.
- Jøsang, A. 2002. The consensus operator for combining beliefs. *Artificial Intelligence* 141(1/2):157–170.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.
- Konieczny, S., and Pérez, R. P. 1999. Merging with integrity constraints. In *ECSQARU'99*, 233–244.
- Konieczny, S., and Pino Pérez, R. 2002a. Merging information under constraints: a logical framework. *Journal of Logic and Computation* 12(5):773–808.
- Konieczny, S., and Pino Pérez, R. 2002b. On the frontier between arbitration and majority. In *KR'02*, 109–118.
- Konieczny, S., and Pino Pérez, R. 2011. Logic based merging. *Journal of Philosophical Logic* 40:239–270.
- Konieczny, S.; Lang, J.; and Marquis, P. 2004. DA² merging operators. *Artificial Intelligence* 157:49–79.
- Konieczny, S.; Marquis, P.; and Schwind, N. 2011. Belief base rationalization for propositional merging. In *IJCAI'11*, 951–956.
- Konieczny, S. 1999. *Sur la logique du changement : révision et fusion de bases de connaissance*. Ph.D. Dissertation, Université de Lille 1.
- Konieczny, S. 2000. On the difference between merging knowledge bases and combining them. In *KR'00*, 135–144.
- Liberatore, P., and Schaerf, M. 1998. Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering* 10:76–90.
- Lin, J. 1996. Integration of weighted knowledge bases. *Artificial Intelligence* 83:363–378.
- Revesz, P. Z. 1997. On the semantics of arbitration. *International Journal of Algebra and Computation* 7:133–160.
- Schwind, N.; Inoue, K.; Bourgne, G.; Konieczny, S.; and Marquis, P. 2015. Belief revision games. In *AAAI'15*, 1590–1596.
- Schwind, N.; Inoue, K.; Bourgne, G.; Konieczny, S.; and Marquis, P. 2016. Is promoting beliefs useful to make them accepted in networks of agents? In *IJCAI'16*, 1237–1243.