Trends in Computational Social Choice

7

Cite as: Patricia Everaere, Sébastien Konieczny, and Pierre Marquis. An Introduction to Belief Merging and its Links with Judgment Aggregation. In Ulle Endriss (editor), *Trends in Computational Social Choice*, chapter 7, pages 123–143. AI Access, 2017.

http://www.illc.uva.nl/COST-IC1205/Book/

CHAPTER 7

An Introduction to Belief Merging and its Links with Judgment Aggregation

Patricia Everaere, Sébastien Konieczny, and Pierre Marquis

7.1 Introduction

Belief merging aims at defining the beliefs of a group of agents from their individual beliefs and some integrity constraints to be respected. This objective appears as non-trivial as soon as some conflicts between the individual beliefs and possibly involving the constraints exist. For this reason, the belief merging issue has raised much attention in artificial intelligence for more than two decades. Many results of various nature have been obtained, including belief merging operators (based on several intuitions), postulates for delineating the rational ones, representation theorems establishing constructive ways of defining some belief merging operators, the identification of the complexity of some operators, some comparisons of their inferential powers, some implementations, and other results connected to several issues which are often addressed in social choice, like strategy-proofness or truth-tracking.

The main goal of this chapter is to give an introduction to propositional belief merging and to stress some of the links and differences between propositional belief merging (BM) and judgment aggregation (JA). While BM and JA have been developed mainly independently, they are two logically-founded theories of logical aggregation, with similar objectives. Accordingly, some connections between the two have already been exploited in some previous works. Thus, Pigozzi (2006) shows how one can take advantage of a BM operator for defining a JA one. Studying what make BM and JA close and what make them different is useful for a better understanding of the pros and the cons of the two theories.

Let us illustrate what a belief merging process is on a simple scenario. Consider three agents 1, 2 and 3, each of them associated with a belief base K_1, K_2 and K_3 (respectively), as given in Figure 7.1 (we do not consider any integrity constraints — stated otherwise, the integrity constraint is a tautology).

The objective of belief merging is to define the beliefs of the group of agents (the merged base) from the individual beliefs of its members. For this example,

K_1	K_2	K_3
$a, b \to c$	a, b	$\neg a$

Figure 7.1: An example of belief merging: $\Delta(K_1, K_2, K_3) = ?$

one cannot simply consider the conjunction (union) of the bases as the result of the merging process since this conjunction is not consistent. A closer look at the bases shows that a is a conflicting piece of belief, while no conflict involves b or $b \to c$. If b and $b \to c$ are kept in the merged base, then c can be inferred at the group level, though none of the agents can draw such a conclusion alone. Accordingly, "new" pieces of belief, i.e., beliefs that no agent can infer alone, can be generated during the merging process. Now, if one wants to go further in the merging process and get some information about a, a majority argument can be used. Indeed, two out of the three agents believe that a is true, and this can be considered as a sufficient reason for considering a also as a piece of belief at the group level. Note that in belief merging adhering to such a majority principle is not mandatory, but there is a subclass of majority merging operators which are based on it. Similarly, there also exists a subclass of so-called arbitration operators, which aim at defining a merged base which is as close as possible to the base of each agent.

In the following, after a presentation of BM and JA which aims at introducing a number of key concepts and postulates, we focus on the relations and differences between BM and JA. For this purpose, we investigate the question of how to define the judgment set of an agent given her belief base and an agenda. On this ground, we show that the beliefs produced by a BM operator and those produced by a JA one can easily be incompatible, even if the two operators satisfy some rationality conditions. Interestingly, in the restricted case when the two approaches are equally informed (i.e., when the agenda is the set of all interpretations), every merging operator can be associated with a judgment aggregation operator, and *vice versa*. We show that some close connections can be established, linking the satisfaction of some postulates by the pairs of operators that correspond to each other.

7.2 On Belief Merging

We consider a propositional language $\mathcal{L}$ defined from a finite set $\mathcal{P}$ of propositional symbols and the usual connectives, including the Boolean constants $\top$ and $\bot$.

An interpretation (or state of the world) ω is a total function from $\mathcal{P}$ to $\{0, 1\}$. The set of all interpretations is denoted by $\mathcal{W}$. An interpretation ω is usually represented by a bit vector whenever a strict total order on $\mathcal{P}$ is specified. It can also be viewed as the complete formula $\bigwedge_{p \in \mathcal{P} | \omega(p)=1} p \wedge \bigwedge_{p \in \mathcal{P} | \omega(p)=0} \neg p$.

The symbol $\models$ denotes the logical entailment relation and $\equiv$ the logical equivalence relation. The set of models of a formula φ is denoted by $[\varphi]$, i.e., $[\varphi] = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$.

A *belief base* K is a finite set of propositional formulae $\{\varphi_1, \ldots, \varphi_k\}$. We denote by $\bigwedge K$ the conjunction of the formulae of K, i.e., $\bigwedge K = \varphi_1 \wedge \ldots \wedge \varphi_k$. In order to simplify the notation, we often identify[1] a base K with the formula $\bigwedge K$. We suppose that each belief base is consistent, and denote by $\mathcal{K}$ the set of all bases.

A *profile* E represents the beliefs of a group of n agents involved in the merging process; formally E is given by a vector $(K_1, \ldots, K_n)$ of belief bases, where K_i is the belief base of agent i (different agents are allowed to exhibit identical bases). The conjunction of all elements of E is denoted $\bigwedge E$, i.e., $\bigwedge E = \bigwedge K_1 \wedge \ldots \wedge \bigwedge K_n$ and if $E = (K_1, \ldots, K_i)$ and $E' = (K'_1, \ldots, K'_j)$, $E \sqcup E'$ denotes the profile $(K_1, \ldots, K_i, K'_1, \ldots, K'_j)$. $\mathcal{E}$ is the set of all profiles. A profile E is said to be consistent if and only if $\bigwedge E$ is consistent.

We denote by E^p the profile $E^p = \underbrace{E \sqcup \ldots \sqcup E}_{p}$. Two profiles $E = (K_1, \ldots, K_n)$ and $E' = (K'_1, \ldots, K'_n)$ are equivalent, denoted $E \equiv E'$, if there exists a permutation π over $\{1, \ldots, n\}$ such that for each $i \in 1, \ldots, n$, we have $K_i \equiv K'_{\pi(i)}$. If $\leqslant$ is a preorder on $\mathcal{W}$ (i.e., a reflexive and transitive relation), then $<$ denotes the associated strict order defined by $\omega < \omega'$ if and only if $\omega \leqslant \omega'$ and $\omega' \nleqslant \omega$. A preorder is *total* if $\forall \omega, \omega' \in \mathcal{W}$, $\omega \leqslant \omega'$ or $\omega' \leqslant \omega$. A preorder that is not total is called *partial*. If $\leqslant$ is a preorder on A, and $B \subseteq A$, then $\min(B, \leqslant) = \{b \in B \mid \nexists a \in B \; a < b\}$.

An *integrity constraint* μ is a consistent formula restricting the possible results of the merging process.

Merging operators are mappings from the set of profiles and the set of propositional formulae (that represent integrity constraints) to the set of bases, i.e. $\Delta : \mathcal{E} \times \mathcal{L} \to \mathcal{K}$. We use the notation $\Delta_\mu(E)$ instead of $\Delta(E, \mu)$. $\Delta(E)$ is short for $\Delta_\top(E)$.

We first present the main logical properties pointed out for characterizing the IC merging operators and recall a representation theorem for them, expressed in terms of preorders on interpretations. Some of these properties had been proposed by Revesz (1997) in order to define *model fitting* operators. They have been extended by Konieczny and Pino Pérez (2002b).

Definition 7.1. *A merging operator* Δ *is an* IC merging operator *if it satisfies the following properties (the so-called IC postulates):*

(IC0) $\Delta_\mu(E) \models \mu$

(IC1) *If* μ *is consistent, then* $\Delta_\mu(E)$ *is consistent*

(IC2) *If* $\bigwedge E \wedge \mu$ *is consistent, then* $\Delta_\mu(E) \equiv \bigwedge E \wedge \mu$

(IC3) *If* $E_1 \equiv E_2$ *and* $\mu_1 \equiv \mu_2$*, then* $\Delta_{\mu_1}(E_1) \equiv \Delta_{\mu_2}(E_2)$

(IC4) *If* $K_1 \models \mu$ *and* $K_2 \models \mu$*, then* $\Delta_\mu((K_1, K_2)) \wedge K_1$ *is consistent if and only if* $\Delta_\mu((K_1, K_2)) \wedge K_2$ *is consistent*

(IC5) $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \models \Delta_\mu(E_1 \sqcup E_2)$

(IC6) *If* $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$ *is consistent, then* $\Delta_\mu(E_1 \sqcup E_2) \models \Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$

(IC7) $\Delta_{\mu_1}(E) \wedge \mu_2 \models \Delta_{\mu_1 \wedge \mu_2}(E)$

[1] This identification is done when the BM operator under consideration is not sensitive to the syntactical representation of the bases. Otherwise, it is important to make a distinction between a base K and the conjunction of its formulae (see e.g., Konieczny et al., 2004).

(IC8) *If* $\Delta_{\mu_1}(E) \wedge \mu_2$ *is consistent, then* $\Delta_{\mu_1 \wedge \mu_2}(E) \models \Delta_{\mu_1}(E)$

The meaning of the postulates is as follows: when satisfied, (IC0) ensures that the merged base satisfies the integrity constraints. (IC1) states that if the integrity constraints are consistent, then the merged base is consistent as well. (IC2) states that the merged base is the conjunction of the belief bases with the integrity constraints when this conjunction is consistent. (IC3) is the principle of irrelevance of syntax, i.e., if two profiles are equivalent and two integrity constraints bases are logically equivalent then the corresponding merged bases are logically equivalent. (IC4) is a fairness postulate: when two belief bases are merged, no preference has to be given to one of them. (IC5) expresses the following idea: if two profiles E_1 and E_2 agree on some models then these models must be chosen if the two profiles are joined. (IC5) and (IC6) together state that if the merged bases corresponding to two profiles agree on some models, then if the two profiles are joined, the models of the corresponding merged base must be those models for which there is an agreement. (IC7) and (IC8) can be viewed as a direct generalization of the (R5-R6) postulates for belief revision (Katsuno and Mendelzon, 1991). They state some conditions about conjunctions of integrity constraints. Actually, they ensure that the notion of *closeness* one wants to capture is well-behaved. If a model ω is chosen in the set of possible models $[\mu]$, then if the set of possible models is narrowed but ω still belongs to the resulting set, it still must be selected. Similar properties to this quite natural requirement appear in different social choice theories.

The IC properties are the basic ones one could expect for BM operators. Some additional requirements can be considered for constraining further the behavior of the merging operators. Especially, two important subclasses of IC merging operators consist of the majority operators and the arbitration operators. First of all, a *majority merging operator* is an IC merging operator that satisfies the following *majority* postulate:

(Maj) $\exists n \; \Delta_{\mu}(E_1 \sqcup E_2^n) \models \Delta_{\mu}(E_2)$

This postulate expresses the fact that if a subgroup is repeated sufficiently many times in a profile then the opinion of this subgroup must prevail. Majority merging operators aim at satisfying the group of agents as a whole. Contrastingly, arbitration operators aim at satisfying each agent of the group as far as possible. Formally, an *arbitration operator* is an IC merging operator that satisfies the following *arbitration* postulate:

(Arb) If $\Delta_{\mu_1}(K_1) \equiv \Delta_{\mu_2}(K_2), \Delta_{\mu_1 \leftrightarrow \neg\mu_2}((K_1, K_2)) \equiv (\mu_1 \leftrightarrow \neg\mu_2), \mu_1 \not\models \mu_2$, and $\mu_2 \not\models \mu_1$, then $\Delta_{\mu_1 \vee \mu_2}((K_1, K_2)) \equiv \Delta_{\mu_1}(K_1)$

This property, which is much more intuitive when it is expressed in a model-theoretical way (cf. condition 8 of a fair syncretic assignment in Definition 7.2), roughly states that "median models" must be preferred.

We now present some representation theorems that give more constructive ways to define rational BM operators than the previous postulates. Such theorems show that each IC merging operator corresponds to a family of preorders on interpretations. First, one needs to define the notion of *syncretic assignment*.

Definition 7.2. *A* syncretic assignment *is a mapping associating with each profile E a total preorder $\leqslant_E$ over $\mathcal{W}$ such that for any profile E, E_1, E_2 and for any belief base K, K' the following conditions hold:*

1. *If $\omega \models E$ and $\omega' \models E$, then $\omega \simeq_E \omega'$*
2. *If $\omega \models E$ and $\omega' \not\models E$, then $\omega <_E \omega'$*
3. *If $E_1 \equiv E_2$, then $\leqslant_{E_1} = \leqslant_{E_2}$*
4. *$\forall \omega \models K \ \exists \omega' \models K' \ \omega' \leqslant_{(K,K')} \omega$*
5. *If $\omega \leqslant_{E_1} \omega'$ and $\omega \leqslant_{E_2} \omega'$, then $\omega \leqslant_{E_1 \sqcup E_2} \omega'$*
6. *If $\omega <_{E_1} \omega'$ and $\omega \leqslant_{E_2} \omega'$, then $\omega <_{E_1 \sqcup E_2} \omega'$*

A **majority syncretic assignment** *is a syncretic assignment which satisfies the following condition:*

7. *If $\omega <_{E_2} \omega'$, then $\exists n \ \omega <_{E_1 \sqcup E_2^n} \omega'$*

A **fair syncretic assignment** *is a syncretic assignment which satisfies the following condition:*

8. *If $\omega <_{K_1} \omega'$, $\omega <_{K_2} \omega''$, and $\omega' \simeq_{(K_1,K_2)} \omega''$, then $\omega <_{(K_1,K_2)} \omega'$*

The two first conditions ensure that the models of the conjunction of the bases from the profile (if any) are the most plausible interpretations for the preorder associated with the profile. The third condition states that two equivalent profiles are associated with the same preorders. These first three conditions are very close to the ones considered in belief revision for defining faithful assignments (Katsuno and Mendelzon, 1991). The fourth condition states that, when merging two belief bases, for each model of the first one, there is a model of the second one that is at least as good as the first one. It ensures that the two bases receive equal treatments in the merging process. The fifth condition states that if an interpretation ω is at least as plausible as an interpretation ω' for a profile E_1 and if ω is at least as plausible as ω' for a profile E_2, and if one then joins the two profiles, then ω must be at least as plausible as ω'. The sixth condition strengthens the previous condition by stating that if an interpretation ω is at least as plausible as an interpretation ω' for a profile E_1 and if ω is strictly more plausible than ω' for a profile E_2, then if the two profiles are joined, ω must be strictly more plausible than ω'. These two conditions are very close the to Pareto conditions in social choice. The seventh condition states that if an interpretation ω is strictly more plausible than an interpretation ω' for a profile E_2, then there is a quorum n of repetitions of the profile E_2 such that ω is more plausible than ω' for the larger profile $E_1 \sqcup E_2^n$. This condition seems to be the weakest form of "majority" condition one could state. Finally, the eighth condition states that "median choices" must be preferred by the group. More precisely, if an interpretation ω is more plausible than an interpretation ω' for a belief base K_1, if ω is more plausible than ω'' for another base K_2, and if ω' and ω'' are equally plausible for the joint profile (K_1, K_2), then ω has to be more plausible than ω' and ω'' for (K_1, K_2).

The following representation theorems have been established:

Proposition 7.1 (Konieczny and Pino Pérez, 2002a). *A merging operator Δ is an IC merging operator (resp. a majority merging operator, an arbitration operator) if and only if there exists a syncretic assignment (resp. a majority syncretic assignment, a fair syncretic assignment) that maps each profile E to a total preorder $\leqslant_E$ over $\mathcal{W}$ such that $mod(\Delta_\mu(E)) = \min(mod(\mu), \leqslant_E)$*

7.3 On Distance-Based Merging Operators

Let us now give some examples of IC merging operators from the family of distance-based merging operators (Konieczny et al., 2004). For such operators, the total preorders $\leqslant_E$ generated by the corresponding assignments are induced from a distance between interpretations and an aggregation function: an interpretation ω is at least as close to E as an interpretation ω', i.e., $\omega \leqslant_E \omega'$, if the (aggregated) distance of ω to E is lower than or equal to the (aggregated) distance of ω' to E. Formally:

Definition 7.3. *A* (pseudo-)distance *between interpretations is a mapping $d : \mathcal{W} \times \mathcal{W} \rightarrow \mathbb{R}^+$ such that for any $\omega_1, \omega_2 \in \mathcal{W}$:*

- $d(\omega_1, \omega_2) = d(\omega_2, \omega_1)$
- $d(\omega_1, \omega_2) = 0$ *if and only if* $\omega_1 = \omega_2$

Typical distances are the Hamming distance d_H, that is the number of propositional letters on which the two interpretations differ and the drastic distance d_D, defined as $d_D(\omega_1, \omega_2) = 0$ if $\omega_1 = \omega_2$, and $= 1$ otherwise.

Definition 7.4. *An* aggregation function *is a mapping*[2] *f from $\mathbb{R}^m$ to $\mathbb{R}$, which satisfies:*

- *if* $x_i \geqslant x'_i$, *then* $f(x_1, ..., x_i, ..., x_m) \geqslant f(x_1, ..., x'_i, ..., x_m)$ **(non-decreasingness)**
- $f(x_1, \ldots, x_m) = 0$ *if* $\forall i,\ x_i = 0$ **(minimality)**
- $f(x) = x$ **(identity)**
- *If σ is a permutation over $\{1, \ldots, m\}$, then* $f(x_1, \ldots, x_m) = f(x_{\sigma(1)}, \ldots, x_{\sigma(m)})$ **(symmetry)**

Some additional properties can also be considered for f, especially:

- *if* $x_i > x'_i$, *then* $f(x_1, ..., x_i, ..., x_m) > f(x_1, ..., x'_i, ..., x_m)$ **(strict non-decreasingness)**
- *If* $f(x_1, \ldots, x_n) \leqslant f(y_1, \ldots, y_n)$, *then* $f(x_1, \ldots, x_n, z) \leqslant f(y_1, \ldots, y_n, z)$ **(composition)**
- *If* $f(x_1, \ldots, x_n, z) \leqslant f(y_1, \ldots, y_n, z)$, *then* $f(x_1, \ldots, x_n) \leqslant f(y_1, \ldots, y_n)$ **(decomposition)**
- *If* $\forall i, z > y_i$, *then* $f(z, x_1, \ldots, x_n) > f(y_1, \ldots, y_{n+1})$ **(strict preference)**

Standard aggregation functions are Σ (sum), $Gmax$ (also referred to as $leximax$), $Gmin$ (also referred to as $leximin$), and Σ^n (sum of the n^{th} powers).

[2] Strictly speaking, it is a family of mappings, one for each integer $m \geqslant 1$.

Definition 7.5. *Let d and f be a distance between interpretations and an aggregation function, respectively. The* distance-based merging operator $\Delta^{d,f}$ *is defined semantically by*

$$[\Delta_{\mu}^{d,f}(E)] = \min([\mu], \leqslant_E^{d,f})$$

where the total preorder $\leqslant_E$ on $\mathcal{W}$ is defined in the following way :

- $\omega \leqslant_E^{d,f} \omega'$ *if and only if* $d^{d,f}(\omega, E) \leqslant d^{d,f}(\omega', E)$
- $d^{d,f}(\omega, (K_1, \ldots, K_n)) = f(d(\omega, K_1), \ldots, d(\omega, K_n))$
- $d(\omega, K) = \min_{\omega' \models K} d(\omega, \omega')$

Example 7.1. *As a matter of illustration, consider the three belief bases: $K_1 = \{\neg a \wedge \neg b \wedge \neg c\}$, $K_2 = \{(a \wedge b) \vee (\neg a \wedge \neg b \wedge c)\}$ and $K_3 = \{a \wedge b \wedge c\}$. We have $[K_1] = \{000\}$, $[K_2] = \{001, 110, 111\}$ and $[K_3] = \{111\}$. Suppose that the integrity constraints are $\mu = (a \vee c) \wedge (a \wedge c \rightarrow b)$.*

The following table reports the merged bases corresponding to the merging of the profile (K_1, K_2, K_3) under μ for some of the most usual distance-based operators. The lines of the table correspond to the available interpretations. In this example, three propositional symbols a, b and c are considered, so there are 8 interpretations (6 being models of μ). The first three columns give the Hamming distance of the models to each base, the last four columns indicate the aggregated distances of the models to E, depending on the chosen aggregation function. The selected interpretations (depending on the chosen aggregation function) are boldfaced.

This example clearly shows that different BM operators can lead to different merged bases. Δ^{Gmin,d_H} and Δ^{Σ,d_H} are majority merging operators, so they tend to select interpretations satisfying a maximal number of bases. Δ^{Gmax,d_H} is an arbitration operator, so it tends to select "median interpretations".

	K_1	K_2	K_3	Δ^{Gmax,d_H}	Δ^{Gmin,d_H}	Δ^{Σ,d_H}	Δ^{Σ^2,d_H}
001	1	0	2	**210**	012	**3**	**5**
011	2	1	1	211	112	4	6
100	1	1	2	211	112	4	6
110	2	0	1	**210**	012	**3**	**5**
111	3	0	0	300	**003**	**3**	9

On this example, 111 is selected by the operators based on $Gmin$ or Σ because they are majoritarian operators, and $\Delta^{Gmin,d_H} \equiv a \wedge b \wedge c$. 001 or 110 are selected by the operators based on $Gmax$ or Σ^2, because these interpretations are more consensual than the other ones. Thus, we have $\Delta^{Gmax,d_H} \equiv (\neg a \wedge \neg b \wedge c) \vee (a \wedge b \wedge \neg c)$.

For usual aggregation functions, whatever the chosen distance, the corresponding distance-based BM operators exhibit good logical properties:

Proposition 7.2 (Konieczny and Pino Pérez, 2002b). *For any distance d, if f is equal to Σ, $Gmax$, $Gmin$, or Σ^n, then $\Delta^{d,f}$ is an IC merging operator.*

More generally, in (Konieczny et al., 2004), a necessary and sufficient condition on the chosen aggregation function f is identified, ensuring that the corresponding distance-based BM operators $\Delta^{d,f}$ are IC ones (whatever the distance d):

Proposition 7.3. *Let d and f be a distance between interpretations and an aggregation function respectively. The operator $\Delta^{d,f}$ satisfies the postulates (IC0-IC8) iff the aggregation function f satisfies composition and decomposition.*

To conclude this introduction to BM, we sketch some alternative approaches. Though we mainly focused on model-based BM operators in this introduction, it must be noted that formula-based BM operators have also been defined (Baral et al., 1991, 1992). The general principle underlying them is to select some preferred consistent subsets of formulae from the union of all the bases of the input profile. One important limitation of the formula-based BM approaches is to possibly forget some important pieces of information available in the input profile, such as the number of bases supporting each formula. On the other hand, inconsistent belief bases can be taken into account easily by such approaches. Formula-based BM operators have been shown to satisfy less postulates than the model-based ones (Konieczny, 2000). Konieczny et al. (2004) generalize the family of distance-based BM operators to so-called DA^2 operators, in order to take advantage of the pros offered by distance-based BM operators and by formula-based operators. Thus, DA^2 operators can deal with inconsistent belief bases and are based on two aggregation functions: the first one is used to extract pieces of belief from inconsistent bases and the second one to aggregate the resulting pieces of belief.

Everaere et al. (2008) have defined and studied conflict-based merging operators. These operators refine the distance-based ones by computing conflicts. They aim at minimizing the conflicts between the beliefs of the agents. Default-based merging operators have also been introduced by Delgrande and Schaub (2007). In this work, inconsistencies are fixed by renaming some propositional symbols, and the merged bases are characterized as those requiring "as few renamings as possible".

Additional merging postulates have also been pointed out in the literature. Let us mention the work of Everaere et al. (2010) where a **Unanimity** postulate for BM operators has been introduced. This postulate can be considered either for formulae or for models: if all agents share a common piece of belief, it should be the case that the merged base also supports this piece of belief. Everaere et al. (2010) also define a **Disjunction** postulate, which is in a certain sense a counterpoint to the arbitration postulate **(Arb)**. This postulate ensures that every logical consequence of the merged base is among the logical consequences of at least one input base. Other properties inspired by similar conditions in social choice have been translated into the BM framework. This led to the definitions of various notions of interest for merging, like truth tracking by Everaere et al. (2007), rationalization by Konieczny et al. (2011), or egalitarianism by Everaere et al. (2014a). Recently, a study of voting properties in the context of BM has also been conducted by Haret et al. (2016).

Merging has also been studied in other representation frameworks than the purely propositional one. When all the pieces of information belonging to the bases do not have the same importance, weighted approaches must be considered. Many frameworks have been defined and studied to take account of the relative plausibility of pieces of belief, including *possibilistic logic* (Dubois et al., 1994) and *ordinal conditional functions* (Spohn, 1987). Thus, Delgrande et al.

(2006) define prioritized merging operators, in order to merge sets of weighted formulae. Benferhat, Dubois, Kaci and Prade point out several merging operators suited to representations in possibilistic logic (Kaci et al., 2000; Benferhat et al., 2002).

Bloch and Lang (2000) define model-based merging operators using maximum as aggregation function ($\Delta^{d,\max}$) and show how the corresponding merged bases can be characterized via a *dilation* process. Gorogiannis and Hunter (2008) extend this approach in order to define other model-based merging operators, based on a dilation process. The interest of the dilation-based approach is that it can be extended to first-order logic without much efforts.

The issue of merging logic programs under ASP semantics has also been considered. The approach to merging given in (Hué et al., 2009) relies on the deletion of a set of formulae in the union of the bases, characterized using a selection function (the idea is close to the one considered in (Konieczny, 2000)). The corresponding operators satisfy only some IC postulates. Let us also mention the work of Delgrande et al. (2009), where the merging operators pointed out are based on the definition of a distance between stable models.

Condotta et al. (2009) studied the merging of qualitative constraint networks. Finally, Coste-Marquis et al. (2007) and Delobelle et al. (2016) study the problem of *merging argumentation frameworks*, where the arguments are distributed among several agents.

7.4 On Judgment Aggregation

Let us now briefly present some definitions and notation used in the following. An *agenda* $X = \{\varphi_1, \ldots, \varphi_m\}$ is a finite, non-empty and totally ordered set of contingent (i.e., consistent but not valid) propositional formulae. A *judgment* on a formula φ_k of X is an element of $D = \{1, 0, \star\}$, where 1 means that φ_k is supported, 0 that $\neg\varphi_k$ is supported, $\star$ that neither φ_k nor $\neg\varphi_k$ is supported. A *judgment set* on X is a mapping γ from X to D, also viewed as an m-vector over D, when the cardinality of X is m, or alternatively as the set of formulae such that $\varphi_k \in \gamma$ when $\gamma(\varphi_k) = 1$, $\neg\varphi_k \in \gamma$ when $\gamma(\varphi_k) = 0$, for every $\varphi_k \in X$. For each φ_k of X, γ is supposed to satisfy $\gamma(\neg\varphi_k) = \neg\gamma(\varphi_k)$, where $\neg\gamma$ is given by $\neg\gamma(\varphi_k) = \star$ if $\gamma(\varphi_k) = \star$, $\neg\gamma(\varphi_k) = 1$ if $\gamma(\varphi_k) = 0$, and $\neg\gamma(\varphi_k) = 0$ if $\gamma(\varphi_k) = 1$.

Judgment sets are often asked to be consistent and complete, where a judgment set is *complete* if $\forall\varphi_k \in X$, $\gamma(\varphi_k) = 0$ or $\gamma(\varphi_k) = 1$, and a judgment set γ on X is *consistent* if the associated formula (judgment) $\widehat{\gamma} = \bigwedge_{\{\varphi_k \in X \mid \gamma(\varphi_k)=1\}} \varphi_k \wedge \bigwedge_{\{\varphi_k \in X \mid \gamma(\varphi_k)=0\}} \neg\varphi_k$ is consistent.

Aggregating judgments consists in associating a set of collective judgment sets with a profile containing n individual judgment sets (one per agent): a *profile* $\Gamma = (\gamma_1, \ldots, \gamma_n)$ of judgment sets on X is a non-empty vector of judgments sets on X. Γ is *consistent* (resp. *complete*) when each judgment set in it is consistent (resp. complete).

For each agenda X, a *JA operator* Ag associates with a consistent profile Γ on X a non-empty set Ag_Γ of collective judgment sets γ_Γ on X, also viewed as a formula (the collective judgment) $\widehat{Ag_\Gamma} = \bigvee_{\gamma_\Gamma \in Ag_\Gamma} \widehat{\gamma_\Gamma}$. For $\varphi_k \in X$, we note $Ag_\Gamma(\varphi_k) =$

1 (resp. $Ag_\Gamma(\varphi_k) = 0$) if and only if $\forall \gamma_\Gamma \in Ag_\Gamma$, $\gamma_\Gamma(\varphi_k) = 1$ (resp. $\forall \gamma_\Gamma \in Ag_\Gamma$, $\gamma_\Gamma(\varphi_k) = 0$), and $Ag_\Gamma(\varphi_k) = \star$ in the remaining case.

When Ag_Γ is a singleton for each Γ, the JA operator Ag is called a resolute JA rule, and it is called an irresolute JA rule (or a JA correspondence in (Lang et al., 2011)) otherwise.

Here are some common properties for JA rules that have been identified in the literature:

Universal domain. The domain of Ag is the set of all consistent profiles.

This property is often relaxed, as in (List and Pettit, 2002), to:

C-universal domain. The domain of Ag is the set of all profiles which are consistent and complete.

Some properties also state that the result should be consistent and complete:

Collective rationality. For any profile Γ in the domain of Ag, Ag_Γ is a set of consistent collective judgment sets.

Collective completeness. For any profile Γ in the domain of Ag, Ag_Γ is a set of complete collective judgment sets.

For obvious equity reasons, agents and issues are expected to play symmetric roles:

Anonymity. For any two profiles $\Gamma = (\gamma_1, \ldots, \gamma_n)$ and $\Gamma' = (\gamma'_1, \ldots, \gamma'_n)$ in the domain of Ag which are permutations one another, we have $Ag_\Gamma = Ag_{\Gamma'}$.

Neutrality. For any φ, φ' in the agenda X and profile Γ in the domain of Ag, if $\forall i$ $\gamma_i(\varphi) = \gamma_i(\varphi')$, then $Ag_\Gamma(\varphi) = Ag_\Gamma(\varphi')$.

A more demanding property is:

Independence. For any φ in the agenda X and profiles $\Gamma = (\gamma_1, \ldots, \gamma_n)$ and $\Gamma' = (\gamma'_1, \ldots, \gamma'_n)$ in the domain of Ag, if $\forall i$ $\gamma_i(\varphi) = \gamma'_i(\varphi)$, then $Ag_\Gamma(\varphi) = Ag_{\Gamma'}(\varphi)$.

Systematicity. For any φ, φ' in the agenda X and profiles $\Gamma = (\gamma_1, \ldots, \gamma_n)$ and $\Gamma' = (\gamma'_1, \ldots, \gamma'_n)$ in the domain of Ag, if $\forall i$ $\gamma_i(\varphi) = \gamma'_i(\varphi')$, then $Ag_\Gamma(\varphi) = Ag_{\Gamma'}(\varphi')$.

Clearly, **Systematicity** is equivalent to **Independence** and **Neutrality**.

The above properties are quite standard ones, but unfortunately they are jointly incompatible:

Proposition 7.4 (List and Pettit, 2002). *There exists no JA rule that satisfy* **C-universal domain**, **Collective rationality**, **Collective resoluteness**, **Systematicity**, *and* **Anonymity**.

This impossibility theorem relies on some strong assumptions. First is the completeness assumptions of the individuals (**C-universal domain**), that can be criticized, since in many cases one cannot reasonably expect all agents to have an opinion on all possible issues; this is also the case of the **Collective completeness** property, that is helpful for making decisions, but forces to make some choices even when it is not possible to do so (Gärdenfors, 2006). Thus the **Collective completeness** requirement imposes sometimes to discriminate further some judgment sets, using additional information not given in the input profile,

Γ	φ_1	φ_2	φ_3	φ_4	φ_5	φ_6
1	1	1	0	1	1	1
2	1	1	1	0	1	1
3	1	1	1	1	0	1
4	1	1	1	1	1	0
5	1	0	1	1	1	1
6	1	0	1	1	1	1

Γ'	φ_1	φ_2	φ_3	φ_4	φ_5	φ_6
$1'$	0	1	1	1	1	1
$2'$	0	1	1	1	1	1
$3'$	0	1	1	1	1	1
$4'$	0	1	1	1	1	1
$5'$	1	0	0	0	0	0
$6'$	1	0	0	0	0	0

Table 7.1: Example of the need of a non-isolated decision on φ_2

and as such, it conflicts with the **Anonymity** and **Neutrality** conditions. Suppose for instance a perfect tie (say, about a unique issue φ in the agenda, with 4 votes for and 4 votes against it), why and how to make a distinction between φ and $\neg\varphi$? The **Systematicity** property is also highly criticizable, as shown by Everaere et al. (2014b). Indeed, it prevents from viewing JA as an optimization process, trying to achieve a best compromise. The following example illustrates this:

Example 7.2. *Let us consider an agenda X composed of the following six formulae: $\varphi_1 = (\neg a \vee \neg b \vee \neg c \vee \neg d \vee \neg e)$, $\varphi_2 = a$, $\varphi_3 = b$, $\varphi_4 = c$, $\varphi_5 = d$, $\varphi_6 = e$. Let us consider the profiles Γ and Γ' on this agenda, as given by Table 7.1. In the (resolute) profile Γ, every formula has a majority of votes, so using simple majority vote all the formulae have to be selected, which would lead to an inconsistent collective judgment set. So (at least) one of the six formulae has to be rejected by the judgment aggregation correspondence. There is a unanimity for accepting φ_1, so it seems sensible to select φ_1 in the result. All the other formulae except φ_2 are quasi-unanimous (they get all votes but one). The least supported formula is φ_2, so we can consider that the most sensible result should be $\gamma_\Gamma = \{\varphi_1, \neg\varphi_2, \varphi_3, \varphi_4, \varphi_5, \varphi_6\}$.*

Consider now the profile Γ'. The simple majority vote leads to a consistent collective judgment set $\gamma_{\Gamma'} = \{\neg\varphi_1, \varphi_2, \varphi_3, \varphi_4, \varphi_5, \varphi_6\}$, which thus appears as the expected result (a requirement of the majority preservation property that we will recall below). So, although the individual judgments for φ_2 are the same ones in the two profiles Γ and Γ', $\neg\varphi_2$ is selected when Γ is considered, whereas φ_2 is selected when Γ' is considered. Since φ_2 gets the same votes pros and cons in the two profiles, no judgment aggregation method satisfying **Systematicity** *can make such a distinction.*

This example illustrates that the individual judgments on an issue should not be considered independently from those for the other issues.

Lang et al. (2011) and Everaere et al. (2014b) study other attractive properties for JA operators, such as **Unanimity** or **Majority preservation**.

Unanimity. For any $\varphi_k \in X$, for any profile Γ in the domain of Ag, if $\exists x \in \{0,1\}$ s.t. $\forall \gamma_i \in \Gamma$, $\gamma_i(\varphi_k) = x$, then for every $\gamma_\Gamma \in Ag_\Gamma$, we have $\gamma_\Gamma(\varphi_k) = x$.

Note that a unanimity condition is not required when $x = \star$, since in this case it makes sense to let the acceptance of φ_k depend on the acceptance of other (logically related) formulae.

Majority preservation. If the judgment set obtained using the majority rule[3] is consistent and complete, then Ag_{Γ} is a singleton which consists of this set.

Majority preservation (Lang et al., 2011) (called strong majority preservation by Slavkovik (2012)) is a very natural property, stating that if the simple majority vote on each issue leads to a consistent judgment set, then the JA operator must output precisely this set. Indeed, it is sensible to stick to the result furnished by a simple majority vote when no doctrinal paradox occurs.

Many JA operators that have been defined in the literature, but most of them, following the requirements given by the **Systematicity/Independence** properties, do not allow to consider the collective judgment sets as the results of some optimization process, aiming at making decision on each issue by taking the context (i.e., the other issues) into account. Nonetheless, Lang et al. (2011) introduce several families of JA operators based on minimization, inspired by operators considered in voting theory or in artificial intelligence. Majority preservation is presented as a natural requirement for such operators. Everaere et al. (2014b) define another family of operators, called ranked majority operators, that exploit the number of votes received by each formula of the agenda.

7.5 BM vs. JA

A fundamental difference between BM and JA concerns the nature of the input and the nature of the output considered in the two settings. Propositional BM mainly considers the beliefs of individual agents from a group. Beliefs are typically encoded by belief bases, i.e., sets of propositional formulae over a finite set of atoms, or, equivalently, by sets of interpretations. Interpretations are independent and mutually conflicting views of the same world, and each agent believes that the true world is one of the interpretations in her beliefs. Belief bases are often supposed to be consistent. In many approaches to BM, the notion of interpretation is a key notion, since the interpretations are the "candidates" of the decision process (in a nutshell, belief merging can be defined as a process which aims at finding the most plausible interpretations given the beliefs of the agents of the group). The selection of the most plausible interpretations relies on a number of principles. The notion of majority can play a role here (i.e., one can focus on interpretations shared by a majority of agents) but this is not mandatory, some BM operators can select interpretations that are rejected by every agent of the group. The belief base provided by each agent is implicitly assumed to contain all the pieces of belief to be considered in the merging process. Modifying a single base by adding/removing some pieces of belief may have a strong impact on the result furnished by the merging operator, which can be logically strengthened, weakened or become logically unrelated with the merged base obtained before the modification.

In JA, a focus is laid on a specific set of issues, encoded as propositional formulae φ, and called the agenda. The input is a set of individual judgments

[3] Several definitions are possible for the majority rule when abstention is allowed. Here, one considers that the majority rule gives 1 (resp. 0) when the number of agents reporting 1 (resp. 0) is strictly greater than the number of agents reporting 0 (resp. 1), and it gives $\star$ otherwise.

(0/reject or 1/accept or $\star$/abstain) from the agents of the group on the formulae of the agenda. If an agent accepts/rejects φ, then she rejects/accepts $\neg\varphi$. No assumption is made on the way individual decisions/judgments are made by each agent (but a consistency condition, stating that an agent cannot accept formulae that are jointly inconsistent, reflecting the fact that she has consistent beliefs). Each individual judgment reflects the epistemic status of each formula from the agenda in the belief base of the corresponding agent (but nothing more). The beliefs of the group are determined from the collective judgment(s) on the agenda, which is the result produced by the JA operator that is considered. The principle of majority plays an important role here. Typically, the judgment of the group on a given formula φ of the agenda is the point of view of the majority on it: if a majority of agents accept/reject it, then this is the case as well for the group. A doctrinal paradox occurs when the conjunction of the formulae accepted by the group is inconsistent. In such a case, the resulting collective judgment must be somehow weakened for recovering consistency, and several approaches can be used to this purpose (e.g., considering that some formulae of the agenda are more important than others). Modifying the agenda by adding/removing some issues in it may have a strong impact on the collective judgment(s) furnished by the JA operator. Typically, no integrity constraints are considered in a JA process, but it could be possible to do it (merely, by replacing the notion of consistency by a notion of consistency with the integrity constraints).

Unsurprisingly, since BM operators and JA operators are based on different inputs and outputs, the sets of properties of interest for BM operators and for JA operators do not coincide. In propositional BM, some sets of postulates characterize the behavior of rational operators, and representation theorems exist. In judgment aggregation, some properties (often inspired by voting theory) have been identified as well. However, these properties typically lead to impossibility theorems, showing that they are not jointly compatible. That said, some judgment aggregation operators have nevertheless been characterized by sets of properties, as the quota operators in (Dietrich and List, 2007) (the work of Grandi and Endriss (2011) gives also a characterization of quota rules using binary aggregation and suited integrity constraints).

7.6 Decision Policies

As explained previously, BM and JA consider different inputs. Notwithstanding the integrity constraints, in BM, the input is a profile of belief bases, representing the beliefs of a group of agents. In JA, the input is composed of answers "yes" (1), "no" (0) or "undetermined" ($\star$) reported by the agents for some issues (those of the agenda), and the input profile is a vector of such answers (alias judgment sets). Of course, agents might use their beliefs to answer the questions, but it is out of the scope of JA operators to specify how individual judgments are obtained.

Let us consider the following question: if the beliefs K_i of an agent i are known, given an issue φ_k, what could be the opinion of the agent on the issue?

Suppose that the agent only believes that $a \wedge b$ is true, and is questioned about a: she will probably answer "yes" to the question because she necessarily believes that a is true. If the question is $\neg b$, she will probably answer "no" because b being false is incompatible with her beliefs. Suppose now that the agent just believes that a is true, and that the question is $a \wedge b$. In this case the agent probably has no opinion on the question (the question is contingent given her beliefs), thus she will probably answer "undetermined".

What is needed to make it formal is a mapping which characterizes the answers (i.e., the judgment set) an agent can give to the issues of the agenda, depending on her belief base. We call such mapping *decision policies*, and our purpose is first to characterize them axiomatically:

Definition 7.6. *A* decision policy $p : \mathcal{L} \times \mathcal{L} \rightarrow \{0, 1, \star\}$ *is a mapping associating an element of* $\{0, 1, \star\}$ *with any pair of non-trivial formulae* (K, φ) *and satisfying:*

1. *if* $K_1 \equiv K_2$*, then* $\forall \varphi$*,* $p(K_1, \varphi) = p(K_2, \varphi)$

2. *if* $\varphi_1 \equiv \varphi_2$*, then* $\forall K$*,* $p(K, \varphi_1) = p(K, \varphi_2)$

3. $p(\varphi, \varphi) = 1$

Conditions 1 and 2 can be viewed as a formal counterpart, respectively, of a neutrality condition and of an anonymity condition for decision policies.

Given an agenda $X = \{\varphi_1, \ldots, \varphi_m\}$ and a belief base K (respectively a profile $E = (K_1, \ldots, K_n)$ of belief bases), every decision policy p induces a judgment set $p_X(K) = (p(K, \varphi_1), \ldots, p(K, \varphi_m))$ (resp. a profile of judgment sets $p_X(E) = (p_X(K_1), \ldots, p_X(K_n))$.

Examples of decision policies are the following ones:

$$p_B(K, \varphi) = \begin{cases} 1 & \text{if } K \models \varphi \\ 0 & \text{if } K \models \neg\varphi \\ \star & \text{otherwise} \end{cases} \qquad p_C(K, \varphi) = \begin{cases} 1 & \text{if } K \wedge \varphi \not\models \bot \\ 0 & \text{otherwise} \end{cases}$$

Using the *belief decision policy* p_B, an agent answers "yes" (resp. "no") to a given issue precisely when it (resp. its negation) is a logical consequence of her belief base; in the remaining case, she answers "undetermined".

Observe that with the *consistency decision policy* p_C it is possible to have together $p_C(K_i, \varphi_k) = 1$ and $p_C(K_i, \neg\varphi_k) = 1$ (for instance, a belief base equivalent to a is consistent with b and with $\neg b$). In order to avoid this problem, some additional conditions must be satisfied:

Definition 7.7. *Let* $p : \mathcal{L} \times \mathcal{L} \rightarrow \{0, 1, \star\}$ *be a decision policy. It is a* rational *decision policy if it satisfies the two following conditions:*

4. *if* $p(K, \varphi) = 1$*, then* $p(K, \neg\varphi) = 0$

5. *If* $K_1 \wedge K_2$ *is consistent and if* $p(K_1, \varphi) = 1$ *then* $p(K_1 \wedge K_2, \varphi) = 1$

It turns out that these two additional conditions fully characterize the belief decision policy:

Proposition 7.5. p *is a rational decision policy if and only if* $p = p_B$*.*

p_B also ensures individual consistency:

Proposition 7.6. *Whatever the belief base* K *and the agenda* X*, if* γ *is the judgment set on* X *induced by* p_B *given* K*, then the associated judgment* $\widehat{\gamma}$ *is consistent.*

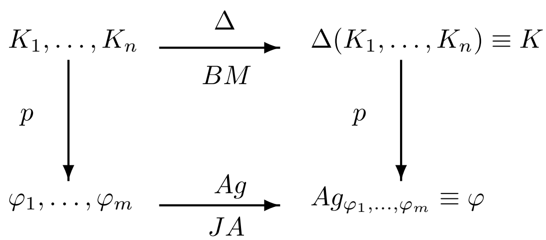

Figure 7.2: BM vs. JA

7.7 Merge-then-Project or Project-then-Aggregate?

Thanks to the rational decision policy $p = p_B$ defined in the previous section, we know how to determine the judgment set of an agent i on any agenda $X = \{\varphi_1, \ldots, \varphi_n\}$ when her belief base K_i is known. This judgment set $(p(K_i, \varphi_1), \ldots, p(K_i, \varphi_m))$ can be viewed as the "projection" of K_i onto X, and such a concept of projection was the missing link between the BM process and the JA one, as illustrated on Figure 7.2.

From it, it appears that there are two ways to define the collective judgment(s) of a group of agents on an agenda, assuming that the beliefs of the agents are known. On the one hand, belief bases can be first merged using a BM operator, then the resulting merged base can be projected onto the agenda. On the other hand, the belief bases can be first projected on the agenda, and then the resulting $0/1/\star$ matrix can be aggregated using a JA operator. The chosen approach typically has a strong influence on the output of the process.

Example 7.3. *Let us illustrate this on a simple example, with 4 agents. Suppose that the set $\mathcal{P}$ of propositional symbols is $\{a, b\}$, so that only 4 interpretations are possible: $\omega_1 = 00$, $\omega_2 = 01$, $\omega_3 = 10$, and $\omega_4 = 11$. Suppose also that the set of models of the belief base of each agent consists of a single interpretation: $[K_1] = \{\omega_1\}$, $[K_2] = \{\omega_2\}$, $[K_3] = \{\omega_3\}$, and $[K_4] = \{\omega_4\}$. Suppose finally that the agenda X consists of the three formulae $\neg a \wedge \neg b$, $\neg a \wedge b$, and $a \wedge \neg b$.*

In order to use the project-then-aggregate approach, one first need to determine the individual judgments on X. They are reported in the following table:

	$\neg a \wedge \neg b$	$\neg a \wedge b$	$a \wedge \neg b$
K_1	1	0	0
K_2	0	1	0
K_3	0	0	1
K_4	0	0	0

Since a majority of agents reject each of the issues, using a JA rule satisfying the **Majority preservation** *condition leads to the collective judgment $\gamma_2 = (0, 0, 0)$ so that $\widehat{\gamma_2} \equiv a \wedge b$ which is consistent. Hence the beliefs of the group will be computed as $a \wedge b$, despite the fact that each agent actually rejects $a \wedge b$.*

This example also shows the importance of the choice of the agenda. For this example, the fact that $a \wedge b$ is accepted or rejected by any agent is fully determined by her individual judgment set on $X = \{\neg a \wedge \neg b, \neg a \wedge b, a \wedge \neg b\}$ (since judgment sets *are supposed to be consistent). Stated otherwise, $a \wedge b$ is a redundant issue given X, hence asking each agent whether she accepts or rejects $a \wedge b$ should be useless when the issues from X have been considered. It turns out that this is not the case. Let us complete the agenda X with the issue $a \wedge b$ and the previous table by the individual judgments of the four agents on $a \wedge b$:*

	$\neg a \wedge \neg b$	$\neg a \wedge b$	$a \wedge \neg b$	$a \wedge b$
K_1	1	0	0	0
K_2	0	1	0	0
K_3	0	0	1	0
K_4	0	0	0	1

Since a majority of agents reject each of the issues, using a JA rule satisfying the **Majority preservation** *condition leads to the collective judgment $\gamma_3 = (0, 0, 0, 0)$ so that $\widehat{\gamma_3} \equiv \bot$ which is inconsistent. This conflict must thus be solved by the JA rule under consideration. Clearly enough, since the instance is fully symmetric, there is no reason here that the beliefs of the group will be computed as $a \wedge b$ again.*

Contrastingly, using the merge-then-project approach (with the distance-based operator $\Delta^{d_H,\Sigma}$), one first compute $\Delta^{d_H,\Sigma}_{\top}((K_1, K_2, K_3, K_4))$ which is equivalent to $\top$, then the projection of this merged base on X gives the collective judgment $\gamma_1 = (\star, \star, \star)$, so that $\widehat{\gamma_1} \equiv \top$ as well, meaning that the group has no information about the issues from X. Note that the same result is obtained if the agenda is completed by the issue $a \wedge b$.

This example illustrates that the processus of judgment aggregation is very sensible to the choice of the issues, and that adding to the agenda an issue (even a redundant one) may have a huge impact on the result. This is strongly related to the problem of manipulating the agenda and as such, it has been studied in the JA literature, for example by Dietrich (2016) and by Lang et al. (2016).

Let us now give an example that illustrates that some information can be lost with the project-then-aggregate approach while it is preserved when the merge-then-project approach is considered instead.

Example 7.4. *Let $K_1 = \{p \vee q\}$, $K_2 = K_1$, and $K_3 = \{p \vee \neg q\}$ and an agenda containing only p. The matrix containing the responses to p is composed of three $\star$, because no agent knows whether p is true or false. If the projection of the bases on p is computed first, then the three projected bases are equivalent to $\top$, and the result of their aggregation will be equivalent to $\top$ as well (the group does not know whether p is true or false). If the three bases are first merged using an IC merging operator, then since their conjunction is consistent, the merged base will be equivalent to p. If this merged base is then projected onto p, the conclusion will be that p is true. Thus, JA typically lead to lose much more information than BM when the input belief bases are not complete ones.*

7.8 More on BM vs. JA

Let us step back to Figure 7.2 and denote by φ_{Ag} the formula obtained following the project-then-aggregate path $Ag \circ p$, and φ_{Δ} the formula obtained following the merge-then-project path $p \circ \Delta$. While the previous example shows that φ_Δ and φ_{Ag} can be distinct, it is important to determine whether some logical connections between φ_Δ and φ_{Ag} can be ensured whenever Δ and Ag both satisfy some rationality properties.

The answer is negative in the general case. Everaere et al. (2015) show that the two resulting formulae are not necessarily jointly consistent. More precisely, for distance-based operators, a negative result has been exhibited: φ_{Ag} and $\varphi_{\Delta^{d,f}}$ are not necessarily jointly consistent even if Ag satisfies **Majority preservation**, d is any normal[4] distance and f is any strictly non-decreasing aggregation function. The significance of the result comes from the fact that usual distances between interpretations (Hamming distance, drastic distance) are normal ones and usual aggregation functions (Σ, $Gmax$, $Gmin$, Σ^n, …) satisfy strict non-decreasingness.

Things are different in the case when the two approaches are equally informed, i.e., when the agenda X gathers all interpretations of $\mathcal{W}$. Under these assumptions, Everaere et al. (2015) have shown that the beliefs of the group computed following the merge-then-project path are the same as the beliefs of the group obtained following the project-then-aggregate path. On this ground, an irresolute JA rule $Ag = Ag^\Delta$ can be defined from a BM operator Δ, and reciprocally, a BM operator $\Delta = \Delta^{Ag}$ can be defined from a irresolute JA rule Ag. Details of the construction are given by Everaere et al. (2015).

A natural question is then to determine whether imposing some rationality conditions on a BM operator Δ leads the induced JA operator $Ag = Ag^\Delta$ to satisfy some rationality conditions, and vice-versa, whether imposing some rationality conditions on a JA operator Ag leads the induced BM operator $\Delta = \Delta^{Ag}$ to satisfy some rationality conditions.

It has been shown that Δ^{Ag} satisfies the postulates **(IC0)**, **(IC1)** and **(IC3)** when Ag satisfies **Universal domain** and **Anonymity**. For getting **(IC2)** for Δ^{Ag}, an additional condition of **Consensuality** must be satisfied by Ag. This condition ensures that if an issue of the agenda is accepted by **all** the agents (i.e., it is a unanimous issue), then the collective judgment set computed by Ag consists exactly of those unanimous issues. Interestingly, **(IC4)** for Δ^{Ag} is not ensured by the **Neutrality** condition on Ag.

Other connections between the satisfaction of some IC postulates for BM and the satisfaction of some conditions on JA operators can be established for the operators in correspondence (Ag^Δ and Δ, and Δ^{Ag} and Ag) in the complete agenda case (i.e., when the agenda is the set of all interpretations). The **Weak consistency** condition on JA operators states that if an issue is accepted by a profile Γ of individual judgment sets and by a profile Γ' of individual judgment sets, then it must be accepted by the union of the two profiles.

The **Consistency** condition strengthens it by stating that if there is at least one issue that is accepted by two profiles Γ and Γ', then each issue that is accepted by the whole profile $\Gamma \sqcup \Gamma'$ should be accepted by each of the two profiles Γ and Γ'. It turns out that those properties correspond respectively to the IC postulates **(IC5)** and **(IC6)**. Quite surprisingly, these conditions have not been considered as standard ones for JA operators (we are only aware of (Lang et al.,

[4]A distance is *normal* if $d(\omega_1, \omega_2) \leqslant d(\omega_3, \omega_4)$ whenever the variables which have different truth values in ω_1 and ω_2 are included into the variables which have different truth values in ω_3 and ω_4.

2011; Slavkovik, 2012; Lang et al., 2016) which point out the **Consistency** condition, under the name "separability").

Everaere et al. (2015) give a list of properties required for a JA operator to induce an IC merging operator in the complete agenda case. A key question is whether these properties can be satisfied by some JA operator. A positive answer to this issue is given when some JA rules $\delta^{RM_\oplus}$ defined by Everaere et al. (2014b) are considered. Roughly, each $\delta^{RM_\oplus}$ rule consists in selecting in the set of all consistent and resolute judgment sets the best ones, where the score of each judgment set is defined as the $\oplus$-aggregation of an m-vector of values (one value per question in the agenda X, reflecting the number of agents supporting the question in the input profile Γ). When the agenda is complete, for any $\oplus$ satisfying strict non-decreasingness, the ranked majority judgment aggregation rule $\delta^{RM_\oplus}$ satisfies **Universal domain**, **Collective rationality**, **Collective resoluteness**, **Anonymity**, **Neutrality**, **Unanimity**, **Consensuality**, and **Majority preservation**. It does not satisfy **Independence**. For $\oplus = \Sigma$, **Weak consistency** and **Consistency** are also satisfied.

7.9 Conclusion

BM and JA are two distinct theories for the aggregation of beliefs. They do not operate on the same inputs, and typically lead to collective beliefs that can be jointly conflicting. When focusing on the case where the inputs are equally informed (i.e., when the agenda is the set of all interpretations), some valuable connections between the two families of operators and between the corresponding rationality postulates can be established nevertheless.

Bibliography

C. Baral, S. Kraus, and J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering*, 3(2):208–220, 1991.

C. Baral, S. Kraus, J. Minker, and V. S. Subrahmanian. Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, 8(1):45–71, 1992.

S. Benferhat, D. Dubois, S. Kaci, and H. Prade. Possibilistic merging and distance-based fusion of propositional information. *Annals of Mathematics and Artificial Intelligence*, 34(1-3):217–252, 2002.

I. Bloch and J. Lang. Towards mathematical morpho-logics. In *Proceedings of the 8th International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU)*, pages 1405–1412, 2000.

J.-F. Condotta, S. Kaci, P. Marquis, and N. Schwind. Merging qualitative constraint networks in a piecewise fashion. In *Proceedings of the 21st International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 605–608, 2009.

S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasquie-Schiex, and P. Marquis. On the merging of dung's argumentation systems. *Artificial Intelligence*, 171:740–753, 2007.

J. Delgrande and T. Schaub. A consistency-based framework for merging knowledge bases. *Journal of Applied Logic*, 5(3):459–477, 2007.

J. Delgrande, D. Dubois, and J. Lang. Iterated revision as prioritized merging. In *Proceedings of the 10th International Conference on Knowledge Representation and Reasoning (KR)*, pages 210–220, 2006.

J. Delgrande, T. Schaub, H. Tompits, and S. Woltran. Merging logic programs under answer set semantics. In *Proceedings of the 25th International Conference on Logic Programming (ICLP)*, pages 160–174, 2009.

J. Delobelle, A. Haret, S. Konieczny, J.-G. Mailly, J. Rossit, and S. Woltran. Merging of abstract argumentation frameworks. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 33–42, 2016.

F. Dietrich. Judgment aggregation and agenda manipulation. *Games and Economic Behavior*, 95:113–136, 2016.

F. Dietrich and C. List. Judgment aggregation by quota rules: Majority voting generalized. *Journal of Theoretical Politics*, 4(19):391–424, 2007.

D. Dubois, J. Lang, and H. Prade. Nonmonotonic reasoning and uncertain reasoning, chapt. possibilistic logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford Science Publications, 1994.

P. Everaere, S. Konieczny, and P. Marquis. The strategy-proofness landscape of merging. *Journal of Artificial Intelligence Research*, 28:49–105, 2007.

P. Everaere, S. Konieczny, and P. Marquis. Conflict-based merging operators. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 348–357, 2008.

P. Everaere, S. Konieczny, and P. Marquis. Disjunctive merging: Quota and Gmin merging operators. *Artificial Intelligence*, 174(12-13):824–849, 2010.

P. Everaere, S. Konieczny, and P. Marquis. On egalitarian belief merging. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 121–130, 2014a.

P. Everaere, S. Konieczny, and P. Marquis. Counting votes for aggregating judgments. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1177–1184, 2014b.

P. Everaere, S. Konieczny, and P. Marquis. Belief merging versus judgment aggregation. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 999–1007, 2015.

P. Gärdenfors. An Arrow-like theorem for voting with logical consequences. *Economics and Philosophy*, 22(2):181–19, 2006.

N. Gorogiannis and A. Hunter. Merging first-order knowledge using dilation operators. In *Proceedings of the 5th International Symposium on Foundations of Information and Knowledge Systems (FoIKS)*, pages 132–150, 2008.

U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 204–209. AAAI Press, 2011.

A. Haret, A. Pfandler, and S. Woltran. Beyond IC postulates: Classification criteria for merging operators. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI)*, pages 372–380, 2016.

J. Hué, O. Papini, and E. Würbel. Merging belief bases represented by logic programs. In *Proceedings of the 10th International Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, pages 371–382, 2009.

S. Kaci, S. Benferhat, D. Dubois, and H. Prade. A principled analysis of merging operations in possibilistic logic. In *Proceedings of the 16th Conference in Uncertainty in Artificial Intelligence (UAI)*, pages 24–31, 2000.

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

S. Konieczny. 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)*, pages 135–144, 2000.

S. Konieczny and R. Pino Pérez. On the frontier between arbitration and majority. In *Proceedings of the 8th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 109–118, 2002a.

S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002b.

S. Konieczny, J. Lang, and P. Marquis. DA^2 merging operators. *Artificial Intelligence*, 157:49–79, 2004.

S. Konieczny, P. Marquis, and N. Schwind. Belief base rationalization for propositional merging. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, pages 951–956, 2011.

J. Lang, G. Pigozzi, M. Slavkovik, and L. van der Torre. Judgment aggregation rules based on minimization. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, pages 238–246, 2011.

J. Lang, M. Slavkovik, and S. Vesic. Agenda separability in judgment aggregation. In *Proceedings of the 30th Conference on Artificial Intelligence (AAAI)*, pages 1016–1022, 2016.

C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18:89–110, 2002.

G. Pigozzi. Belief merging and the discursive dilemma: An argument-based account to paradoxes of judgment aggregation. *Synthese*, 152(2):285–298, 2006.

P. Z. Revesz. On the semantics of arbitration. *International Journal of Algebra and Computation*, 7(2):133–160, 1997.

M. Slavkovik. *Judgment Aggregation in Multiagent Systems*. PhD thesis, Luxemburg University, 2012.

W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. *Causation in Decision, Belief Change, and Statistics*, 2:105–134, 1987.