



Belief merging in absence of reliability information

Paolo Liberatore¹ 

Received: 10 January 2021 / Accepted: 18 May 2022 / Published online: 5 July 2022

© The Author(s) 2022

Abstract

Merging beliefs depends on the relative reliability of their sources. When this information is absent, assuming equal reliability is unwarranted. The solution proposed in this article is that every reliability profile is possible, and only what holds according to all of them is accepted. Alternatively, one source is completely reliable, but which one is not specified. These two cases motivate two existing forms of merging: maxcons-based merging and disjunctive merging.

Keywords Belief merging · Knowledge representation · Nonmonotonic reasoning · Artificial intelligence

1 Introduction

Most of the literature on belief merging concerns sources of the information of equal reliability (Chopra et al. 2006; Everaere et al. 2010a; Konieczny and Pino Pérez 2011; Lin and Mendelzon 1999). Such a scenario occurs, but not especially often. Two identical temperature sensors produce readings that are equally likely to be close to the actual value, but a difference in made, age, or position changes their reliability. Two experts hardly have the very same knowledge, experience and ability. The reliability of two databases on a certain area may depend on factors that are unknown when merging them.

Merging under equal and unequal reliability are two scenarios, but a third exists: absent reliability. Most previous work in belief merging is about the first (Chopra et al. 2006; Everaere et al. 2010a, 2020; Haret et al. 2020; Konieczny and Pino Pérez 2011, 2002a; Lin and Mendelzon 1999); some is about the second (Cholvy 1998; Konieczny et al. 2004; Lin 1996; Revesz 1997); this one is about the third.

The difference between equal and absent reliability is clear when its implications on some examples are shown.

✉ Paolo Liberatore
liberato@diag.uniroma1.it

¹ DIAG, Sapienza University of Rome, Rome, Italy

Example	Equal reliability	Absent reliability
Two experts	They have the very same knowledge, experience and ability	They differ in knowledge, experience and ability, but how much of these they possess is unknown
Two sensors	They are of the same kind and are in the same condition (temperature sensors located next to each other, distance sensors with the same orientation)	They are of different kind, or are in different conditions, and which one is more reliable in the current situation is not known
Two databases	They cover the very same domain, and are equally likely to be correct	They cover different domains, so that a certain piece of information may have been crucial to one but a detail in the second

The assumption of equal reliability is quite strong in the example of the two experts; rather, there may be some reason to believe one more than the other; not knowing who, this scenario falls in the case of absent reliability. For the two sensors and the two databases equal reliability is not unlikely, but so is absent reliability.

If reliability is absent, can it be assumed equal?

When merging preferences, yes. When merging beliefs, no.

Merging preferences (List 2013; Lang 2004; Mata Díaz and Pino Pérez 2017) aims at obtaining a result that best reflects the collective opinion of a group. A common premise is that all members of the group have the same weight on the final decision, as formalized by the condition of anonymity. In lack of information telling otherwise, equal weights are a valid assumption.

A technical example shows why not when merging beliefs instead. Three scenarios are possible: *A*, *B* and *C*; two sources of information rank their unlikeliness on a scale from 0 to 3, with 0 being the most likely and 3 the least (unlikeliness scales are common in belief revision (Darwiche and Pearl 1997; Katsuno and Mendelzon 1991; Rott 2006), in spite of likeliness being more intuitive). The first source grades *A* as the most unlikely scenario, the second as the most likely; numerically, unlikeliness are 3 and 0. Both sources grade *B* as kind of likely (1), and *C* in the opposite way of *A* (0 and 3).

Scenario	Unlikeliness according to the first source	Unlikeliness according to the second source
<i>A</i>	3	0
<i>B</i>	1	1
<i>C</i>	0	3

Two different cases are considered: in the first case, the first source is twice as reliable as the second; in the second, no reliability information is present.

If the first source is twice as reliable as the first, the overall unlikeliness of the three scenarios are $2 \cdot 3 + 0 = 6$, $2 \cdot 1 + 1 = 3$ and $2 \cdot 0 + 3 = 3$. The minimum is 3: the least unlikely scenarios are *B* and *C*.

Scenario	Unlikeliness according to the first source, weighted	Unlikeliness according to the second source, weighed
<i>A</i>	$2 \cdot 3$	0
<i>B</i>	$2 \cdot 1$	1
<i>C</i>	$2 \cdot 0$	3

If some fact *x* holds in *B* but not in *C*, its truth is uncertain since it differs in the two most likely scenarios.

The second considered case is when reliability information is absent. The table of overall unlikeliness can no longer be computed since it requires not only the unlikeliness of the scenarios according to the sources but also the reliability of the sources. A tempting solution is: “since the relative reliability of the sources is absent, it is assumed equal”. This allows to compute the table again, this time with multipliers 1 and 1.

Scenario	Unlikeliness according to the first source, weighted	Unlikeliness according to the second source, weighed
<i>A</i>	3	0
<i>B</i>	1	1
<i>C</i>	0	3

The most likely scenario is now *B*, and *B* only: its overall unlikeliness 2 beats those of *A* and of *C*, both 3. If the fact *x* holds in *B* but not in *C*, it is deemed true.

The two cases differ both in their initial information and in their conclusions. In the first case, reliability information is present, and *x* is not concluded. In the second case, reliability information is absent, but *x* is concluded. Starting from more information leads to less information.

Knowing that *x* is true is more information than not knowing the value of *x*. That the first source is twice as reliable as the second is more information than an unspecified relative reliability.

Information is not just different. It is strictly more in one case than in the other: knowing that the first source is twice more reliable than the second is strictly more

information than no given reliability; knowing that x is true is strictly more information than not knowing it. The case starting from strictly more information ends up with strictly less information.

This example is inspired by the “penny z ” of Popper (1959, pp. 425–426): a coin is initially assumed fair in lack of information indicating otherwise; adding the confirmation of fairness does not change its probability of falling heads or tails. In the interpretation of probability as degree of belief (Hájek 2012), the probabilities are the epistemic state. Adding information should alter the epistemic state, but the addition of fairness (which is new information) changes nothing.

This example was used by Popper against the subjective interpretation of probabilities, but relies on the principle of indifference: events of unknown probability are assumed equally probable (Keynes 1921; Shackel 2007). The Bertrand paradox (Bertrand 1889; Keynes 1921; Shackel 2007) shows it is problematic; the coin example by Popper (1959) shows another contradictory aspect of it (Gärdenfors and Sahlin 1982).

The belief merging version of the principle of indifference is the assumption of equal reliability in lack of information about the relative reliability of the sources. In the subjective interpretation of probability, the probability of an event is the degree of belief in that event happening (Hájek 2012); in belief merging, the weight of a source is the likeliness of the formulae it provides being true, or at least close to truth (Cholvy 1998; Darwiche and Marquis 2004; Konieczny et al. 2004; Lin 1996; Revesz 1997). The event “formula F is true in the real world” provides a qualitative connection of probability with merging. The principle of indifference translates into the assumption of equal reliability.

The probability version of sources of unknown reliability is the lack of knowledge of the probability of events. Economists distinguish between risk (known probability) and Knightian uncertainty (unknown probability) (Nishimura and Ozaki 2007). An often-used example is the urn containing twenty yellow balls and forty balls of another color, which may be either blue or green; these forty balls are either all blue or all green, but which of the two is not known. This scenario involves both risk (the probability of yellow or not yellow is known) and Knightian uncertainty (the presence of blue balls is unknown).

This urn suggests a way to deal with the problem in belief merging. The probability of drawing a yellow ball is always one third, but assuming the same for blue and green is as if the urn contained twenty balls for each color. This is acceptable for a single drawn, but not in general. The probability of drawing two balls of the same color (putting the first ball back in the urn) under the assumption of equal probability is $\frac{1}{3}$ instead of $\frac{1}{3} \cdot \frac{1}{3} + \frac{2}{3} \cdot \frac{2}{3}$. The first value is obtained by selecting from the nine possible outcomes of probability $\frac{1}{9}$ each (random first ball and random second ball) only the three where the balls have the same color: $3 \cdot \frac{1}{9} = \frac{1}{3}$. The second value can be obtained by considering the second drawn not independent of the first, but also by calculating the probability under the assumption of forty blue balls: the probability of the two balls being both yellow is $\frac{1}{3} \cdot \frac{1}{3}$, that of being both blue is $\frac{2}{3} \cdot \frac{2}{3}$. Importantly, the very same value is obtained for forty green balls instead. Not only this probability holds in both cases, it resists the addition of information. It holds even if it is later discovered

that the urn is made in a factory that normally produces forty green balls, and the blue ball version is a rare collector's edition.

In terms of belief merging, two sources of unspecified reliability may be equally reliable, or one may be more reliable than the other. All cases are considered, and only what holds in all of them is taken. This is analogous to reasoning from multiple probability distributions (Halpern and Tuttle 1993).

How does this solution work in the example of the three scenarios A , B and C respectively ranked $[3, 0]$, $[1, 1]$ and $[0, 3]$? Scenario A is preferred if the second source is much more reliable than the first, scenario B if they are equally reliable and scenario C if the first is much more reliable than the second. If reliability information is absent, none of the three scenarios can be considered more likely than the others. If the first source is much more reliable than the second, scenario A is much less likely than the others. More information (the relative reliability of the sources) leads to more information (from all three scenarios to only B and C).

A result in this article is that the disjunction of all maxcons (Ammoura et al. 2015; Baral et al. 1992; Benferhat et al. 1997; Brewka 1989; Dubois et al. 2016; Grant and Hunter 2011; Konieczny and Pino Pérez 2011) is the result of merging formulae of unknown reliability using the drastic distance. This result invalidates the view that maxcons are unsuitable for merging since they do not take into account the distribution of information (Konieczny 2000; Konieczny and Pino Pérez 2011). This may be the case under equal or otherwise specified reliability, but maxcons do exactly what they should when reliability information is absent.

Another result is a motivation for disjunctive merging, the kind that only selects models of the formulae (Everaere et al. 2010a; Liberatore and Schaerf 1998). It results from assuming that one of the sources is completely reliable, but which one is not specified.

Technically, merging is defined by selecting the models at a minimal weighted distance from the formulae provided by the sources. The drastic and the Hamming distances are considered as two relevant examples. After Sect. 2 fixes the formal language used and other notions related to merging, Sect. 3 defines merging when reliability is completely or partially absent; this definition is based on the concept of weighted distance between models and formulae. Sect. 4 shows results about a property of models that makes them relevant to merging. Sect. 5 and Sect. 6 analyze the case of the drastic and the Hamming distance. Sect. 7 shows which postulates are satisfied, while Sect. 8 concentrates on a specific condition of merging. Sect. 9 considers alternative ways of merging: sum of powers, leximax and leximin. Sect. 10 briefly considers the case of sources providing more than one formula. Sect. 11 discusses the results obtained in this article. An appendix contains all proofs of theorems and lemmas.

2 Preliminaries

The formulae in this article are propositional over a finite alphabet. Models are represented by the set of literals they satisfy; for example, $I = \{a, \neg b, c\}$ is the model assigning false to b and true to a and c . The notation $I \models F$ indicates that the model

I satisfies the formula F . The same symbol is also used to indicate that all models of a set satisfy the formula; for example, $\{I, J\} \models F$ tells that both I and J satisfy F . The set of all models that satisfy a formula F is denoted $\text{mod}(F)$.

The formulae to merge are denoted F_1, \dots, F_m ; a single other formula μ contains all integrity constraints—what is known for certainty. Contrary to most previous studies in belief merging, the formulae to be merged are not assumed equally reliable, nor they are assumed to have a certain relative reliability either. The aim of merging is to draw as many conclusions as possible.

Merging is formalized by a function from μ and F_1, \dots, F_m to something that represents information. In this article, this function produces the propositional interpretations that model the scenarios that are considered possible as the result of merging. In other words, the codomain of this function is the set of all sets of propositional models over the given alphabet. As an example, merging $x \wedge z$ and $y \wedge \neg z$ under the integrity constraint $x \wedge y$ may result in two scenarios considered possible: one where x, y and z are all true and another where x and y are true while z is false; the result of the function is the set of the two models $\{x, y, z\}$ and $\{x, y, \neg z\}$.

Definition 1 A merging operator is a function Δ from propositional formulae μ and F_1, \dots, F_m to sets of propositional models that satisfy μ .

The most general situation is when the reliability of the formulae is absent. Other cases are: the formulae are equally reliable; one is much more reliable than the others; none is so (this case was suggested by a referee). These three cases are the implicit assumptions of respectively the usual definition of merging, of disjunctive merging and of merging by majority. Unless otherwise specified, merging with absent reliability means that no reliability information is present at all, not even qualitatively like in the case of no formula being much more reliable than the others.

Merging is often based on a distance measure between models. This is a function from pairs of models to non-negative integers. If I and J are two models, $d(I, J)$ is a non-negative integer that tells how much they differ. This integer is zero if I and J coincide, otherwise it is greater than zero.

Two intuitive and commonly used distances are the drastic and the Hamming distance. The drastic distance is defined by $dd(I, J) = 0$ if $I = J$ and $dd(I, J) = 1$ otherwise. The Hamming distance $dh(I, J)$ is the number of literals assigned different truth values by I and J ; for example, $dh(\{a, \neg b, c\}, \{\neg a, \neg b, \neg c\}) = 2$, since the two models differ on a and c . Other distances can be defined; they are assumed to satisfy $d(I, I) = 0$ and $d(I, J) > 0$ if $I \neq J$.

3 Merge by weights

In this article, merging is done by minimizing the weighted distance of the models obeying the integrity constraints from the formulae to be merged. The integrity constraints are denoted μ , the formulae to be merged F_1, \dots, F_m . This is the basic settings for belief merging, where each source provides exactly one formula F_i ; the case of multiple formulae is considered in a following section. Formulae μ and F_1, \dots, F_m are propositional over a finite alphabet.

Merging is based on the distance between models, denoted by $d(I, J)$. Two intuitive and commonly used distances are the drastic and the Hamming distance, defined in the previous section. Distance extends from models to formulae: regardless of which distance is used, $d(I, F)$ is the minimal value of $d(I, J)$ for every $J \models F$. It further extends from a formula to a list of them: the distance between a model and a list of formulae is the array of integers $d(I, F_1, \dots, F_m) = [d(I, F_1), \dots, d(I, F_m)]$.

Merging by weighted distance was the historically first way of integrating formulae coming from sources of different reliability (Revesz 1997). Given a vector of positive integers $W = [w_1, \dots, w_m]$, the weighted distance of a model I from the formulae F_1, \dots, F_m is $W \cdot d(I, F_1, \dots, F_m)$, where the dot stands as usual for the scalar product:

$$[w_1, \dots, w_m] \cdot d(I, F_1, \dots, F_m) = \sum_{1 \leq i \leq m} w_i d(I, F_i)$$

This product defines a single integer telling the aggregated distance from I to the formulae F_i , weighted by the relative reliability of each as represented by the integer w_i . Merging selects the models satisfying the integrity constraints μ that have minimal weighted distance from the formulae.

$$\Delta_{\mu}^{d,W}(F_1, \dots, F_m) = \{I \models \mu \mid W \cdot d(I, F_1, \dots, F_m) \text{ is minimal}\}$$

This function depends on two parameters: a model-to-model distance d and a vector of weights $W = [w_1, \dots, w_m]$.

Fixed weights are used when the relative reliability of the sources is given. Weights $W = [1, \dots, 1]$ make the scalar product the same as a sum, and weighted merge the same as the usual operators based on the sum of the drastic and Hamming distances. In the notation by Konieczny and Pino Pérez (2011):

$$\begin{aligned} \Delta^{dd,[1,\dots,1]} &= \Delta^{d_D,\Sigma} \\ \Delta^{dh,[1,\dots,1]} &= \Delta^{d_h,\Sigma} \end{aligned}$$

The dh distance was first used in belief revision by Dalal (1988); for this reason, it is sometimes called ‘‘Dalal distance’’. Revesz (1993, 1997) used it with weights for belief merging, followed by Lin (1996) and Lin and Mendelzon (1999). Weights reflect the reliability of the sources: the distance from a formula of high weight affects the total more than the distance from a formula of low weight.

When reliability is absent, all possible weight vectors are considered. The set of all weight vectors of positive integers is the focus of this article:

$$W_{\exists} = \{[w_1, \dots, w_m] \mid w_i \in \mathbf{N}, w_i > 0\}$$

Nevertheless, other sets of weight vectors are considered. In some scenarios a source is correct; the others only provide refining information. For example, a cardiologist, a pulmonologist and an allergist may have contrasting opinions about the state of a

patient; if the symptoms are caused by a heart disease then the cardiologist is likely to be right on everything, for example on the reason of the breathing problems, even if that contradicts the pulmonologist and the allergist; their opinion only provide some additional insight. In the same way, if the symptoms are caused by an allergy, the allergist is likely right on everything. The same for the pulmonologist. In these cases, one source is totally correct, but which one is unknown.

$$W_a = \{[a, 1, \dots, 1], [1, a, 1, \dots, 1], \dots, [1, \dots, 1, a]\}$$

The value of a for scenarios like that of the three doctors depends on the maximal possible distance between a model and a formula. For the drastic distance, $a = m + 1$ suffices, where m is the number of formulae to be merged. For the Hamming distance, $a = nm + 1$, where n is the number of variables. The opposite case is that of no source deemed much more reliable than the others. It can be formalized by bounding all weights by a constant.

Finally, merging with fixed weights falls into this generalization as the set comprising a single vector. For example, equal reliability is captured by:

$$W_{=} = \{[1, \dots, 1]\}$$

In all these cases, a set of weights W represents all possible reliability the sources are considered to have. Three relevant such sets are W_{\exists} , W_a or $W_{=}$. The set $W_{=}$ is for equally reliable sources; W_{\exists} is the other extreme: no reliability information on the sources is present. Every $W \in W$ is an encoding of the reliability of the sources. All of these are plausible alternatives. Every scenario (every model) that is possible when merging with some $W \in W$ is possible when merging with $W_{=}$:

$$\Delta_{\mu}^{d,W_{=}}(F_1, \dots, F_m) = \bigcup_{W \in W} \Delta_{\mu}^{d,W}(F_1, \dots, F_m)$$

Merging on a set of weights generates all models obtained by merging with one of these weights. This is different from obtaining a single ordering on the models as done by Benferhat et al. (2014) to solve the related problem of commensurability that occurs when the sources themselves assess the reliability of the formulae they provide.

4 Dominance

If a model is farther from every formula than another, the latter is always preferred to the former regardless of the weights. The second model *dominates* the first. Despite the seeming triviality of the concept, a number of relevant results follow:

- if a model has minimal weighted distance for some weights, it is not strictly dominated by another;
- if the codomain is binary, the converse also holds: a model that is not strictly dominated by another has minimal weighted distance for some weights;
- for a ternary codomain, there exist two formulae such that their merge does not include an undominated model;

Dominance could be defined over models with respect to formulae, but is simpler to formalize over vectors of integers. It can then be carried over to the distance vectors of two models.

Definition 2 A vector of integers D dominates another D' , denoted $D \leq D'$, if every element of D is less than or equal to the element of the same index in D' . Strict dominance is the strict part of this ordering: $D < D'$ holds if both $D \leq D'$ and $D' \not\leq D$ hold.

The dominance between the distance vectors of two models is the same as the weak Pareto dominance used in multi-objective decision making (Giagkiozis and Fleming 2014) when the objectives to minimize are the distances between the models and the formulae.

If the distance vector of a model is strictly dominated by that of another, the first is never minimal regardless of the weights. This fact holds because weights are strictly positive.

Lemma 1 For every distance d , vector of weights $W \in W_{\exists}$ and model I of μ , if $I \in \Delta_{\mu}^{d,W}(F_1, \dots, F_m)$ then $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ holds for no model J of μ .

This result is almost trivial, since merging selects the models that have a minimal value of the sum of the distances, each multiplied by a positive weight. The converse does not hold in general, but does in a relevant case: when $d(I, F_i)$ can only be 0 or 1, or more generally when the codomain of d has size two.

Lemma 2 If the codomain of the distance function d is a subset of cardinality two of \mathbf{N} , I is a model of μ , $d(I, F_1, \dots, F_m)$ is not strictly dominated by the vector of distances of any other model of μ , then there exists W such that $I \in \Delta_{\mu}^{d,W}(F_1, \dots, F_m)$.

The last two lemmas imply that the minimal models with arbitrary weights W_{\exists} according to a distance of binary codomain are exactly the models whose distance vectors are not dominated by others. Since dominance is the same as weak Pareto dominance, these minimal models are exactly the Pareto set (Giagkiozis and Fleming 2014). This is therefore the result of merging, but only when reliability information is completely absent and the codomain of the distance function is binary.

Theorem 1 If the codomain of the distance d is a subset of cardinality two of \mathbf{N} , then $\Delta_{\mu}^{d,W_{\exists}}(F_1, \dots, F_m)$ is the set of all models of μ of minimal distance vector according to the dominance ordering.

The next question is whether this condition holds for every fixed-size codomain, or whether a codomain of size three is sufficient for making some undominated model to be excluded from merge. The latter is indeed the case in general. A preliminary lemma will be useful in the sequel.

Lemma 3 If μ has three models of distance $[3, 0]$, $[2, 2]$ and $[0, 3]$ from F_1 and F_2 , then $\Delta_{\mu}^{d,W_{\exists}}(F_1, F_2)$ does not contain the model at distance $[2, 2]$.

This is almost the proof of the claim, but the codomain of the Hamming distance has size unbounded, not three. However, a slightly different distance function fixes the problem.

Theorem 2 *For some distance d with codomain of size three, there exists I, μ and F_1, \dots, F_m such that $I \models \mu$ and $I \notin \Delta_\mu^{d, W\exists}(F_1, \dots, F_m)$, but $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ does not hold for any $J \models \mu$.*

5 Drastic distance

Merging with all possible weights and the drastic distance dd generates all models of all maximal subsets of F_1, \dots, F_m that are consistent with μ . This is proved in three steps:

- dominance with the drastic distance is the same as the containment of the set of formulae F_1, \dots, F_m satisfied by the models;
- the models of the maxcons are the models that are minimal according to that containment;
- therefore, the models of the maxcons are exactly the undominated models; by the results in the previous section, they are the models of minimal weighted distance according to some weights.

Maximal consistent subsets (maxcons) have a general definition over lists of sets of formulae, but what is necessary for this article is only the version with a list of two sets, the first comprising a single consistent formula μ and the second F_1, \dots, F_m . With this limitation, the (possibly non-maximal) consistent subsets and the maximal consistent subsets are defined as:

$$\begin{aligned} \text{con}_\mu(F_1, \dots, F_m) &= \{S \subseteq \{\mu, F_1, \dots, F_m\} \mid \mu \in S \text{ and } S \not\subseteq \perp\} \\ \text{maxcon}_\mu(F_1, \dots, F_m) &= \{S \in \text{con}_\mu(F_1, \dots, F_m) \mid \nexists S' \in \text{con}_\mu(F_1, \dots, F_m) . S \subset S'\} \end{aligned}$$

Since μ is consistent, these sets cannot be empty. To establish the correspondence between models and maxcons, the subset of formulae satisfied by a model is needed.

Definition 3 The set of formulae satisfied by a model I is denoted $\text{subsat}(I, F_1, \dots, F_m) = \{F_i \mid I \models F_i\}$.

The basic brick in the proof construction is that dominance of the drastic distance vectors is the same as containment of the subsets of formulae satisfied by models.

Lemma 4 *For every pair of models I and J and every formulae F_1, \dots, F_m , the following two conditions are equivalent:*

$$\begin{aligned} dd(I, F_1, \dots, F_m) &\leq dd(J, F_1, \dots, F_m) \\ \text{subsat}(J, F_1, \dots, F_m) &\subseteq \text{subsat}(I, F_1, \dots, F_m) \end{aligned}$$

This lemma links the dominance ordering under dd and the containment of subsat . The next links the latter with the maxcons.

Lemma 5 *A model I of μ satisfies some element of $\text{maxcon}_\mu(F_1, \dots, F_m)$ if and only if $\text{subsat}(I, F_1, \dots, F_m) \subset \text{subsat}(J, F_1, \dots, F_m)$ holds for no model J of μ .*

The lemma is the final link of the connection between maxcons and dominance under the drastic distance.

Theorem 3 *For every consistent formulae μ, F_1, \dots, F_m , the following equality holds:*

$$\Delta_\mu^{dd, W\exists}(F_1, \dots, F_m) = \bigcup_{S \in \text{maxcon}_\mu(F_1, \dots, F_m)} \text{mod}(\wedge S)$$

Maxcons have been long used in belief revision (Fagin et al. 1983; Baral et al. 1992; Benferhat et al. 1997; Konieczny and Pino Pérez 2011; Ammoura et al. 2015; Grant and Hunter 2011; Dubois et al. 2016) and nonmonotonic reasoning (Rescher and Manor 1970; Brewka 1989; Ginsberg 1986). Yet, they are sometimes dismissed as “unsuitable for merging” because they do not take into account the distribution of information among the sources (Konieczny 2000; Konieczny and Pino Pérez 2011). This criticism is grounded in the assumption of equal or given reliability. This theorem blocks it from extending to absent reliability; maxcons are weighted merge with the drastic distance when reliability information is completely absent. Not only they are suitable for merging, they deal with the common situation where the credibility of the sources cannot be assessed.

An example clarifies why. If reliability information is completely absent, a formula $\neg x$ provided by two sources cannot beat a formula x provided by one source, since the one source may be much more reliable than the two. Merging by maxcons collects as many formulae as possible while retaining consistency; each maxcon may come from the most reliable sources, making the number of formulae itself irrelevant.

This is not the case when some reliability information is present. A source cannot beat all others when the formulae have comparable reliability, as formalized by weights bounded by a constant. It also cannot when all formulae have the same reliability. This difference supports distinguishing absent reliability from equal reliability.

6 Hamming distance

The Hamming distance dh has a codomain of more than two elements. Therefore, the previous results about binary codomains do not apply. Some existence results are proved:

- every given set of distance vectors is obtainable from some formulae μ, F_1, \dots, F_m ;
- for some μ, F, F' , merging with all possible weight vectors is not equivalent to merging with subexponentially many weight vectors.

Merging with the Hamming distance does not have a simple equivalent form like for the drastic distance, which selects the models that are not strictly dominated by others. The same does not hold in general: a model that is undominated may still be excluded in the merging. This was proved abstractly by three distance vectors $[3, 0]$,

[2, 2] e [0, 3]. With the Hamming distance, these distance vectors can be obtained from concrete formulae. So can every set of distance vectors, actually. This existence result is analogous to a similar one for maxcons (Liberatore 2015, Lemma 4.6).

Lemma 6 *Given an arbitrary set of distance vectors of m elements each, all bounded by an integer n , for some formulae μ and F_1, \dots, F_m over nm variables the vectors of Hamming distances from the models of μ to F_1, \dots, F_m are exactly the given set of distance vectors.*

This theorem allows for an easy way of building counterexamples: rather than providing d and μ, F_1, \dots, F_m that have a certain property, that property is shown directly on the set of distance vectors. This method was already used to prove that some undominated models are not selected by merging, for some distance. In particular, it shows this being the case for the Hamming distance.

Another application is the proof that exponentially many weight vectors have to be considered when merging. The definition itself requires all weight vectors to be taken into account: a model is selected if and only if it is selected by at least one of the infinitely many weight vectors in W_\exists . The next lemma shows that at least exponentially many have to be considered. It will be later proved that exponentially many suffice.

Lemma 7 *There exist three formulae μ, F, F' on an alphabet of six variables such that every W_r such that $\Delta_\mu^{dh, W_r}(F, F') = \Delta_\mu^{dh, W_\exists}(F, F')$ contains at least two weight vectors.*

This lemma shows a pair of formulae that requires at least two weight vectors. This technical result has an abstract implication: since each weight vector encodes a specific way to compare the reliability of the formulae, merging with absent reliability cannot be reduced to merging with any specific reliability degree of the formulae. This makes sense, as reliability is not objective, like for example the values of variables, but subjective, since it is the strength of believing that a formula is true. If merging were always possible with a single weight vector, that weight vector could be considered as part of reality rather than beliefs.

The lemma requires two formulae of six variables. The claim actually holds for three variables, but the proof would be ad-hoc rather than simply referring to a previous lemma, and is therefore omitted. The claim does not hold for two variables, as proved by exhaustive analysis on the four possible models.

The construction in the lemma can be replicated over many distinct alphabets of six variables each. Each alphabet doubles the number of necessary weight vectors, leading to exponentiality.

Lemma 8 *There exists μ, F_1, \dots, F_m such that the size of every W_r for which $\Delta_\mu^{dh, W_r}(F, F') = \Delta_\mu^{dh, W_\exists}(F, F')$ is exponential in the size of the formulae.*

This result relies on an unbounded number of formulae to be merged. With two formulae, a number of weight vectors linear in the number of variables suffices.

7 Postulates

Merging with unknown weights depends on the distance function and the set of weight vectors. Some postulates for belief merging hold for all sets of weights vectors (IC0-IC2 and IC7), others only for some including W_{\exists} (IC3-IC6), and one does not hold for W_{\exists} (IC8). Merging with W_{\exists} cannot be expressed as a preorder, not even a partial one.

Postulates IC0-8 (Konieczny and Pino Pérez 2002a) cannot all hold, since a merging operator satisfying all of them can be expressed as a selection of models of μ that are minimal according to some total preorder that depends only on F_1, \dots, F_m . Actually, not even a partial preorder expresses merging with all possible weight vectors.

Theorem 4 *No partial preorder \leq depending on F_1 and F_2 only is such that $\Delta_{\mu}^{dh, W_{\exists}}(F_1, F_2) = \min(\text{mod}(\mu), \leq)$.*

A consequence of this theorem is that merging with all possible weight vectors W_{\exists} does not satisfy all postulates, since that would imply that merging could be expressed by a preorder. Some postulates are not satisfied. Others are.

Some postulates hold for every set of weight vectors, others only for some. Some postulates hold only if the distance function satisfies the triangle inequality, others hold even if $d(I, F)$ is not defined in terms of a distance among models $d(I, J)$. The latter requires $d(I, F) \in \mathbf{N}$ and $d(I, F) = 0$ if and only if $I \models F$. In the following summary, this case is described as “a model-formula distance”. In this section, E is sometimes used in place of F_1, \dots, F_m following the notation by Konieczny and Pino Pérez (2011). This simplifies some formulae.

- IC0 $\Delta_{\mu}^{d, W \cdot}(E) \subseteq \text{mod}(\mu)$
holds for every model-formula distance and non-empty set of weight vectors
- IC1 if μ is consistent, then $\Delta_{\mu}(E)$ is not empty
holds for every model-formula distance and non-empty set of weight vectors
- IC2 if $\bigwedge E$ is consistent with μ , then $\Delta_{\mu}^{d, W \cdot}(E) = \text{mod}(\mu) \cap \text{mod}(\bigwedge E)$
holds for every model-formula distance and non-empty set of weight vectors
- IC3 if $E_1 \equiv E_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}^{d, W \cdot}(E_1) = \Delta_{\mu_2}^{d, W \cdot}(E_2)$
holds for every model-formula distance and non-empty set of weight vectors that contains every permutation of every vector it contains (W_{\exists} has this property, as well as W_a for every $a \in \mathbf{N}$ with $a > 0$).
- IC4 if $F_1 \models \mu$ and $F_2 \models \mu$ then $\Delta_{\mu}^{d, W \cdot}(F_1, F_2) \cap \text{mod}(F_1)$ is not empty if and only if $\Delta_{\mu}^{d, W \cdot}(F_1, F_2) \cap \text{mod}(F_2)$ is not empty.
holds if $W \cdot$ contains every permutation of every vector it contains and d satisfies the triangle inequality: $d(I, K) + d(K, J) \geq d(I, J)$ (both dd and dh have this property); if any of these two conditions do not hold, a counterexample shows that the postulate does not hold
- IC5 $\Delta_{\mu}^{d, W \cdot'}(F_1, \dots, F_k) \cap \Delta_{\mu}^{d, W \cdot''}(F_{k+1}, \dots, F_m) \subseteq \Delta_{\mu}^{d, W \cdot}(F_1, \dots, F_m)$
requires $W \cdot$ to be the Cartesian product of two sets of weight vectors $W \cdot'$ and $W \cdot''$ whose vectors have size k and $m - k$, respectively
- IC6 if $\Delta_{\mu}^{d, W \cdot'}(F_1, \dots, F_k) \cap \Delta_{\mu}^{d, W \cdot''}(F_{k+1}, \dots, F_m)$ is not empty, it contains $\Delta_{\mu}^{d, W \cdot}(F_1, \dots, F_m)$

requires $W.$ to be the Cartesian product of two sets of weight vectors $W.'$ and W'' whose vectors have size k and $m - k$, respectively

IC7 $\text{mod}(\mu') \cap \Delta_{\mu}^{d,W.}(E) \subseteq \Delta_{\mu \wedge \mu'}^{d,W.}(E)$

holds for every model-formula distance and non-empty set of weight vectors

IC8 if $\text{mod}(\mu') \cap \Delta_{\mu}^{d,W.}(E)$ is not empty, then $\Delta_{\mu \wedge \mu'}^{d,W.}(E) \subseteq \Delta_{\mu}^{d,W.}(E)$

does not hold for the Hamming distance dh and the set of all weight vectors W_{\exists}

The formal proofs of these claims follow. First, postulates IC0, IC1, IC2 and IC7 hold for every non-empty set of weight vectors $W.$ and model-formula distance.

Lemma 9 *For every model-formula distance d and non-empty set of weight vectors $W.$, the merging operator $\Delta^{d,W.}$ satisfies postulates IC0, IC1, IC2 and IC7.*

Postulate IC3 includes the case where the order of the formulae is changed. This affects the weight vectors: they must be allowed to change their internal order accordingly.

Lemma 10 *If $W.$ contains every permutation of every vector it contains, then IC3 holds. For some set of weight vectors that does not include a permutation of one of its elements, IC3 does not hold.*

These lemmas do not require $d(I, F)$ to be defined in terms of a distance between models $d(I, J)$. The next one does, and additionally needs the triangle inequality.

Lemma 11 *If $W.$ contains every permutation of every vector it contains and d satisfies the triangle inequality $\forall I, J, K. d(I, K) + d(K, J) \geq d(I, J)$, then IC4 holds. For some set of weight vectors that does not include a permutation of one of its elements IC4 does not hold. The same for some distance not satisfying the triangle inequality.*

Since W_{\exists} is symmetric and both dd and dh satisfy the triangle inequality, Postulate IC4 holds in these two cases. Actually, for W_{\exists} the distance does not matter, and $\Delta_{\mu}^{d,W_{\exists}}(F_1, \dots, F_m)$ always contains some models of every F_i that is consistent with μ . If the maximal value of the distance from F_1 and from F_2 is k , the weight vectors $[k + 1, 1]$ and $[1, k + 1]$ suffice. The first guarantees that every model of F_1 is always better than one of $\neg F_1$, no matter how close the second is to F_2 . The same for the second weight vector.

This lemma shows an effect of the triangle inequality on belief merging. It is a quite natural requirement and is obeyed by both the drastic and the Hamming distance, but is mostly useless in belief merging (Konieczny and Pino Pérez 2011). Besides proving that a certain merging operator satisfies an additional postulate (Konieczny and Pino Pérez 2002a), so far it only seemed to affect the infinite-alphabet case (Chacón and Pino Pérez 2006) and the application of belief revision to case-based reasoning (Cojan and Lieber 2012).

Postulates IC5 and IC6 require special care even to be formulated. Informally, they tell that merging $F_1, \dots, F_k, F_{k+1}, \dots, F_m$ is the same as merging F_1, \dots, F_k , merging F_{k+1}, \dots, F_m and then conjoining the two results if they do not conflict. This is simple to express if no weights are involved, otherwise each of these three mergings

is defined over its set of weights. If these are unrelated, like $[1, \dots, 1, 1 \dots, 1]$ for the overall merge and $[10, 1, \dots, 1]$ and $[1, \dots, 10]$ for merging the two parts, the three results cannot be expected to be coherent.

This is why Postulates IC5 and IC6 cannot be said to be obeyed plain and simple. Rather, they are satisfied only when the sets of weights are related in the appropriate way.

Lemma 12 *If $\Delta_{\mu}^{d, W'}. (F_1, \dots, F_k) \cap \Delta_{\mu}^{d, W''}. (F_{k+1}, \dots, F_m)$ is not empty, it coincides with $\Delta_{\mu}^{d, W}. (F_1, \dots, F_m)$, where $W.$ is the Cartesian product of $W'.$ and $W''.$ (postulates IC5 and IC6).*

IC8 does not hold. The following counterexample shows that for the Hamming distance dh and the set of all weight vectors W_{\exists} .

Theorem 5 *There exist μ, μ', F_1 and F_2 such that $\text{mod}(\mu') \cap \Delta_{\mu}^{dh, W_{\exists}}(F_1, F_2)$ is not empty but $\Delta_{\mu \wedge \mu'}^{dh, W_{\exists}}(F_1, F_2) \not\subseteq \Delta_{\mu}^{dh, W_{\exists}}(F_1, F_2)$.*

This counterexample completes the analysis of the basic postulates IC0-IC8. Two additional ones exist: majority and arbitration. The first tells that a formula repeated enough times is entailed by the result of merging; the second was initially defined as the irrelevance of the number of repetitions, and has a newer definition that is difficult to summarize in words.

Majority does not hold with W_{\exists} . Not that it should. No matter how many times a formula is repeated, regardless of how many sources supports it, its negation may come from a single source that is more reliable than all the others together. When reliability is uncertain, this case has to be taken into account. It is not even uncommon in practice: many commonly held belief are in fact false.

Many commonly held beliefs are in fact false: Napoleon was short (Dunan 1963); diamonds had been typical gemstones for engagement rings since a long time (Epstein 1982) the red telephone is a telephone line, and one of its end is in the White House (Clavin 2013); meteorites are always hot when they reach the Earth's surface; flowering sunflowers turn to follow the sun (only the gems do); the Nazis issued an ultimatum before the Ardeatine massacre (something even witnesses of the time believe) (Mazzoni 2003, p. 155); fans in closed rooms kill people (many people in Korea believed this). A page on Wikipedia lists more than a hundred of commonly believed facts that are in fact false (Wikipedia 2017b). The material was enough for a 26-episodes TV show (Wikipedia 2017a).

A view of belief merging is that it formalizes the process of information aggregation by human agents. The above scenarios indicate that unanimity is often a driving mechanism of believing: hearing and reading many times that Napoleon was short leads to believing he was without questioning. Yet, unanimity is not majority. A single person with a funny haircut on TV may at least cast a doubt.

All of this shows that no matter how many times a fact is repeated, when no reliability information is present, it may still be falsified by a single reliable source. This is what the following theorem formally proves.

Theorem 6 *There exists F_1, F_2 such that $\Delta_{true}^{d, W_{\exists}}(F_1, F_2, \dots, F_2) \not\subseteq \text{mod}(F_2)$, where F_2 is repeated an arbitrary number of times.*

Majority does not hold in the most unconstrained case W_{\exists} where weights are arbitrary. This does not mean that majority never applies. It means that it does not apply when no information about the reliability of the sources is present. In many other cases, it applies. When sources are considered equally reliable, it applies. It holds for $W_{=} = \{[1, \dots, 1]\}$, which formalizes exactly this situation. The operator $\Delta_{\mu}^{d, W_{=}}(F_1, \dots, F_m)$ coincides with the operator $\Delta_{\mu}^{d, \Sigma}(F_1, \dots, F_m)$ defined by Konieczny and Pino Pérez (2011), which satisfies majority (Konieczny and Pino Pérez 2002a).

Where is the boundary between majority and non-majority operators? The majority condition tells that no source is arbitrarily more reliable than the others. Since reliability is formalized by weights, it tells that no weight is arbitrarily large.

Theorem 7 *If all weights in the vectors in W are lower than a constant and d is an arbitrary distance, for every μ and $F_1, \dots, F_o, F_{o+1}, \dots, F_m$, there exists n such that $\Delta_{\mu}^{d, W_{\cdot}}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m) \subseteq \Delta_{\mu}^{d, W_{\cdot}}(F_{o+1}, \dots, F_m)$, where the formulae from F_{o+1} to F_m are repeated n times.*

Weights bounded by a constant indicate that no source is ever considered arbitrarily more reliable than the others. This is a relevant case. So it is the case where a source may be much more reliable than the others, as motivated above by the example of the commonly believed facts that are in fact false and previously by the example of the three doctors.

Arbitration was initially defined as the opposite condition of irrelevance of the number of repetitions (Konieczny and Pino Pérez 1998; Meyer 2001). This property holds for W_{\exists} . The following theorem proves an equivalent formulation of it.

Lemma 13 *For every μ, F_1, \dots, F_m it holds:*

$$\Delta_{\mu}^{d, W_{\exists}}(F_1, \dots, F_m) = \Delta_{\mu}^{d, W_{\exists}}(F_1, \dots, F_m, F_m)$$

A newer version of the arbitration postulate is expressed in terms of the preorder between models as: if $I <_{F_1} J, I <_{F_2} J'$ and $J \equiv_{F_1, F_2} J'$ then $I <_{F_1, F_2} J$. As proved by Theorem 4, merging with absent reliability cannot be expressed as a preorder, total or otherwise. The expression of the postulate in terms of formulae is even more convoluted, and is not clear whether it makes sense when merging is not expressible in terms of a preorder.

8 One reliable source

The case of one reliable, unspecified source is captured by the set of weights W_a when the number a is large enough: more than the maximal distance between a model and the formulae. For these weight vectors, a single source may take over all other ones. Such a situation is not unlikely in practice, as exemplified by the scenario of the three doctors: the one specialized in the field of the actual illness is almost certainly right, but the actual illness is debated. Another example is that of the facts commonly held

true, where the opinion of a real expert may confute them regardless of how many people believe them.

The assumption of one reliable source gives rise and motivates a condition already in the literature: the disjunctive property. Technically, merging by the weight vectors W_a with a sufficiently large a ensures the disjunctive property. Conceptually, the disjunctive property formalizes the assumption of one reliable source.

The disjunctive property was defined on two formulae as Postulate 7 by Liberatore and Schaerf (1998) and later generalized to an arbitrary number of formulae with integrity constraints by Everaere et al. (2010a). In terms of models, it has a simple and intuitive expression. Every model is a possible state of the world; merging only selects the ones that at least one of the sources considers possible. In formulae, a model I is in the result of merging only if $I \models F_i$ for at least one of the merged formulae F_i . Since I must also satisfy the integrity constraints μ , this requirement is lifted when none of the formulae F_i is consistent with μ .

Definition 4 A merging operator Δ is disjunctive if it satisfies the disjunctive property: $\Delta_\mu(F_1, \dots, F_m) \subseteq \text{mod}(F_1 \vee \dots \vee F_m)$ holds if at least one of the formulae F_1, \dots, F_m is consistent with μ .

This condition is not satisfied by $\Delta^{dh, W_=}$. As a result, is not satisfied by Δ^{dh, W_\exists} either, since $W_= \subset W_\exists$.

The disjunctive property fails in the case of equal reliability (formalized by $W_=$) and completely absent reliability information (formalized by W_\exists). This is one part of the claim that the disjunctive property is a formalization of the assumption of one reliable source. The other is that it succeeds when one unspecified source is much more reliable than the others, formalized by W_a with a large value of a .

Definition 5 (Liberatore and Schaerf 1998)

Merging by closest pairs of models is defined from the ordering between pairs of models $\langle I, J \rangle \leq_{dh} \langle I', J' \rangle$ if and only if $dh(I, J) \leq dh(I', J')$ by selecting the models in all minimal pairs:

$$F_1 \Delta_D F_2 = \{I, J \mid \langle I, J \rangle \in \min(\{\langle K, L \rangle \mid K \in \text{mod}(F_1), \\ L \in \text{mod}(F_2)\}, \leq_{dh})\}$$

This definition is framed in the general framework of merging with the set of weights $W_{n+1} = \{[1, n+1], [n+1, 1]\}$, the specific form of W_a when a is the number of variables increased by one and merging is between two formulae. Since n is the maximal value of the Hamming distance, this set W_a characterizes exactly the assumption that one of the two formulae is reliable, but no information about which is present. This assumption leads to merging by closest pairs of models.

Theorem 8 *For every pair of satisfiable formulae F_1 and F_2 over an alphabet of n variables, it holds $F_1 \Delta_D F_2 = \Delta_{true}^{dh, W_{n+1}}(F_1, F_2)$.*

A disjunctive operator on m formulae is obtained similarly when all formulae are consistent and the integrity constraints are void: $\mu = true$.

Theorem 9 *For every distance d bounded by k , if F_1, \dots, F_m are satisfiable then $\Delta_{true}^{d, W_{km}}(F_1, \dots, F_m)$ is a disjunctive merging operator.*

The weight vectors in these theorems provide an alternative view of the disjunctive property. Rather than being a principle by itself, it is a formalization of the assumption that a single source is fully reliable, but which one is not specified. What the other sources tell is kept into account, but not as much as contradicting the reliable source. In terms of formulae, the other formulae help in selecting some of the models of the reliable formula, but do not drive the choice outside the set of these models. The result of this selection is always a group of these models, a subset of the models of the reliable formula. However, which formula is reliable is not specified. It may be every one of them. For each one, merging may only select some of its models. Overall, only the models of the formulae can be in the result of merging.

This mechanism interprets the principle of indifference in belief merging in the right way: rather than assuming that all sources are equally reliable, one of them is taken as completely right, but this is done for each of them in turn. Indifference is realized by symmetry, not equality.

Many interesting operators are not disjunctive (Konieczny and Pino Pérez 2011; Everaere et al. 2010a). An operator may not be disjunctive because it interprets the principle of indifference in a different way, or because it does not follow the principle of indifference at all. Indifference is not a universal rule. It is the formalization of the absence of reliability information. When the credibility of the sources are given, or is believed to be equal or comparable, the principle of indifference does not apply. Forcing it on all operators would be a gross mistake.

9 Other aggregator functions

Merging selects models of minimal sum of weighted distances. The sum was historically the first way of combining distances. Others were later invented. Few of the properties studied in this article change when switching to other mechanisms: sum of

powers, leximax and leximin ordering (Konieczny and Pino Pérez 2011). The main differences are that leximax produces all undominated models even if the codomain of the distance function is not binary and that leximin produces the models of maximal-size maxcons instead of all maxcons when using the drastic distance.

9.1 Merging by sum of powers

Instead of adding the weighted distances, merging by sum of powers adds their weighted powers (Konieczny and Pino Pérez 2002b). The power could be the square (power 2) or an arbitrary positive integer n . The ordering between two weighted vector distances is:

$$[v_1, \dots, v_m] <^n [u_1, \dots, u_m] \text{ iff } \sum_{1 \leq i \leq m} w_i v_i^n < \sum_{1 \leq i \leq m} w_i u_i^n$$

This order defines merging by a single weight vector:

$$\begin{aligned} \Delta_{\mu}^{d,W,\Sigma^n}(F_1, \dots, F_m) &= \{I \in \text{mod}(\mu) \mid \nexists J \in \text{mod}(\mu). \\ &W \cdot d(I, F_1, \dots, F_m) <^n W \cdot d(J, F_1, \dots, F_m)\} \end{aligned}$$

Merging by a set of weight vectors is the union of merging by each:

$$\Delta_{\mu}^{d,W,\Sigma^n}(F_1, \dots, F_m) = \bigcup_{W \in W} \Delta_{\mu}^{d,W,\Sigma^n}(F_1, \dots, F_m)$$

Given any distance function d , the function defined by $d'(I, F_i) = d(I, F_i)^n$ is also a distance function. It is binary if and only if d is binary. Therefore, all results involving arbitrary distance functions or arbitrary binary distance functions carry over from merging by sum to merging by sum of powers:

- Lemma 1: merging only produces models that are not strictly dominated by others;
- Theorem 1: if the distance function has binary codomain, merging produces exactly the models that are not strictly dominated by others;
- Theorem 3: merging by the drastic distance produces the union of the models of all maxcons.

The latter holds because dd is binary, which implies that $d'(I, F_i) = dd(I, F_i)^n$ is binary as well. Therefore, merging produces exactly the undominated models. This is the same as the result of merging by the drastic distance without powering the distances. The latter is proved by Theorem 3 to be the union of the models of all maxcons.

The results requiring specific values may not carry over when powering the distances. For example, the proof of Theorem 2 involves three models with vector distances $[3, 0]$, $[2, 2]$ and $[0, 3]$, where the second is not dominated by the others but is not in the result of merging. This is not the case when squaring the distances, as $[4, 4]$ is less than $[9, 0]$ and $[0, 9]$. Yet, this lost property is found in the distance vectors $[4, 0]$, $[3, 3]$ and $[0, 4]$.

Theorem 10 For some distance d with codomain of size three, there exists I, μ and F_1, \dots, F_m such that $I \models \mu$ and $I \notin \Delta_\mu^{d, W_{\exists}, \Sigma^2}(F_1, \dots, F_m)$ hold, but $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ does not hold for any $J \models \mu$.

The preconditions of this theorem are satisfied by the Hamming distance, since this distance produces every possible set of distance vectors by Theorem 6. As a result, merging by squared distances may not generate some models of μ that are not dominated by others.

9.2 Leximax merging

The leximax ordering is the lexicographic order between two vectors sorted in descending order (Konieczny and Pino Pérez 2002a). The vectors $W \cdot d(I, F_1, \dots, F_m)$ are sorted so that each element is less than or equal to the previous, and compared according to the lexicographic order. The minimal models form the result of leximax merging with a single weight vector $\Delta_\mu^{d, W, \text{leximax}}(F_1, \dots, F_m)$. The result of merging with a set of weight vectors $\Delta_\mu^{d, W, \text{leximax}}(F_1, \dots, F_m)$ is the union of merging with each.

The dominance ordering is maintained when sorting vectors in descending order. If the first vector is less than or equal to the second, it remains so after sorting both.

Lemma 14 If $v = [v_1, \dots, v_m]$ and $u = [u_1, \dots, u_m]$ are two vectors of integers such that $v_i \leq u_i$ holds for every index i , the same holds for the result of sorting v and u in descending order.

This result only proves that the ordering before sorting implies that after. The following results require the ordering to be strict. This case is covered by the following lemma.

Lemma 15 If $v = [v_1, \dots, v_m]$ and $u = [u_1, \dots, u_m]$ are two vectors of integers such that $v_i \leq u_i$ holds for every index i and $v_i < u_i$ for some index i , the same holds for the result of sorting v and u in descending order.

This result allows extending Lemma 1 to leximax merging: it does not generate any model dominated by another.

Lemma 16 For every distance d , vector of weights $W \in W_{\exists}$ and model I , if $I \in \Delta_\mu^{d, W, \text{leximax}}(F_1, \dots, F_m)$ then $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ holds for no model J of μ .

Leximax merging generates exactly the undominated models. This is analogous to Lemma 2, but does not require the codomain of the distance function to be binary.

Lemma 17 If I is a model of μ and $d(I, F_1, \dots, F_m)$ is not strictly dominated by the vector of distances of any other model of μ , then there exists W such that $I \in \Delta_\mu^{d, W, \text{leximax}}(F_1, \dots, F_m)$.

Combining the last two results: leximax merging generates exactly the models of μ that are not dominated by others. This links leximax to maxcons.

Theorem 11 *For every consistent formulae μ, F_1, \dots, F_m , the following equality holds:*

$$\Delta_{\mu}^{d, W_{\exists}, \text{leximax}}(F_1, \dots, F_m) = \bigcup_{S \in \text{maxcon}_{\mu}(F_1, \dots, F_m)} \text{mod}(\wedge S)$$

9.3 Leximin merging

The leximin ordering is similar to the leximax ordering, but sorts vectors in ascending order instead of descending (Evertaere et al. 2010a). Leximin merging generates the minimal weighted vectors $\Delta_{\mu}^{d, W, \text{leximin}}(F_1, \dots, F_m)$ according to this ordering. The union of these for all $W \in W$ defines $\Delta_{\mu}^{d, W, \text{leximin}}(F_1, \dots, F_m)$.

A model dominated by another is not selected by leximin merging. Its distance vector is greater than that of the other. This property weathers multiplying each distance by the same weight. It also weathers sorting the two vectors in descending order by Lemma 15. It again weathers inverting the order of the two vectors. Therefore, the dominated model is greater than another in the leximin order, and is therefore not selected by leximin merging.

The converse is however not the case even if the distance function has binary codomain. Yet, maxcons are still related to leximin merging with the drastic distance.

The counterexample is based on three models of distance vectors $[1, 1, 0, 0]$, $[0, 0, 1, 0]$ and $[0, 0, 0, 1]$. Multiplying these vectors by the weights $[w_1, w_2, w_3, w_4]$ results in $[w_1, w_2, 0, 0]$, $[0, 0, w_3, 0]$ and $[0, 0, 0, w_4]$. Since all weights are larger than zero, ordering these vectors produces $[0, 0, w_1, w_2]$, $[0, 0, 0, w_3]$ and $[0, 0, 0, w_4]$ if $w_1 \leq w_2$, otherwise $[0, 0, w_2, w_1]$, $[0, 0, 0, w_3]$ and $[0, 0, 0, w_4]$. Regardless, the first vector is not minimal because both w_1 and w_2 are larger than zero.

Model selection is primarily based on the length of the initial string of zeros. The drastic distance gives zero when the model satisfies the formula. Therefore, the number of zeros is the number of satisfied formulae. The selected models are therefore those of the largest-size maxcons.

Theorem 12 *For every consistent formulae μ, F_1, \dots, F_m , the following equality holds*

$$\Delta_{\mu}^{dd, W_{\exists}, \text{leximin}}(F_1, \dots, F_m) = \bigcup_{S \in \text{cardcons}_{\mu}(F_1, \dots, F_m)} \text{mod}(\wedge S)$$

where

$$\text{cardcons}_{\mu}(F_1, \dots, F_m) = \{S \in \text{maxcon}_{\mu}(F_1, \dots, F_m) \mid \nexists S' \in \text{maxcon}_{\mu}(F_1, \dots, F_m) \cdot |S| < |S'|\}$$

This theorem implies that leximin merging does not generate all models of all maxcons. Therefore, it may not produce all undominated models.

10 Sources providing multiple formulae

The previous sections are about combining a number of independent formulae. This is the basic problem of belief merging: each formula comes from a different source; therefore, their reliabilities are independent. This is formalized by the weights being unconstrained in the set W_{\exists} .

When a source provides more than one formula, each of them is as reliable as its source. The same mechanism employing the weighted sum of the drastic or Hamming distance can be used, but the weights are associated to the sources rather than to the formulae. All formulae from the same source have the same reliability and therefore the same weight. This condition is close in spirit to the unit partitions by Booth and Hunter (2018).

Technically, each source is represented by a set of formulae S_i . Its reliability is encoded by a positive integer w_i . Given a set $\{S_1, \dots, S_m\}$ of such sources, merging is done by selecting the minimal models of the integrity constraints μ according to this evaluation:

$$v(I) = \sum_{S_i} w_i \sum_{F_i \in S_i} d(I, F_i)$$

This is the DA^2 operator (Konieczny et al. 2004) with the sum as intra-source aggregation and the weighted sum as the inter-source aggregation.

The sum is subject to the problem of manipulation: a source may provide the same formula multiple times in order to influence the final result (Chopra et al. 2006); this is a problem especially when merging preferences, but not when merging beliefs from sources of unspecified reliability. Even if a source provides the same formula a thousand times, one of the considered alternatives is that the weight of this source is a thousand times smaller than the others, making such a manipulation ineffective.

The only technical result in this section is that merging with the drastic distance is not the same as the union of the models of the maxcons. This is proved by the following sources with $\mu = true$.

$$S_1 = \{x, y, z\}$$

$$S_2 = \{\neg x, \neg y\}$$

$$S_3 = \{\neg x, \neg z\}$$

One of the maxcons of $\{x, y, z, \neg x, \neg y, \neg z\}$ is $\{x, \neg y, \neg z\}$, which is not obtained when merging with all possible weights. Intuitively, to include the formula x from S_1 in the result, that formula needs to count at least twice as much as each formula $\neg x$ from S_2 and S_3 , but this implies the same for y and z , which excludes $\neg y$ and $\neg z$.

Formally, let the weights of the sources be w_1, w_2, w_3 . The weighted distances of some relevant models are:

$$\begin{aligned} I &= \{x, \neg y, \neg z\} & v(I) &= w_1 2 + w_2 1 + w_3 1 = w_1 2 + w_2 + w_3 \\ J &= \{\neg x, \neg y, \neg z\} & v(J) &= w_1 3 + w_2 0 + w_3 0 = w_1 3 \\ K &= \{x, y, z\} & v(K) &= w_1 0 + w_2 2 + w_3 2 = w_2 2 + w_3 2 \end{aligned}$$

In order for I to be minimal, $v(I)$ must be less than or equal to $v(J)$ and $v(K)$:

$$\begin{aligned} w_1 2 + w_2 + w_3 &\leq w_1 3 \\ w_1 2 + w_2 + w_3 &\leq w_2 2 + w_3 2 \end{aligned}$$

This system of inequalities is infeasible. The first implies $w_2 + w_3 \leq w_1$, which makes the right-hand side of the second become less than or equal to $w_1 2$, while it should instead be greater than or equal to $w_1 2 + w_2 + w_3$, and therefore greater than $w_1 2$. This proves that $\{x, \neg y, \neg z\}$ is not a minimal model for any weight vector.

A similar example shows that merging does not generate the models of the conjunctions of the formulae of each source. Let $S_1 = \{x, y\}$ and $S_2 = \{\neg x, \neg y\}$. The only maxcons of $\{\wedge S_1, \wedge S_2\}$ are $x \wedge y$ and $\neg x \wedge \neg y$, but merging with all possible weights selects all four models over x and y , since this is the result when $W = [1, 1]$.

11 Conclusions

Sometimes the information to be merged comes from sources of equal reliability. In such cases, merging with equal weights is correct. But when no reliability information on the sources is present, assuming weights equal is unwarranted. The difference is not only conceptual but also technical. Theorem 4 shows that merging with absent reliability cannot in general be reduced to a preorder among models, not even a partial one.

A result emerged in the study of this setting is a motivation for merging by maxcons (Baral et al. 1992). This mechanism has sometimes been considered unsuitable for merging because it disregards the distribution of information among sources (Konieczny 2000; Konieczny and Pino Pérez 2011). Such a distribution is important when the sources have the same reliability, or more generally their reliability is given. It is not when reliability information is absent. Theorem 3 shows that merging with maxcons is the same as merging with the drastic distance in absence of any reliability assessment on the sources. The number of repetitions of a formula is irrelevant to this kind of merging—as it should. A formula only occurring once may come from a very reliable source, while its negation is supported only by unreliable sources. Without any assumption on the reliability of the sources, this is a situation to take into account. At the same time, Theorem 3 is limited to merging by the drastic distance. It does not apply when a finer distance would provide a more informative result of merging.

This article not only backs merging by maxcons, but more generally merging by the (MI) postulate: the number of repetitions of a formula is irrelevant to merging.

Of course, there are many cases where this postulate should not hold; when their reliability is equal, two sources providing a formula give twice the support for it; the same if no source is arbitrarily more reliable than the others, as formalized by weights bounded by a constant. But when reliability information is completely absent, every numeric evaluation is irrelevant, including doubling the support for a formula like in this case. As already discussed by Meyer (2001), Postulate (MI) may sometimes be right; it is inconsistent with the other postulates IC0-IC8, but the fault is on them. While Meyer blames Postulate 4 of merging without integrity constraints (Konieczny and Pino Pérez 1998), merging with absent reliability conflicts with IC8.

The most significant outcome of this article is a motivation of choices made in the past, uncovering the implicit assumptions they are based on: maxcons come from completely absent reliability information; the disjunctive property comes from the assumption that one formula is totally reliable, but which one is not specified.

A minor technical contribution of this article is a case for the triangle inequality of the distance function. This property had only a couple of applications in belief revision and merging so far (Chacón and Pino Pérez 2006; Cojan and Lieber 2012; Konieczny and Pino Pérez 2002a), but is generally not required (Konieczny and Pino Pérez 2011). The new consequence of it shown in this article is that it allows satisfying Postulate IC4 when merging in absence of reliability information.

Most results in this article are on merging based on the weighted distances of models from the formulae to merge. The overall picture does not change much when switching to other ways of combining distances like the sum of powers and the leximax and the leximin ordering.

A comparison with related work follows.

The weighted distance from a set of formulae was first used for merging by Revesz (1997), and investigated by Lin (1996). Lin and Mendelzon (1999) and Konieczny and Pino Pérez (1998) used the unweighted sum. These articles assume either equal or fixed weights, not varying weights like this one.

Benferhat et al. (2014) consider the related problem of commensurability: when the sources themselves assess the reliability of the formulae they provide, they may not use the same scale; this is related to a similar issue in social choice theory. Their study and the present one differ in formalism (ranked bases instead of formulae with distance functions), but they share the principle of considering a set of alternative reliability assessments. There is however an early point of departure: Benferhat et al. 2014 distill a single preorder and then select the models that are minimal according to it; Theorem 4 shows that the same cannot be done in the settings of the present article. A point of contact is the case of the drastic distance: Theorem 1 could be alternatively proved from a result by Benferhat et al. (2014, Propositions 1,2,8).

That reliability information may be partially absent is mentioned by Konieczny (2004) as the starting assumption of his model of belief merging as a game: “The hypothesis for those operators is that all the sources are a priori reliable, or that we know that some sources are less reliable than the others, but without knowing which ones.” The approach taken is however very different, as it proceeds by iteratively assessing the reliability of the sources based on the others.

A related question is what changes between the synthetic and epistemic view of belief merging (Everaere et al. 2010b). When attempting to establish the truth (the

epistemic view of belief merging), a single opinion from an expert may take over many other ones. When forming a unified opinion of a group (the synthetic view), deciding by majority may look the only way to proceed. As a matter of facts, majority influence research (Gardikiotis 2011) shows otherwise. A minority view may end up prevailing. An example are trial juries, where the opinion of few jurors sometimes forms the final judgment.

Accepting what is true according to all possible relative reliabilities is analogous to drawing the consequences that hold in all probability measures in a set (Halpern and Tuttle 1993), and can be seen as the formal logic version of the “worst scenario” in economics: “*the firm may not be certain about the “relative plausibility” of these boom probabilities. [...] if the firm acts in accordance with certain sensible axioms, then its behavior can be characterized as being uncertainty-averse: when the firm evaluates its position, it will use a probability corresponding to the “worst” scenario*” (Nishimura and Ozaki 2007). Belief revision and merging aim at the most knowledge that can be justifiably and consistently obtained; therefore, minimal knowledge takes the place of the least profit, and the worst scenario for a formula is one where it is false.

Acknowledgements The author thanks the reviewers for their valuable indications on the previous versions of this article.

Funding Open access funding provided by Università degli Studi di Roma La Sapienza within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Proofs

Lemma 1 *For every distance d , vector of weights $W \in W_{\exists}$ and model I of μ , if $I \in \Delta_{\mu}^{d,W}(F_1, \dots, F_m)$ then $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ holds for no model J of μ .*

Proof The claim is proved in the opposite direction: $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ entails $I \notin \Delta_{\mu}^{d,W}(F_1, \dots, F_m)$.

Since weights are all strictly positive, $d(J, F_1, \dots, F_m) \leq d(I, F_1, \dots, F_m)$ entails $W \cdot d(J, F_1, \dots, F_m) \leq W \cdot d(I, F_1, \dots, F_m)$ and $d(I, F_1, \dots, F_m) \not\leq d(J, F_1, \dots, F_m)$ entails $W \cdot d(I, F_1, \dots, F_m) \not\leq W \cdot d(J, F_1, \dots, F_m)$. These two consequences together are $W \cdot d(J, F_1, \dots, F_m) < W \cdot d(I, F_1, \dots, F_m)$, which proves that I is not a model of minimal distance weighted by W , and is not therefore in $\Delta_{\mu}^{d,W}(F_1, \dots, F_m)$. \square

Lemma 2 *If the codomain of the distance function d is a subset of cardinality two of \mathbf{N} , I is a model of μ , $d(I, F_1, \dots, F_m)$ is not strictly dominated by the vector of distances of any other model of μ , then there exists W such that $I \in \Delta_\mu^{d,W}(F_1, \dots, F_m)$.*

Proof Since d is a distance function, it holds $d(I, I) = 0$ for every model I . Therefore, one of the two values of its codomain is 0. Since this codomain is a subset of \mathbf{N} , the other value b is greater than zero.

The weight vector W is $[w_1, \dots, w_m]$, where $w_i = m + 1$ if $d(I, F_i) = 0$, and $w_i = 1$ otherwise.

For every other model J of μ , the weighted distance of J is proved to be greater than or equal to that of I . Two cases are possible: either $d(J, F_i) = 0$ for every F_i such that $d(I, F_i) = 0$, or this is not the case for at least one formula F_i .

The first case is that $d(J, F_i) \leq d(I, F_i)$ holds for every F_i , which implies $d(J, F_1, \dots, F_m) \leq d(I, F_1, \dots, F_m)$. Since J does not dominate I by assumption, $d(I, F_1, \dots, F_m) \not\leq d(J, F_1, \dots, F_m)$ is false, which means that $d(I, F_1, \dots, F_m) \leq d(J, F_1, \dots, F_m)$ is true. The distance vectors of I and J are the same. Therefore, multiplying both by W produces the same result. This proves that I is minimal.

The second case is that $d(I, F_i) = 0$ and $d(J, F_i) = b$ holds for some F_i . If $k > 0$ is the number of formulae F_i such that $d(I, F_i) = 0$, the weighted distance of I is:

$$W \cdot d(I, F_1, \dots, F_m) = (m + 1)k0 + 1(m - k)b = (m - k)b < mb$$

Only $d(J, F_i) = b$ is known, the distance of J from the other formulae may be either 0 or b . Assuming it is 0 for all of them leads to the minimal possible weighted distance, which is:

- one formula has distance b ; since $d(I, F_i) = 0$, the weight is $m + 1$;
- the other formulae have all distance 0.

The weighted distance of J is therefore:

$$W \cdot d(J, F_1, \dots, F_m) \geq (m + 1)1b = mb + b > mb$$

The weighted distance of I is proved above to be less than mb . □

Theorem 1 *If the codomain of the distance d is a subset of cardinality two of \mathbf{N} , then $\Delta_\mu^{d,W\exists}(F_1, \dots, F_m)$ is the set of all models of μ of minimal distance vector according to the dominance ordering.*

Proof Lemma 1 proves that models of minimal weighted distance are never strictly dominated by any other model of μ . By Lemma 2, if the codomain of d is binary, every model that is not strictly dominated has a weight vector W that makes its weighted distance minimal. Since $W\exists$ contains all weight vectors, the claim is proved. □

Lemma 3 *If μ has three models of distance $[3, 0]$, $[2, 2]$ and $[0, 3]$ from F_1 and F_2 , then $\Delta_\mu^{d,W\exists}(F_1, F_2)$ does not contain the model at distance $[2, 2]$.*

Proof If the model at distance [2, 2] were in the result of merging, it would be minimal. This implies the following set of linear inequalities for some $W = [w_1, w_2]$.

$$\begin{aligned} w_1 2 + w_2 2 &\leq w_1 3 + w_2 0 \\ w_1 2 + w_2 2 &\leq w_1 0 + w_2 3 \end{aligned}$$

The first implies $w_2 2 \leq w_1$, the second $w_1 2 \leq w_2$: each weight is at least twice the other. No positive values may satisfy this condition. \square

Theorem 2 For some distance d with codomain of size three, there exists I, μ and F_1, \dots, F_m such that $I \models \mu$ and $I \notin \Delta_\mu^{d, W \exists}(F_1, \dots, F_m)$, but $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ does not hold for any $J \models \mu$.

Proof This is shown on the codomain $\{0, 2, 3\}$ and a formula μ with three models of distance vectors $[0, 3], [2, 2] \in [3, 0]$ from F_1 and F_2 . That such formulae exist is later proved by Lemma 6 for the Hamming distance. To obtain the right codomain $\{0, 2, 3\}$ the distance is modified by setting $dh^l(I, F_i) = 3$ for every model K such that $dh(K, F_i) \notin \{0, 2, 3\}$. This change does not affect the distance of the three considered models.

None of the three distance vectors is strictly dominated by another. However, the previous lemma shows that $[2, 2]$ is not minimal for any weight vector. \square

Lemma 4 For every pair of models I and J and every formulae F_1, \dots, F_m , the following two conditions are equivalent:

$$\begin{aligned} dd(I, F_1, \dots, F_m) &\leq dd(J, F_1, \dots, F_m) \\ \text{subsat}(J, F_1, \dots, F_m) &\subseteq \text{subsat}(I, F_1, \dots, F_m) \end{aligned}$$

Proof The inequality $dd(I, F_1, \dots, F_m) \leq dd(J, F_1, \dots, F_m)$ is the same as $dd(I, F_i) \leq dd(J, F_i)$ for each index i . The only possible values for $dd(I, F_i)$ are 0 when $I \models F_i$ and 1 when $I \not\models F_i$. The same holds for $dd(J, F_i)$. As a result, $dd(I, F_i) \leq dd(J, F_i)$ holds if and only if $J \models F_i$ implies $I \models F_i$. Since this is the case for every index i , all formulae satisfied by J are also satisfied by I . Since $\text{subsat}(J, F_1, \dots, F_m)$ is the set of formulae satisfied by J and $\text{subsat}(I, F_1, \dots, F_m)$ is the set of formulae satisfied by I , the claim follows. \square

Lemma 5 A model I of μ satisfies some element of $\text{maxcon}_\mu(F_1, \dots, F_m)$ if and only if $\text{subsat}(I, F_1, \dots, F_m) \subset \text{subsat}(J, F_1, \dots, F_m)$ holds for no model J of μ .

Proof Let I be a model of μ such that $\text{subsat}(I, F_1, \dots, F_m) \subset \text{subsat}(J, F_1, \dots, F_m)$ holds for no model J of μ . The set $\text{subsat}(I, F_1, \dots, F_m)$ is proved to be a maxcon. This set is consistent with μ because I satisfies both. It is also maximally so. Otherwise, $\text{subsat}(I, F_1, \dots, F_m) \cup \{\mu, F_i\}$ would be consistent for some $F_i \notin \text{subsat}(I, F_1, \dots, F_m)$. Consistency implies the existence of a model $J \models \text{subsat}(I, F_1, \dots, F_m) \cup \{\mu, F_i\}$. Since J satisfies all these formulae, $\text{subsat}(J, F_1, \dots, F_m)$ contains all of them: $\text{subsat}(I, F_1, \dots, F_m) \cup \{F_i\} \subseteq$

$subsat(J, F_1, \dots, F_m)$. This implies $subsat(I, F_1, \dots, F_m) \subset subsat(J, F_1, \dots, F_m)$ for a model J that also satisfies μ , contrary to assumption.

Let J be a model of μ such that $subsat(I, F_1, \dots, F_m) \subset subsat(J, F_1, \dots, F_m)$. The claim is that I is not in any maxcon. By contradiction, let M be such a maxcon. Since all its formulae satisfy I , it holds $M \subseteq subsat(I, F_1, \dots, F_m)$. By assumption, this set is strictly contained in $subsat(J, F_1, \dots, F_m)$ for some $J \models \mu$. Since J satisfies both $subsat(J, F_1, \dots, F_m)$ and μ , this other set $subsat(J, F_1, \dots, F_m)$ is consistent with μ , contradicting the assumption that M is maximally consistent with μ . \square

Theorem 3 For every consistent formulae μ, F_1, \dots, F_m , the following equality holds:

$$\Delta_{\mu}^{dd, W\exists}(F_1, \dots, F_m) = \bigcup_{S \in \maxcon_{\mu}(F_1, \dots, F_m)} \text{mod}(\wedge S)$$

Proof All elements S of $\maxcon_{\mu}(F_1, \dots, F_m)$ contain μ by definition. Therefore, all conjunctions $\wedge S$ satisfy μ . All their models I satisfy μ . By Lemma 5, these are exactly the models such that $subsat(I, F_1, \dots, F_m) \subset subsat(J, F_1, \dots, F_m)$ holds for no other model J of μ . This is equivalent to $dd(J, F_1, \dots, F_m) < dd(I, F_1, \dots, F_m)$ by Lemma 4. Therefore, these models I are the models of μ that are not strictly dominated by other models of μ . Since the codomain of dd is binary, these are the models of $\Delta_{\mu}^{dd, W\exists}(F_1, \dots, F_m)$ by Lemma 1. \square

Lemma 6 Given an arbitrary set of distance vectors of m elements each, all bounded by an integer n , for some formulae μ and F_1, \dots, F_m over nm variables the vectors of Hamming distances from the models of μ to F_1, \dots, F_m are exactly the given set of distance vectors.

Proof Formulae μ and F_1, \dots, F_m are built over the set of variables $\{x_j^i \mid 1 \leq j \leq n, 1 \leq i \leq m\}$. Each formula F_i is a conjunction of some of them.

$$F_i = x_1^i \wedge \dots \wedge x_n^i$$

Given a model I , its closest model of F_i has all variables x_1^i, \dots, x_n^i positive and the same evaluation of I on the other variables. Therefore, $dh(I, F_i)$ is the number of variables x_1^i, \dots, x_n^i assigned false by I .

For each distance vector $[d_1, \dots, d_m]$ among the given ones, μ has the following model:

$$\bigcup_{1 \leq i \leq m} \{\neg x_j^i \mid 1 \leq j \leq d_i\} \cup \{x_j^i \mid d_i < j \leq n\}$$

For each i , this model has d_i negative variables among x_1^i, \dots, x_n^i ; therefore, $dh(I, F_i) = d_i$. As a result, $dh(I, F_1, \dots, F_m) = [d_1, \dots, d_m]$. Since μ has one such model for each of the given distance vectors, the claim is proved. \square

Lemma 7 *There exist three formulae μ, F, F' on an alphabet of six variables such that every W_r such that $\Delta_\mu^{dh, W_r}(F, F') = \Delta_\mu^{dh, W_\exists}(F, F')$ contains at least two weight vectors.*

Proof By Lemma 6, given distances [3, 0], [1, 1] and [0, 3], there exist formulae μ, F_1, F_2 over six variables such that the three models of μ have these distance vectors.

All three distance vectors are minimal for some $W \in W_\exists$. In particular, the first two are minimal for $W = [2, 4]$, the third is minimal for $W = [4, 1]$. This proves that $\Delta_\mu^{dh, W_\exists}(F_1, \dots, F_m)$ contains all three models of μ .

Contrary to the claim, a single weight vector is assumed to produce the same result. Since the model at distance [3, 0] is minimal, its weighted distance is less than or equal to that of the model at distance [1, 1], and the same for [0, 3]:

$$\begin{aligned} [w_1, w_2] \cdot [3, 0] &\leq [w_1, w_2] \cdot [1, 1] \\ [w_1, w_2] \cdot [0, 3] &\leq [w_1, w_2] \cdot [1, 1] \end{aligned}$$

Expressing the two vector products explicitly:

$$\begin{aligned} w_1 \cdot 3 &\leq w_1 + w_2 \\ w_2 \cdot 3 &\leq w_1 + w_2 \end{aligned}$$

Since all weights are positive, the left-hand and right-hand sides of these inequalities can be added, leading to $w_1 \cdot 3 + w_2 \cdot 3 \leq w_1 + w_2$, which is impossible for strictly positive integers. This proves that no single weight vector produces the same merging of W_\exists . □

Lemma 8 *There exists μ, F_1, \dots, F_m such that the size of every W_r for which $\Delta_\mu^{dh, W_r}(F, F') = \Delta_\mu^{dh, W_\exists}(F, F')$ is exponential in the size of the formulae.*

Proof By Lemma 7, there exists formulae μ, F, F' on six variables X such that W_\exists is only equivalent to sets of weight vectors of cardinality greater than or equal to two. Since the variables are six, these three formulae are equivalent to formulae of size at most 2^6 , a constant.

This construction is replicated on m disjoint alphabets X_1, \dots, X_m of six variables each, giving m triples μ_i, F_i, F'_i of formulae with no shared variables among different triples and size bounded by a constant each. These formulae are conjoined: $\mu = \mu_1 \wedge \dots \wedge \mu_m, F = F_1 \wedge \dots \wedge F_m$ and $F' = F'_1 \wedge \dots \wedge F'_m$.

Since the triples are on different variables, the models of μ are combinations of models of each μ_i , the distance between a model and F is the sum of the distance from every F_i , and the same for F' . As a result, the models of μ at minimal weighted distance from F and F' are combinations of the models of each μ_i at minimal weighted distance from F_i and F'_i . For each triple, all sets of weight vectors W_r that make $\Delta_\mu^{dh, W_r}(F_i, F'_i)$ equal to $\Delta_{\mu_i}^{dh, W_\exists}(F_i, F'_i)$ contain at least two weight vectors by Lemma 7. The weight vectors that make $\Delta_\mu^{dh, W_r}(F, F')$ equal to $\Delta_\mu^{dh, W_\exists}(F, F')$ are their combinations, and are therefore exponentially many in m . □

Theorem 4 No partial preorder \leq depending on F_1 and F_2 only is such that $\Delta_{\mu}^{dh, W\exists}(F_1, F_2) = \min(\text{mod}(\mu), \leq)$.

Proof By Lemma 6, for every set of distance vectors there exists μ, F_1 and F_2 such that the models of μ have these Hamming distance vectors from F_1 and F_2 . The distance vectors that prove the claim are $[3, 0]$, $[2, 2]$ and $[0, 3]$. Their corresponding models of μ are denoted I, J and K .

Let μ' be the formula satisfied only by the models I and J , the ones at distance $[3, 0]$ and $[2, 2]$ from F_1 and F_2 . They are both minimal with weights $W = [2, 1]$. As a result, $I \not\prec J$. By symmetry, $K \not\prec J$. The ordering \leq is the same since by assumption it does not depend on μ but only on F_1 and F_2 , which are the same. A consequence of $I \not\prec J$ and $K \not\prec J$ is that J is a minimal model of μ . However, it is not in the result of merging with constraints μ as proved by Lemma 3. \square

Lemma 9 For every model-formula distance d and non-empty set of weight vectors $W.$, the merging operator $\Delta^{d, W.}$ satisfies postulates IC0, IC1, IC2 and IC7.

Proof The claim is proved one postulate at time.

IC0 $\Delta_{\mu}^{d, W.}(E) \subseteq \text{mod}(\mu)$

by definition, $\Delta_{\mu}^{d, W.}(E)$ is a subset of the models of μ ;

IC1 if μ is consistent, then $\Delta_{\mu}^{d, W.}(E)$ is not empty

by assumption, $W.$ contains at least a vector of weights W ; for this vector, $\Delta_{\mu}^{d, W.}(E)$ is the set of models of μ at minimal weighted distance from F_1, \dots, F_m ; if μ is consistent, it has at least a minimal model;

IC2 if $\wedge E$ is consistent with μ , then $\Delta_{\mu}^{d, W.}(E) = \text{mod}(\mu) \cap \text{mod}(\wedge E)$

since $d(I, F_i) = 0$ when $I \models F_i$, the distance vectors of the models of $\wedge E$ are $[0, \dots, 0]$; regardless of the weights, the weighted distance is zero, and therefore minimal; all other models have a strictly positive distance; since weights are strictly positive, their weighted distance is greater than zero;

IC7 $\text{mod}(\mu') \cap \Delta_{\mu}^{d, W.}(E) \subseteq \Delta_{\mu \wedge \mu'}^{d, W.}(E)$

the models in $\text{mod}(\mu') \cap \Delta_{\mu}^{d, W.}(E)$, if any, are the models that satisfy μ' , and also satisfy μ and no other model of μ has a lower distance from E weighted by some $W \in W.$; each such model satisfies $\mu \wedge \mu'$, and no other model of $\mu \wedge \mu'$ has lower distance weighted by W , since the models of $\mu \wedge \mu'$ are a subset of those of μ . \square

Lemma 10 If $W.$ contains every permutation of every vector it contains, then IC3 holds. For some set of weight vectors that does not include a permutation of one of its elements, IC3 does not hold.

Proof Postulate IC3 is: if $E_1 \equiv E_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}^{d, W.}(E_1) = \Delta_{\mu_2}^{d, W.}(E_2)$, where profiles are equivalent if there exists a bijection such that the associated formulae are equivalent.

Since the definition of merging only involves the set of models of μ_1 and not its syntax, the result is the same when switching to an equivalent formula μ_2 .

The same holds for the formulae: the result of merging does not change if a formula F_i is replaced by an equivalent one. This proves the claim when the bijection links

each formula of the first profile to the one of the same index of the second. This is generalized to arbitrary bijections by showing that the result of merging does not change when swapping the position of two arbitrary formulae.

Let $F_1, \dots, F_i, \dots, F_j, \dots, F_m$ and $F_1, \dots, F_j, \dots, F_i, \dots, F_m$ be the two profiles.

A model I is in the result of merging the first profile if there exists a weight vector $W = [w_1, \dots, w_i, \dots, w_j, \dots, w_m]$ in W . such that $W \cdot d(F_1, \dots, F_i, \dots, F_j, \dots, F_m)$ is minimal.

The weight vector $W' = [w_1, \dots, w_j, \dots, w_i, \dots, w_m]$ is obtained by swapping the weights of index i and j in W . Since W . contains the permutation of every weight vector it contains, and it contains W , it also contains W' .

The distance from every model I to the first profile according to W and to the second according to W' are the same:

$$\begin{aligned} &W \cdot d(I, F_1, \dots, F_i, \dots, F_j, \dots, F_m) \\ &= [w_1, \dots, w_i, \dots, w_j, \dots, w_m] \cdot d(I, F_1, \dots, F_i, \dots, F_j, \dots, F_m) \\ &= w_1d(I, F_1) + \dots + w_id(I, F_i) + \dots + w_jd(I, F_j) + \dots \\ &\quad + w_md(I, F_m) \\ &W' \cdot d(I, F_1, \dots, F_j, \dots, F_i, \dots, F_m) \\ &= [w_1, \dots, w_j, \dots, w_i, \dots, w_m] \cdot d(I, F_1, \dots, F_j, \dots, F_i, \dots, F_m) \\ &= w_1d(I, F_1) + \dots + w_jd(j, F_j) + \dots + w_id(I, F_i) + \dots \\ &\quad + w_md(I, F_m) \end{aligned}$$

As a result, the minimal models are also the same; therefore, the results of merging are also the same.

An example of a set of weight vectors that does not include the permutation of every vector it contains is $\{[1, 2]\}$. Merging the profile made of x and $\neg x$ produces $\neg x$ while merging the profile made of $\neg x$ and x produces x , using $\mu = true$ in both cases. □

Lemma 11 *If W . contains every permutation of every vector it contains and d satisfies the triangle inequality $\forall I, J, K. d(I, K) + d(K, J) \geq d(I, J)$, then IC4 holds. For some set of weight vectors that does not include a permutation of one of its elements IC4 does not hold. The same for some distance not satisfying the triangle inequality.*

Proof Postulate IC4 is: if $F_1 \models \mu$ and $F_2 \models \mu$ then $\Delta_{\mu}^{d, W} \cdot (F_1, F_2) \cap \text{mod}(F_1)$ is not empty if and only if $\Delta_{\mu}^{d, W} \cdot (F_1, F_2) \cap \text{mod}(F_2)$ is not empty.

This postulate does not hold in general. For example, the single weight vector $[2, 1]$ with the drastic or Hamming distance and $\mu = true$, $F_1 = x$ and $F_2 = \neg x$ would select the model of F_1 only. Both distances satisfy the triangle inequality.

This counterexample suggests that the postulate holds if the set W . has some sort of symmetry: if it contains a weight vector, it also contains all its permutations. This

is however not sufficient, as shown by the following counterexample:

$$\begin{aligned}
 ds(I, J) &= \begin{cases} 0 & \text{if } dh(I, J) = 0 \\ 1 & \text{if } dh(I, J) = 1 \\ 2 & \text{if } 2 \leq dh(I, J) \leq 4 \\ 5 & \text{if } 5 \leq dh(I, J) \end{cases} \\
 W &= \{[5, 2], [2, 5]\} \\
 F_1 &= x_1 \wedge x_2 \wedge x_3 \wedge x_4 \wedge x_5 \\
 F_2 &= \neg x_1 \wedge \neg x_2 \wedge \neg x_3 \wedge \neg x_4 \wedge \neg x_5 \\
 \mu &= F_1 \vee F_2 \vee (x_1 \wedge \neg x_2 \wedge \neg x_3 \wedge \neg x_4 \wedge \neg x_5)
 \end{aligned}$$

The distance ds may look unnatural, but has a rationale: instead of measuring the distance between models by the exact number of differing variables, it roughly approximates it by aggregating certain groups of consecutive values into one, so that only finitely many different distances exist.

The models of μ have distance vectors $[0, 5]$, $[2, 1]$ and $[5, 0]$. The first and the third are the models of F_1 and F_2 , respectively. The weight vector $[5, 2]$ turns these distance vectors into the weighted distances $[5, 2] \cdot [0, 5] = 10$, $[5, 2] \cdot [2, 1] = 12$, $[5, 2] \cdot [5, 0] = 25$; only the model of F_1 is minimal. For the weight vector $[2, 5]$: $[2, 5] \cdot [0, 5] = 25$, $[2, 5] \cdot [2, 1] = 9$, $[2, 5] \cdot [5, 0] = 10$; the only minimal model is the second, which is not a model of F_2 . This is a case in which both F_1 and F_2 imply μ and $\Delta_{\mu}^{ds, W} \cdot (F_1, F_2)$ contains some models of F_1 but none of F_2 .

Note that Lemma 6 does not apply to this case. It tells how to obtain certain distance vectors with formulae μ, F_1, F_2 , but these do not necessarily obey $F_1 \models \mu$ and $F_2 \models \mu$. To the contrary, the proof of the lemma involves formulae F_1 and F_2 that have models that falsify μ .

Postulate IC4 requires not only W . to be symmetric, but also d to satisfy the triangle inequality: for every three models I, J and K , it holds $d(I, K) + d(K, J) \geq d(I, J)$.

Since $\Delta_{\mu}^{d, W} \cdot (F_1, F_2) \cap \text{mod}(F_1)$ is not empty, there exists a weight vector $[a, b]$ and a model I of F_1 with distance vector $[0, c]$ such that $[a, b] \cdot [0, c]$ is minimal (the zero is because $I \models F_1$).

By definition, $d(I, F_2) = c$ implies $d(I, J) = c$ for some $J \in \text{mod}(F_2)$. This implies $d(J, F_1) \leq d(J, I) = c$; if $d(J, F_1) < c$ then $d(J, K) < c$ for some $K \in \text{mod}(F_1)$, which implies $d(K, F_2) < c = d(I, F_2)$, contradicting the assumption that I is minimal; therefore, $d(J, F_1) = c$.

Since J satisfies F_2 , it also satisfies μ . It is therefore a candidate for being in the result of merging. If $a < b$, then the weighted distance of J is $[a, b] \cdot [c, 0] = ac < bc = [a, b] \cdot [0, c]$. Since $[0, c]$ is the distance vector of I , this contradicts the assumption that I is minimal for weights $[a, b]$. This proves $a \geq b$.

Model J is now proved to have minimal distance weighted by $[b, a]$. The weighted distance of J is $[b, a] \cdot [c, 0] = bc$. Contrary to the claim, let K be a model with distance vector $[e, f]$ such that $[b, a] \cdot [e, f] < bc$.

The triangular property implies $e + f \geq c$. In details: $e + f < c$ implies the existence of two models I' and J' of respectively F_1 and F_2 such that $d(K, I') = e$, $d(K, J') = f$ and $d(I', J') \leq e + f < c$. This contradicts the assumption of

minimality of I . This property $e + f \geq c$, together with $a \geq b$, makes the following inequalities valid:

$$\begin{aligned}
 [b, a] \cdot [e, f] &= be + af \\
 &= be + bf + (a - b)f \\
 &= b(e + f) + (a - b)f \\
 &\geq bc + (a - b)f \\
 &\geq bc
 \end{aligned}$$

Contrary to what assumed, $[b, a] \cdot [e, f] \geq bc$. This proves that no such model K may exist, and that J has minimal distance weighted by $[b, a]$. Since $J \models F_2$, the intersection $\Delta_{\mu}^{d,W} \cdot (F_1, F_2) \cap \text{mod}(F_2)$ is proved not empty as required. \square

Lemma 12 *If $\Delta_{\mu}^{d,W'} \cdot (F_1, \dots, F_k) \cap \Delta_{\mu}^{d,W''} \cdot (F_{k+1}, \dots, F_m)$ is not empty, it coincides with $\Delta_{\mu}^{d,W} \cdot (F_1, \dots, F_m)$, where W is the Cartesian product of W' and W'' (postulates IC5 and IC6).*

Proof Let I be a model of both $\Delta_{\mu}^{d,W'} \cdot (F_1, \dots, F_k)$ and $\Delta_{\mu}^{d,W''} \cdot (F_{k+1}, \dots, F_m)$. By assumption, there exist $W' \in W'$ and $W'' \in W''$ such that the distance vector $d(I, F_1, \dots, F_k)$ weighted by W' is minimal among the models of μ , and the distance vector $d(I, F_{k+1}, \dots, F_m)$ weighted by W'' is minimal among the models of μ . This is equivalent to $d(I, F_1, \dots, F_k, F_{k+1}, \dots, F_m)$ being minimal when weighted by $W'W''$; this is the vector obtained by concatenating W' and W'' , and is therefore in W .

In the other way around, a model that is not minimal on its weighted distance to F_1, \dots, F_k or to F_{k+1}, \dots, F_m is not minimal on its weighted distance to $F_1, \dots, F_k, F_{k+1}, \dots, F_m$. \square

Theorem 5 *There exist μ, μ', F_1 and F_2 such that $\text{mod}(\mu') \cap \Delta_{\mu}^{dh,W\exists} \cdot (F_1, F_2)$ is not empty but $\Delta_{\mu \wedge \mu'}^{dh,W\exists} \cdot (F_1, F_2) \not\subseteq \Delta_{\mu}^{dh,W\exists} \cdot (F_1, F_2)$.*

Proof Let μ, F_1 and F_2 be such that μ has three models with distance vectors $d(I, F_1, F_2) = [1, 0]$, $d(J, F_1, F_2) = [0, 1]$ and $d(K, F_1, F_2) = [0, 2]$. Such formulae exist thanks to Lemma 6.

The models of $\Delta_{\mu}^{dh,W\exists} \cdot (F_1, F_2)$ are I and J , since these two models have minimal distance weighted by $[1, 1]$. Since J dominates K , by Lemma 1 K is not in the result of merging for any weights.

Let μ' be the formula with models I and K . Since I is also in $\Delta_{\mu}^{dh,W\exists} \cdot (F_1, F_2)$, this set contains a model of μ' , as required. When merging under constraints $\mu \wedge \mu'$, model K is minimal with weights $[2, 1]$, since it and the other model I of $\mu \wedge \mu'$ have both weighted distance 2. \square

Theorem 6 *There exists F_1, F_2 such that $\Delta_{true}^{d,W\exists} \cdot (F_1, F_2, \dots, F_2) \not\subseteq \text{mod}(F_2)$, where F_2 is repeated an arbitrary number of times.*

Proof The formulae are $F_1 = a$ and $F_2 = \neg a$. For every number of repetitions n , there exists W such that $\Delta_\mu^{d,\{W\}}(F_1, F_2, \dots, F_2)$ contains the model $\{a\}$, which does not satisfy F_2 . In particular, the weight vector is $W = [n, 1, \dots, 1]$. The weighted distance of $\{a\}$ from the formulae is n , the same as the weighted distance of the only other model $\{\neg a\}$. As a result, $\{a\}$ is minimal. \square

Theorem 7 *If all weights in the vectors in W . are lower than a constant and d is an arbitrary distance, for every μ and $F_1, \dots, F_o, F_{o+1}, \dots, F_m$, there exists n such that $\Delta_\mu^{d,W}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m) \subseteq \Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$, where the formulae from F_{o+1} to F_m are repeated n times.*

Proof By definition, $\Delta_\mu^{d,W}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m)$ is the union of $\Delta_\mu^{d,W}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m)$ for all weight vectors $W \in W$. The same holds for $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$. The claim is proved by showing a number n that makes $\Delta_\mu^{d,W}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m)$ contained in $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$ for every $W \in W$.

The set $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$ comprises all models of μ at minimal weighted distance from F_{o+1}, \dots, F_m . This minimal weighted distance is denoted by b . The minimal weighted distance from these models to F_1, \dots, F_o is denoted by a ; this is the minimal weighted distance to F_1, \dots, F_o from the models of $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$ only, not from all models of μ .

By definition, the weighted distance from a model to $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$ is the same as its weighted distance to F_1, \dots, F_o plus its weighted distance to $F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$, which is n times its weighted distance from F_{o+1}, \dots, F_m . As a result, some models of $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$ are at weighted distance $a + nb$ from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$.

The set $\Delta_\mu^{d,W}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m)$ comprises all models of μ at minimal weighted distance from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$, which is equal to the weighted distance from F_1, \dots, F_o plus n times the weighted distance from F_{o+1}, \dots, F_m .

If one of these models is not in $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$, its weighted distance from F_{o+1}, \dots, F_m is at least $b + 1$, since b is the minimal weighted distance from the models of μ to F_{o+1}, \dots, F_m . Therefore, its weighted distance from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$ is at least $n(b + 1)$.

Since this model is at minimal weighted distance from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$, all other models of μ are at least at the same weighted distance from these formulae. This includes the models that are the closest to F_{o+1}, \dots, F_m , which have been proved to be at weighed distance $a + nb$ from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$. The inequality $n(b + 1) \leq a + nb$ follows.

This is the same as $nb + n \leq a + nb$. Removing the common addends from both sides results in $n \leq a$. In summary, if a model of μ at minimal weighted distance from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$ is not at minimal weighted distance from F_{o+1}, \dots, F_m then $n \leq a$. In the other way around, if $n > a$ then all models of μ at minimal weighted distance from $F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m$ are also at minimal weighted distance from F_{o+1}, \dots, F_m .

The conclusion is that $\Delta_\mu^{d,W}(F_1, \dots, F_o, F_{o+1}, \dots, F_{o+1}, \dots, F_m, \dots, F_m)$ is a subset of $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$ if $n > a$, where a is the minimal weighted distance of a model of $\Delta_\mu^{d,W}(F_{o+1}, \dots, F_m)$ from F_1, \dots, F_o . Since the weights in W are bounded by a constant, only a finite number of weight vectors are in W . Therefore, the maximal value of a across all $W \in W$ is finite. Every n larger than it satisfies the claim. \square

Lemma 13 For every μ, F_1, \dots, F_m it holds:

$$\Delta_\mu^{d,W\exists}(F_1, \dots, F_m) = \Delta_\mu^{d,W\exists}(F_1, \dots, F_m, F_m)$$

Proof By definition, $\Delta_\mu^{d,W\exists}(F_1, \dots, F_m)$ is the union of $\Delta_\mu^{d,W}(F_1, \dots, F_m)$ for every $W \in W\exists$, and the same when F_m is duplicated. The claim is proved by showing that for each $W = [w_1, \dots, w_{m-1}, w_m]$ there exists $W' = [w'_1, \dots, w'_{m-1}, w'_m, w'_m]$ such that $\Delta_\mu^{d,W}(F_1, \dots, F_m)$ is equal to $\Delta_\mu^{d,W'}(F_1, \dots, F_m, F_m)$, and vice versa.

The distance of a model from F_1, \dots, F_m weighted by $W = [w_1, \dots, w_{m-1}, w_m]$ is exactly half of the distance of the same model from F_1, \dots, F_m, F_m weighted by $W' = [2w_1, \dots, 2w_{m-1}, w_m, w_m]$, since each distance is multiplied by two. Therefore, the minimal models are the same.

Vice versa, the distance of a model from F_1, \dots, F_m, F_m weighted by $W' = [w_1, \dots, w_{m-1}, w_m, w'_m]$ is exactly the same as the distance of the same model from F_1, \dots, F_m weighted by $W = [w_1, \dots, w_{m-1}, w_m + w'_m]$. In this case, the weighted distances are exactly the same, and the minimal models coincide. \square

Theorem 8 For every pair of satisfiable formulae F_1 and F_2 over an alphabet of n variables, it holds $F_1 \Delta_D F_2 = \Delta_{true}^{dh,W_{n+1}}(F_1, F_2)$.

Proof By definition, $I \in F_1 \Delta_D F_2$ if and only if $I \models F_1$ and there exists $J \models F_2$ such that $\langle I, J \rangle$ is minimal according to \leq_{dh} , or the same with F_1 and F_2 swapped. What is now proved is that the first condition is equivalent to $I \in \Delta_{true}^{dh,[n+1,1]}(F_1, F_2)$. By symmetry, the condition with the two formulae swapped is equivalent to $I \in \Delta_{true}^{dh,[1,n+1]}(F_1, F_2)$.

The relevant cases are: $I \models F_1$ and $\langle I, J \rangle$ is minimal for some $J \models F_2$, $I \models F_1$ and $\langle I, J \rangle$ is minimal for no $J \models F_2$, and $I \not\models F_1$. The claim holds if $I \in \Delta_{true}^{dh,W_{n+1}}(F_1, F_2)$ holds exactly in the first case.

- $I \models F_1$ and $\langle I, J \rangle$ is minimal for some $J \models F_2$; since $I \models F_1$, the distance from I to F_1 is zero: $dh(I, F_1) = 0$; therefore, the weighted distance from I to the formulae is $[n + 1, 1] \cdot [0, dh(I, F_2)] = dh(I, F_2)$, which is at most n ; the negation of the claim is that the weighted distance $(n + 1)dh(K, F_1) + 1dh(K, F_2)$ of some other model K is less than that; for it being less than n implies $dh(K, F_1) = 0$; as a result, the weighted distance of K is $dh(K, F_2)$; if it were less than the weighted distance of I then $dh(K, F_2) < dh(I, F_2)$; by definition, this means that there exists K' such that $dh(K, K')$ is less than $dh(I, I')$ for every $I' \models F_2$, including $I' = J$; this implies $dh(K, K') < dh(I, J)$, contrary to the assumption that $\langle I, J \rangle$ is minimal;

2. $I \models F_1$ and $\langle I, J \rangle$ is minimal for no $J \models F_2$; by assumption, there exists K, K' such that $K \models F_1, K' \models F_2$ and $dh(K, K') < dh(I, J)$ for every $J \models F_2$; this implies $dh(K, F_2) < dh(I, F_2)$; since both I and K satisfy F_1 , it also holds $dh(I, F_1) = dh(K, F_1) = 0$; as a result, the weighted distances of these models are $[n+1, 1] \cdot [0, dh(I, F_2)] = dh(I, F_2)$ and $[n+1, 1] \cdot [0, dh(K, F_2)] = dh(K, F_2)$; since $dh(K, F_2) < dh(I, F_2)$, the model I is not at a minimal weighted distance;
3. $I \not\models F_1$; since F_1 is by assumption satisfiable, it has a model K ; since $dh(K, F_1) = 0$, the weighted distance for this model is $[n+1, 1] \cdot [dh(K, F_1), dh(K, F_2)] = dh(K, F_2)$, which is at most n ; the weighted distance of I is instead $[n+1, 1] \cdot [dh(I, F_1), dh(K, F_2)] = (n+1)dh(I, F_1) + dh(I, F_2)$, which is greater than n since $dh(I, F_1) > 0$.

Since I has minimal weighted distance from F_1 and F_2 in the first case but not in the second and in the third, the claim is proved. □

Theorem 9 For every distance d bounded by k , if F_1, \dots, F_m are satisfiable then $\Delta_{true}^{d, W_{km}}(F_1, \dots, F_m)$ is a disjunctive merging operator.

Proof Let I be a model satisfying no formula F_i . The disjunctive property holds if I is not in $\Delta_{true}^{d, W_{km}}(F_1, \dots, F_m)$. This holds if I is not in $\Delta_{true}^{d, W}(F_1, \dots, F_m)$ for any $W \in W_{km}$.

By assumption, I does not satisfy any of the formulae. Therefore, its distance vector is greater than or equal to $[1, \dots, 1]$. Multiplying $[1, \dots, 1]$ by W results in $km + (m - 1)$.

Since W is in W_{km} , one of its elements is km . Let i be its index. Since F_i is satisfiable, it has a model J . The distance vector of J is at most $[k, \dots, k, 0, k, \dots, k]$ where 0 is at index i . The result of multiplying it by W is $(m - 1)k$.

The upper bound for the weighted distance of J is $(m - 1)k$, which is less than $km + (m - 1)$, the lower bound of the weighted distance of minimal I . This proves that I is not minimal. □

Theorem 10 For some distance d with codomain of size three, there exists I, μ and F_1, \dots, F_m such that $I \models \mu$ and $I \notin \Delta_{\mu}^{d, W_{\exists}, \Sigma^2}(F_1, \dots, F_m)$ hold, but $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ does not hold for any $J \models \mu$.

Proof The proof is based on three models of distance vectors $[4, 0], [3, 3]$ and $[0, 4]$. Let $w = [w_1, w_2]$ be an arbitrary weight vector. The squared distance vectors are $[16, 0], [9, 9]$ and $[0, 16]$. Multiplying them by the weight vector $[w_1, w_2]$ gives $16w_1, 9w_1 + 9w_2$ and $16w_2$.

In order for the second model to be generated by merging, the second number has to be lower than or equal to both the first and the third.

$$\begin{aligned} 9w_1 + 9w_2 &\leq 16w_1 \\ 9w_1 + 9w_2 &\leq 16w_2 \end{aligned}$$

These inequalities are the same as $9w_2 \leq 7w_1$ and $9w_1 \leq 7w_2$, or $w_2 \leq \frac{7}{9}w_1$ and $w_1 \leq \frac{7}{9}w_2$. These imply $w_2 < w_1$ and $w_1 < w_2$, which are impossible together. □

Lemma 14 *If $v = [v_1, \dots, v_m]$ and $u = [u_1, \dots, u_m]$ are two vectors of integers such that $v_i \leq u_i$ holds for every index i , the same holds for the result of sorting v and u in descending order.*

Proof Two vectors can be sorted by Bubblesort, which compares and possibly swaps pairs of consecutive elements (Astrachan 2003). Running the algorithm in parallel on the two vectors iterates over the same basic step: if v_i is less than v_{i+1} , these two elements are swapped; the same for u_i and u_{i+1} .

If none of the two pairs is swapped, the vectors do not change; therefore, each element of v is still less than or equal to the corresponding element of u . If both pairs are swapped, the condition still holds because the element corresponding to v_i is still u_i and that corresponding to v_{i+1} is still u_{i+1} .

The same is proved when only one of the two pairs is swapped. The swap is assumed done on v and not on u ; the converse case is symmetric. The result of swapping is the following.

$$\begin{aligned} & [v_1, \dots, v_{i+1}, v_i, \dots, v_m] \\ & [u_1, \dots, u_i, u_{i+1}, \dots, u_m] \end{aligned}$$

Since the two elements of v are swapped, they were not in the requested order: $v_i < v_{i+1}$ holds. Since the two elements of u are not swapped, they are already in the requested order: $u_i \geq u_{i+1}$. Before the swap, each element of v was less than or equal to the corresponding elements of u before the swap: $v_i \leq u_i$ and $v_{i+1} \leq u_{i+1}$.

The claim is the same after the swap: $v_{i+1} \leq u_i$ and $v_i \leq u_{i+1}$. The first is a consequence of $v_{i+1} \leq u_{i+1}$ and $u_{i+1} \leq u_i$. The second is a consequence of $v_i < v_{i+1}$ and $v_{i+1} \leq u_{i+1}$.

The conclusion is that the basic step of Bubblesort keeps each element of v less than or equal to the corresponding elements of u . This condition holds for the two ordered vectors since they result from iterating this step. □

Lemma 15 *If $v = [v_1, \dots, v_m]$ and $u = [u_1, \dots, u_m]$ are two vectors of integers such that $v_i \leq u_i$ holds for every index i and $v_i < u_i$ for some index i , the same holds for the result of sorting v and u in descending order.*

Proof Lemma 14 proves that each element of the first sorted vector is less than or equal to the corresponding one of the second. The claim requires the comparison to be strict for at least one element.

Proof is by contradiction. If the ordering is not strict, the two sorted vectors are the same. This implies that their sum is the same. Since sorting only changes the order among the elements of the vectors, this is also the case for the two unsorted vectors. Their sum cannot be the same since $v_i \leq u_i$ holds for all elements and $v_i < u_i$ for at least one. □

Lemma 16 *For every distance d , vector of weights $W \in W_{\exists}$ and model I , if $I \in \Delta_{\mu}^{d,W,leximax}(F_1, \dots, F_m)$ then $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ holds for no model J of μ .*

Proof The claim is proved by showing that $d(J, F_1, \dots, F_m) < d(I, F_1, \dots, F_m)$ forbids I from being in $\Delta_{\mu}^{d,W,leximax}(F_1, \dots, F_m)$ for any $W \in W_{\exists}$.

Let $W = [w_1, \dots, w_m]$ be an arbitrary element of W_{\exists} . Since all weights are greater than zero, $d(J, F_i) \leq d(I, F_i)$ is the same as $w_i d(J, F_i) \leq w_i d(I, F_i)$. As a result, multiplying the elements of $d(J, F_1, \dots, F_m)$ and $d(I, F_1, \dots, F_m)$ by their respective weights in $[w_1, \dots, w_m]$ results in two vectors V and U such that $V < U$.

Let V' be the result of sorting V in descending order and U' the result of sorting U . Lemma 15 proves that $V < U$ implies $V' < U'$. Every element of V' is less than or equal to the corresponding element of U' , and one is strictly so. As a result, V' is less than U' in the lexicographic order. As a result, I is not minimal, and is therefore not in $\Delta_{\mu}^{d,W,leximax}(F_1, \dots, F_m)$. This holds for every set of weights $W \in W_{\exists}$. \square

Lemma 17 *If I is a model of μ and $d(I, F_1, \dots, F_m)$ is not strictly dominated by the vector of distances of any other model of μ , then there exists W such that $I \in \Delta_{\mu}^{d,W,leximax}(F_1, \dots, F_m)$.*

Proof Since I is not strictly dominated by any other model of μ , for every other model J of μ two cases are possible: either $I \leq J$ or there exists i such that $d(I, F_i) < d(J, F_i)$. In the first case, I is always less than or equal to J according to the leximax ordering regardless of the weights thanks to Lemma 14. For the models J of the second kind, a vector of weights W making I less than all of them is shown.

Each distance $d(I, F_i)$ such that $d(I, F_i) < d(J, F_i)$ for at least one such model J may be zero or greater than zero. Let the ones greater than zero be x, y and z . The other distances are not important, except that their maximum value plus one is denoted v .

The distance vector of I therefore comprises three kinds of elements: the ones such that $d(I, F_i) < d(J, F_i)$ is not the case for any J , the ones such that $d(I, F_i) = 0$ and $d(I, F_i) < d(J, F_i)$ holds for some model J , and the ones such that $d(I, F_i) < d(J, F_i)$ holds for some J and $d(I, F_i)$ is either x, y , or z .

Regarding the distance vectors of the models J , all that is known is that $d(J, F_i)$ is strictly greater than $d(I, F_i)$ for some index i . The remaining elements are unknown, but they are not necessary anyway.

The distance vectors can be rearranged as follows.

I	[$< v$...	$< v$	0	...	0	x	y	z]
J	[≥ 1]
J'	[≥ 1]
J''	[$> x$]
...	[$> y$]
...	[$> z$]
W	[1	...	1	$2vxyz$...	$2vxyz$	vyz	vxz	uxy]

The last line of the table is a weight vector. Multiplying the distances each by its weight produces the following table, where $u = vxyz$.

I	[$< v$	\dots	$< v$	0	\dots	0	u	u	u]
J	[$\geq 2u$]
J'	[$\geq 2u$]
J''	[$> u$]
\dots	[$> u$]
\dots	[$> u$]

Since v is one plus the maximum of some nonnegative numbers, it is larger than zero. Since x, y and z are also larger than zero, $u = vxyz$ is larger than v . As a result, the maximum element of the vector of I is u .

Every other vector contains at least an element strictly greater than u . Its maximum element is therefore strictly greater than u . It is therefore strictly greater than the vector of I according to the lexicmax ordering. \square

Theorem 11 *For every consistent formulae μ, F_1, \dots, F_m , the following equality holds:*

$$\Delta_{\mu}^{d, W_{\exists}, \text{leximax}}(F_1, \dots, F_m) = \bigcup_{S \in \text{maxcon}_{\mu}(F_1, \dots, F_m)} \text{mod}(\wedge S)$$

Proof Lemma 16 proves that lexicmax merging does not select models of μ dominated by others. Lemma 17 proves that lexicmax merging selects all models of μ that are not dominated by others. Overall, lexicmax merging selects exactly the models of μ that are not dominated by others.

This is the same selection made by merging with the drastic distance and the sum as the aggregation function. As a result, $\Delta_{\mu}^{d, W_{\exists}, \text{leximax}}(F_1, \dots, F_m)$ is the same as $\Delta_{\mu}^{dd, W_{\exists}}(F_1, \dots, F_m)$. Theorem 3 proves that the latter is the same as $\bigcup_{S \in \text{maxcon}_{\mu}(F_1, \dots, F_m)} \text{mod}(\wedge S)$. \square

Theorem 12 *For every consistent formulae μ, F_1, \dots, F_m , the following equality holds*

$$\Delta_{\mu}^{dd, W_{\exists}, \text{leximin}}(F_1, \dots, F_m) = \bigcup_{S \in \text{cardcons}_{\mu}(F_1, \dots, F_m)} \text{mod}(\wedge S)$$

where

$$\begin{aligned} \text{cardcons}_{\mu}(F_1, \dots, F_m) \\ = \{S \in \text{maxcon}_{\mu}(F_1, \dots, F_m) \mid \nexists S' \in \text{maxcon}_{\mu}(F_1, \dots, F_m) \cdot |S| < |S'|\} \end{aligned}$$

Proof If a model of μ is dominated by another, the same ordering weathers multiplying their weight distances by the weight vector, sorting them in descending order and inverting the order of their elements. As a result, the dominated model is not minimal according to leximin. Leximin merging only produces undominated models. Not all of them, however.

The drastic distance $dd(I, F_i)$ is 0 if $I \models F_i$ and 1 if $I \not\models F_i$. Multiplying a distance vector $dd(I, F_1, \dots, F_m)$ by the weights and sorting the result in ascending order produces $[0, \dots, 0, w, w', \dots]$, where the number of zeros is the number of formulae satisfied by I . All following weights are strictly greater than zero because they are weights.

If another model J of μ satisfies more formulae F_1, \dots, F_m than I , its weighted and sorted distance vector is $[0, \dots, 0, w'', w''', \dots]$, where the number of zeros is larger than that of the vector of I . As a result, it contains a zero where the vector of I contains w . It is strictly smaller than that according to the leximin ordering. This implies that leximin merging does not generate I .

Leximin merging instead generates I if no other model μ satisfies more formulae than I . This is proved by the weight vector $[1, \dots, 1]$. The weighted and sorted distance vector of I is $[0, \dots, 0, 1, \dots, 1]$. That of any other model of μ is a vector $[0, \dots, 0, 1, \dots, 1]$ with a larger or equal number of zeros. As a result, I is minimal according to the leximin ordering. Leximin merging generates it. \square

References

- Ammoura, M., Raddaoui, B., Salhi, Y., & Oukacha, B. (2015). On measuring inconsistency using maximal consistent sets. In S. Destercke & T. Denoëux (Eds.), *Symbolic and quantitative approaches to reasoning with uncertainty* (pp. 267–276). Springer.
- Astrachan, O. (2003). Bubble sort: an archaeological algorithmic analysis. In *Proceedings of the 34th SIGCSE technical symposium on computer science education, SIGCSE 2003*, (pp. 1–5). ACM.
- Baral, C., Kraus, S., Minker, J., & Subrahmanian, V. (1992). Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, 8(1), 45–71.
- Benferhat, S., Dubois, D., & Prade, H. (1997). Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study part 1: The flat case. *Studia Logica*, 58(1), 17–45.
- Benferhat, S., Lagrue, S., & Rossit, J. (2014). Sum-based weighted belief base merging: From commensurable to incommensurable framework. *Journal of Automated Reasoning*, 55(9), 2083–2108.
- Bertrand, J. (1889). *Calcul des probabilités*. Gauthier-Villars.
- Booth, R., & Hunter, A. (2018). Trust as a precursor to belief revision. *Journal of Artificial Intelligence Research*, 61, 699–722.
- Brewka, G. (1989). Preferred subtheories: an extended logical framework for default reasoning. In *Proceedings of the eleventh international joint conference on artificial intelligence (IJCAI'89)* (pp. 1043–1048).
- Chacón, J., & Pino Pérez, R. (2006). Merging operators: Beyond the finite case. *Information Fusion*, 7(1), 41–60.
- Cholvy, L. (1998). Reasoning about data provided by federated deductive databases. *Journal of Intelligent Information Systems*, 10(1), 49–80.
- Chopra, S., Ghose, A., & Meyer, T. (2006). Social choice theory, belief merging, and strategy-proofness. *Information Fusion*, 7(1), 61–79.
- Clavin, T. (2013). There never was such a thing as a red phone in the White House. *Smithsonian Magazine*.
- Cojan, J., & Lieber, J. (2012). Belief revision-based case-based reasoning. In *Proceedings of the ECAI-2012 workshop SAMAI: Similarity and analogy-based methods in AI* (pp. 33–39).
- Dalal, M. (1988). Investigations into a theory of knowledge base revision: Preliminary report. In *Proceedings of the seventh national conference on artificial intelligence (AAAI'88)* (pp. 475–479).
- Darwiche, A., & Marquis, P. (2004). Compiling propositional weighted bases. *Artificial Intelligence*, 157(1), 81–113.
- Darwiche, A., & Pearl, J. (1997). On the logic of iterated belief revision. *Artificial Intelligence Journal*, 89(1–2), 1–29.

- Dubois, D., Liu, W., Ma, J., & Prade, H. (2016). The basic principles of uncertain information fusion. An organised review of merging rules in different representation frameworks. *Information Fusion*, 32, 12–39.
- Dunan, M. (1963). La taille de Napoléon. *Revue de l'Institut Napoléon* (pp. 178–179).
- Epstein, E. (1982). Have you ever tried to sell a diamond? *The Atlantic*.
- Everaere, P., Konieczny, S., & Marquis, P. (2010). Disjunctive merging: Quota and Gmin merging operators. *Artificial Intelligence Journal*, 174(12–13), 824–849.
- Everaere, P., Konieczny, S., & Marquis, P. (2010b). The epistemic view of belief merging: Can we track the truth? In *Proceedings of the nineteenth European conference on artificial intelligence (ECAI 2010)* (pp. 621–626). IOS Press.
- Everaere, P., Konieczny, S., & Marquis, P. (2020). Belief merging operators as maximum likelihood estimators. In *Proceedings of the twenty-ninth international joint conference on artificial intelligence (IJCAI 2020)*.
- Fagin, R., Ullman, J. D., & Vardi, M. Y. (1983). On the semantics of updates in databases. In *Proceedings of the second ACM SIGACT SIGMOD symposium on principles of database systems (PODS'83)* (pp. 352–365).
- Gärdenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk-taking, and decision making. *Synthese*, 53, 361–386.
- Gardikiotis, A. (2011). Minority influence. *Social and personality psychology compass*, 5(9), 679–693.
- Giagkiozis, I., & Fleming, P. (2014). Pareto front estimation for decision making. *Evolutionary Computation*, 22(4), 651–678.
- Ginsberg, M. L. (1986). Conterfactuals. *Artificial Intelligence*, 30, 35–79.
- Grant, J., & Hunter, A. (2011). Measuring the good and the bad in inconsistent information. In *Proceedings of the twenty-second international joint conference on artificial intelligence (IJCAI 2011)* (p. 2632).
- Hájek, A. (2012). Interpretations of probability. In E. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.
- Halpern, J., & Tuttle, M. (1993). Knowledge, probability, and adversaries. *Journal of the ACM*, 40(4), 917–960.
- Haret, A., Lackner, M. P. A., & Wallner, J. (2020). Proportional belief merging. In *Proceedings of the thirty-fourth AAAI conference on artificial intelligence (AAAI 2020)* (pp. 2822–2829).
- Katsuno, H., & Mendelzon, A. O. (1991). Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52, 263–294.
- Keynes, J. (1921). *A treatise on probability*. Macmillan and Company.
- Konieczny, S. (2000). On the difference between merging knowledge bases and combining them. In *Proceedings of the seventh international conference on principles of knowledge representation and reasoning (KR 2000)* (pp. 135–144).
- Konieczny, S. (2004). Belief base merging as a game. *Journal of Applied Non-Classical Logics*, 14(3), 275–294.
- Konieczny, S., Lang, J., & Marquis, P. (2004). DA^2 merging operators. *Artificial Intelligence*, 157(1–2), 49–79.
- Konieczny, S., & Pino Pérez, R. (1998). On the logic of merging. In *Proceedings of the sixth international conference on principles of knowledge representation and reasoning (KR'98)* (pp. 488–498).
- Konieczny, S., & Pino Pérez, R. (2002). Merging information under constraints: A logical framework. *Journal of Logic and Computation*, 12(5), 773.
- Konieczny, S., & Pino Pérez, R. (2002b). On the frontier between arbitration and majority. In *Proceedings of the eighth international conference on principles of knowledge representation and reasoning (KR 2002)* (pp. 109–120). Morgan Kaufmann.
- Konieczny, S., & Pino Pérez, R. (2011). Logic based merging. *Journal of Philosophical Logic*, 40(2), 239–270.
- Lang, J. (2004). Logical preference representation and combinatorial vote. *Annals of Mathematics and Artificial Intelligence*, 42(1), 37–71.
- Liberatore, P. (2015). Belief merging by examples. *ACM Transactions on Computational Logic*, 17(2):9:1–9:38.
- Liberatore, P., & Schaerf, M. (1998). Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering*, 10(1), 76–90.
- Lin, J. (1996). Integration of weighted knowledge bases. *Artificial Intelligence*, 83(2), 363–378.
- Lin, J., & Mendelzon, A. (1999). *Knowledge base merging by majority* (pp. 195–218). Springer.

- List, C. (2013). Social choice theory. In E. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University.
- Mata Díaz, A., & Pino Pérez, R. (2017). Impossibility in belief merging. *Artificial Intelligence*, 251, 1–34.
- Mazzoni, G. (2003). *Si può credere a un testimone? Il Mulino*.
- Meyer, T. (2001). On the semantics of combination operations. *Journal of Applied Non-Classical Logics*, 11(1–2), 59–84.
- Nishimura, K., & Ozaki, H. (2007). Irreversible investment and Knightian uncertainty. *Journal of Economic Theory*, 136(1), 668–694.
- Popper, K. (1959). *The logic of scientific discovery*. Routledge.
- Rescher, N., & Manor, R. (1970). On inference from inconsistent premisses. *Theory and Decision*, 1(2), 179–217.
- Revesz, P. (1997). On the semantics of arbitration. *International Journal of Algebra and Computation*, 7, 133–160.
- Revesz, P. Z. (1993). On the semantics of theory change: Arbitration between old and new information. In *Proceedings of the twelfth ACM SIGACT SIGMOD SIGART symposium on principles of database systems (PODS'93)* (pp. 71–82).
- Rott, H. (2006). Shifting priorities: Simple representations for twenty-seven iterated theory change operators. *Modality matters: Twenty-five essays in honour of Krister Segerberg, number 53 in uppsala philosophical studies* (pp. 359–384). Department of Philosophy.
- Shackel, N. (2007). Bertrand's paradox and the principle of indifference. *Philosophy of Science*, 74(2), 150–175.
- Wikipedia. (2017a). Adam ruins everything. <https://en.wikipedia.org/w/index.php?title=Adam+Ruins+Everything>.
- Wikipedia. (2017b). List of common misconceptions. <https://en.wikipedia.org/w/index.php?title=List+of+common+misconceptions>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.