Simon D'Alfonso

Belief Merging with the Aim of Truthlikeness

dalfonso@unimelb.edu.au, +61401826787

Noname manuscript No.
(will be inserted by the editor)

Belief Merging with the Aim of Truthlikeness

First Author

Received: date / Accepted: date

Abstract The merging/fusion of belief/data collections in propositional logic form is a topic that has received due attention within the domains of database and AI research. A distinction can be made between two types of scenarios to which the process of merging can be applied. In the first type, the collections represent preferences, such as the voting choices of a group of people, that need to be aggregated so as to give a consistent result that in some way best represents the collective judgement of the group. In the second type, the collections represent factual data that is to be aggregated with an aim of obtaining a result that maximises factual correctness. After introducing a general framework for belief merging via some prominent literature on the topic, this paper then introduces and considers a method for belief merging with the second type of scenario in mind. Its suitability is corroborated by demonstrating how it can be seen as a special case of a merging procedure that combines aggregation of probabilities and maximisation of expected truthlikeness.

Keywords Belief Merging · Information Fusion · Truthlikeness · Judgment Aggregation

1 Introduction

Propositional belief merging (logic-based information fusion) concerns the process of aggregating possibly conflicting propositional belief bases/logical databases into one consistent result.[1] Over the last few decades it is a topic that has received significant attention within the domains of database and AI research and is also relevant to the topics of judgement aggregation and social choice theory [5].

A fair portion of this work focuses on merging as a merging of group preferences, with the aim of obtaining a synthesis that best represents the group's

author

[1] See [7] for a good overview.

preferences according to some criteria. For example, suppose that a group of four people, Tom, Richard, Harry and Bob, are deciding what they should do for the night. They can do anywhere between none or all of the following actions: go out to dinner (A), watch a movie (B) and go to see some live music (C). Tom and Richard want to do A and B and are happy to go either way for C. Harry wants to do A, does not want to do C and is happy to go either way for B. Bob prefers to do none of these things. Out of the nine possible scenarios arising from these three choices, which should be selected to best satisfy the preferences of this group?

This paper on the other hand has an epistemic focus and is interested in the merging of factual data from multiple sources with an aim of obtaining results that best approximate the true state of the world. For example, suppose that Tom, Richard and Harry are given a true/false questionnaire with three questions (Q_1, Q_2, Q_3). In actuality, the correct answers are as follows: (Q_1, false), (Q_2, true) and (Q_3, true). Tom answers (Q_1, false), (Q_2, true) and (Q_3, true). Richard answers (Q_1, false), (Q_2, false) and does not answer for Q_3. Harry answers (Q_1, false), (Q_2, true) and (Q_3, false). How should the group's responses be merged with the aim of obtaining an optimal result; if not the completely correct trio of answers (Q_1, false), (Q_2, true) and (Q_3, true), then at least a trio with as many correct answers as possible given the group's input?

Whilst there has been some work on investigating this epistemic perspective of merging, it has tended to be in terms of epistemic social choice theory, involving probabilistic analyses of truth-tracking (arriving at the correct answer) given the reliability of belief sources and concerned with results such as the Condorcet Jury Theorem [18,6]. This paper approaches the matter of belief merging and truth conduciveness via the idea of truthlikeness [17], a concept from the philosophy of science that concerns measuring the extent to which a theory approaches the truth.

The relevance of truthlikeness to science as a truth-seeking enterprise and the collaborative nature of science motivate investigation into such a merging operator. For example, members of a team of individuals investigating a certain scientific problem might each conduct some experiments and generate some results pertaining to some hypothesis. When the time comes to combine these individual results a merging operator conducive to truthlikeness is apt. Such an operator could also be applied to any non-scientific scenarios where truth matters; for example, combining eyewitness accounts pertaining to some event. Beyond simply dividing possible merge results into the one model that represents the true state of the world and others that do not, a truthlikeness approach to merging will assess all possible merge results in terms of their degree of closeness to the true state of the world.

Thus in this paper we are interested in investigating the following question: given the aim to obtain results that maximise truthlikeness, what type of belief merging operator should be employed? Furthermore, what are the properties of such an operator and how might it differ from operators already proposed in the literature? We begin in Section 2 by giving a formal frame-

work for belief merging and some basic belief merging operators that have been established in the literature, before introducing the new operator proposed in this paper and investigating some of its formal properties. In Section 3 the notions of truthlikeness and estimated truthlikeness are explained and a formal truthlikeness measure is adopted. In Section 4 we look at the new merging operator in terms of its truthlikeness conduciveness and demonstrate how it can be seen as a special case of a procedure that combines probabilities and maximises estimated truthlikeness.

2 A formal framework for belief merging

Following the framework given in [9,10], some belief merging operators will now be presented and their basic properties discussed. We start off by establishing some preliminary definitions:

- We are working with a standard propositional logic language that has a finite set of atomic propositions $\mathcal{P}$. An interpretation assigns a value from $\{0,1\}$ to each atom $p \in \mathcal{P}$. The set of all interpretations/possible states[2] for a given logical space is denoted by $\mathcal{W}$.
- We will be using a model-based approach and represent an agent's belief base K semantically as the set of its models (mod); that is, the interpretations which satisfy it.
- A belief profile $E = \{K_1, ..., K_n\}$ is a multiset[3] of belief bases.
- A merge operator $\Delta(E)$ takes a belief profile E as input and returns a set of interpretations.[4]

Let $\leq_E$ denote a binary relation on $\mathcal{W}$. The set of interpretations resulting from the merge is defined as:

$$\Delta(E) = \min(\text{mod}(), \leq_E)$$

where $\min(\text{mod}(), \leq_E)$ represents the set of interpretations that are minimal with respect to $\leq_E$. $\leq_E$ is characterised by a measure of distance between an interpretation w and a belief profile E:

$$w \leq_E w' \text{ if and only if } d(w, E) \leq d(w', E)$$

$d(w, E)$ is defined by calculating a distance $d(w, K)$ between w and each belief base K in E and then using some function to aggregate these calculations. Here the distance $d(w, K)$ is defined using the min distance, which given some distance measure between interpretations is the minimal distance between the interpretation w and the models of the belief base K:

[2] These terms are used interchangeably.

[3] A generalisation of the concept of a set that allows multiple instances of an element.

[4] One can also add integrity constraints IC to the operation, such that a merge operator $\Delta_{IC}(E)$ takes a belief profile E as input and returns a set of interpretations that satisfies IC (i.e. IC is used to denote a possibly empty set of formulas that must be logical consequences of the merging result).

$$d(w, K) = \min_{w' \in K} d(w, w')$$

Finally, a function is required to calculate for each interpretation w the distance $d(w, E)$ based on the values of the distances $d(w, K)$ for each belief base K:

$$d(w, E) = \oplus_{K \in E} d(w, K)$$

Three aggregation functions investigated in [9] are:

- sum (Σ) - the total sum of $d(w, K)$ for each $K \in E$. In this case, a minimal $d(w, E)$ is simply a numerical minimum.
- max (Max) - the highest $d(w, K)$ out of each $K \in E$. In this case, a minimal $d(w, E)$ is simply a numerical minimum.
- leximax (GMax) - a tuple consisting of each $d(w, K)$ positioned from highest to lowest. In this case, a minimal $d(w, E)$ is determined by the minimum first position value, followed by the minimum second position value and so on.

The notation $\Delta^{d,k,\oplus}$ is used to denote the model-based merge operator that uses d as the distance between interpretations (*i.e.* $d(w, w')$), k as the distance between interpretations and belief bases (*i.e.* $d(w, K)$) and $\oplus$ as the belief profile aggregation function. For the remainder of this paper d will generally be omitted as it is set to the Hamming distance, which is the number of atomic evaluations on which two interpretations differ. As well as being the default choice for belief merging implementations, we implement this one particular distance function as it is the one used for the type of truthlikeness measures we will be bringing in later to connect with merging. Also, for the rest of this section k will be omitted as it is set to the min distance.

Example 1 provides a demonstration of these methods, which are particularly suited to aggregating preferences or opinions. The propositional logical space to be used throughout this paper consists of the atomic formulas $\{p_1, p_2, p_3\}$ and is depicted in Table 1, with interpretations w_1 to w_8.

	p_1	p_2	p_3
w_1	1	1	1
w_2	1	1	0
w_3	1	0	1
w_4	1	0	0
w_5	0	1	1
w_6	0	1	0
w_7	0	0	1
w_8	0	0	0

Table 1 The eight interpretations/possible worlds in the logical space based on the three atomic formulas p_1, p_2 and p_3.

Example 1 Mirroring the first scenario given in the introduction, in this example there are four belief bases such that:

- K_1 (Tom) $= \{w_1, w_2\}$
- K_2 (Richard) $= \{w_1, w_2\}$
- K_3 (Harry) $= \{w_2, w_4\}$
- K_4 (Bob) $= \{w_8\}$
- $E = \{K_1, K_2, K_3, K_4\}$

Table 2 contains a tabulation of the minimal distance between each interpretation and each belief base and the sum, max and leximax aggregation function results, with the results in bold being those that are selected.

State	K_1	K_2	K_3	K_4	Sum	Max	Leximax
w_1	0	0	1	3	4	3	(3,1,0,0)
w_2	0	0	0	2	**2**	2	(2,0,0,0)
w_3	1	1	1	2	5	2	(2,1,1,1)
w_4	1	1	0	1	3	**1**	**(1,1,1,0)**
w_5	1	1	2	2	6	2	(2,2,1,1)
w_6	1	1	1	1	4	**1**	(1,1,1,1)
w_7	2	2	2	1	7	2	(2,2,2,1)
w_8	2	2	1	0	5	2	(2,2,1,0)

Table 2 Distances and sum, max and leximax results

Thus:

- $\Delta^{\Sigma}(E) = \{w_2\}$
- $\Delta^{\mathrm{Max}}(E) = \{w_4, w_6\}$
- $\Delta^{G\mathrm{Max}}(E) = \{w_4\}$

In terms of the first scenario given in the introduction:

- According to $\Delta^{\Sigma}(E)$, the group should go out to dinner, watch a movie and not go to see some live music.
- According to $\Delta^{\mathrm{Max}}(E)$, the group should not go to see some live music and either go to dinner or watch a movie but not both.
- According to $\Delta^{G\mathrm{Max}}(E)$, the group should go out to dinner, not watch a movie and not go to see some live music.

□

Δ^{Σ} is termed a majority merging operator and the result of Δ^{Σ} operators can be seen as the election of the most popular choices amongst the integrity constraints. In our example above three out of the four belief bases had w_2 amongst its selections hence w_2 as the result. Δ^{Max} tries to minimise the more remote distances and $\Delta^{G\mathrm{Max}}$ captures this arbitration characteristic but does so in a more fine-grained way; given the presence of the relatively distant w_8

amongst the belief bases Δ^{Max} gives the more compromising $\{w_4, w_6\}$ and $\Delta^{G\text{Max}}$ narrows this down. Δ^{Σ} is the most democratic of the methods in the sense of majority rule though if there are belief base sources whose preferences have a significantly large distance from the result of Δ^{Σ} such that it would jeopardise their participation then $\Delta^{G\text{Max}}$ might be a more suitable choice [9, p. 785].

2.1 Merging using average distance

As has been established, a distance function must be used to calculate the distance between each belief base and each interpretation. However rather than using the min distance for $d(w, K)$, another option introduced in this paper is to use the average distance between interpretations and belief bases. That is:

$$d_{\text{avg}}(w, K) = \tfrac{1}{|K|} \sum_{v \in K} d(w, w')$$

Example 2 In this example there are four belief bases such that:

- $K_1 = \{w_1, w_6\}$
- $K_2 = \{w_2, w_3, w_5\}$
- $K_3 = \{w_2, w_5\}$
- $K_4 = \{w_5, w_6\}$
- $E = \{K_1, K_2, K_3, K_4\}$

Table 3 contains a tabulation of the average distance between each interpretation and each belief base and the sum aggregation function result.

State	K_1	K_2	K_3	K_4	Sum
w_1	2/2	3/3	2/2	3/2	4.5
w_2	2/2	4/3	2/2	3/2	4.83
w_3	4/2	4/3	4/2	5/2	7.83
w_4	4/2	5/3	4/2	5/2	8.17
w_5	2/2	4/3	2/2	1/2	**3.83**
w_6	2/2	5/3	2/2	1/2	4.17
w_7	2/2	5/3	4/2	3/2	6.17
w_8	4/2	6/3	4/2	3/2	7.5

Table 3 Distances and sum results using d_{avg}

The result is $\Delta^{\text{avg},\Sigma}(E) = \{w_5\}$.
□

The result of Example 2 represents the interpretation that is on average closest to each model of each belief base. Naturally we are prompted to next consider how and in what kinds of situations such a method could be useful.

Unlike the three operators given earlier, $\Delta^{\text{avg},\Sigma}$ is not suited to aggregating preferences or opinions. Rather, as will be supported by findings in the following sections, usage of this operator is suited to cases, such as the following two, where the belief bases represent something like factual theories as to what the true state of the world is and the goal of merging is to obtain a result that maximises truthlikeness:

1. The merging of theories/eyewitness accounts from independent sources pertaining to the truth of some matter. For example, in terms of Example 2, each K could represent the beliefs of a meteorologist concerning the truth of p_1 (it will be hot tomorrow), p_2 (it will rain tomorrow) and p_3 (it will be windy tomorrow). The meteorologist corresponding to K_1 believes that it will rain tomorrow and that either it will be both hot and windy or neither. K_2 believes that exactly any two of the three conditions will obtain and so on. The result of merging these four bases suggests that it will be rainy and windy tomorrow but it will not be hot.
2. The merging of scientific experiment results from independent sources pertaining to some hypothesis.

Thus in Example 2, we take each belief base to represent something like a theory as to what the true state of the world is; according to K_1 it is either w_1 or w_6, according to K_2 it is one of w_2, w_3 or w_5 and so on.

2.2 Properties of $\Delta^{\text{avg},\Sigma}$

Having introduced this new merging operator, it is now time to investigate some of its properties. To begin with, it is worth noting the equivalence between $\Delta^{\text{avg},\Sigma}$ and $\Delta^{\text{min},\Sigma}$ when belief bases are restricted to certain sets of models.

As introduced in [3], let *c-statement* (conjunctive statement) denote a propositional statement in conjunctive normal form, with each of the c conjuncts being either an atom or negated atom.[5] For example, $p_1 \wedge p_2$ is a c-statement whereas $p_1 \wedge (p_2 \vee p_3)$ is not. When the statement is the tautology, $c = 0$. With this class of statements, we have the following theorem:

Theorem 1 *If E consists only of belief bases corresponding to a c-statement then $\Delta^{avg,\Sigma}(E)$ and $\Delta^{min,\Sigma}(E)$ are equivalent.*

Proof See Appendix, Theorem 1.

In developing their merging framework, Konieczny and Pérez [8] provide a list of six postulates for merging operators.[6] Here we will determine which of these postulates are and which are not satisfied by this new operator, a task that will provide some important insights into its nature. According to Konieczny and Pérez, Δ is a merging operator if and only if it satisfies the following postulates:

[5] These statements are used in the basic feature approach to truthlikeness [3].

[6] This list was expanded to nine when incorporating integrity constraints [9].

(A1) $\Delta(E)$ is consistent.
(A2) If E is consistent then $\Delta(E) \equiv \wedge E$.
(A3) If $E_1 \equiv E_2$ then $\Delta(E_1) \equiv \Delta(E_2)$.
(A4) If $K_1 \wedge K_2$ is not consistent, then $\Delta(\{K_1, K_2\}) \nvdash K_1$.
(A5) $\Delta(E_1) \wedge \Delta(E_2) \vdash \Delta(E_1 \sqcup E_2)$.[7]
(A6) If $\Delta(E_1) \wedge \Delta(E_2)$ is consistent, then $\Delta(E_1 \sqcup E_2) \vdash \Delta(E_1) \wedge \Delta(E_2)$.

$\Delta^{\min,\Sigma}$ and $\Delta^{\min,GMax}$ are both examples of such merging operators.

The construction of $\Delta^{\text{avg},\Sigma}$ guarantees that it satisfies (A1). (A2) means that if the conjunction of the belief bases is consistent (i.e. there is a set of models common to all of the bases), then the result is simply the conjunction of the belief bases. Interestingly, $\Delta^{\text{avg},\Sigma}$ does not satisfy (A2); example 3 provides a counterexample.

Example 3

- $K_1 = \{w_1, w_2\}$
- $K_2 = \{w_1, w_4, w_6\}$
- $E = \{K_1, K_2\}$
- $\Delta^{\text{avg},\Sigma}(E) = \{w_2\}$

□

(A3) simply means that syntax is irrelevant, a postulate that is trivially satisfied given that $\Delta^{\text{avg},\Sigma}$ takes a model-based approach.

(A4) is a fairness postulate, which means that when two bases are merged, the operator should provide a result that favours neither of them. This postulate is also not satisfied by $\Delta^{\text{avg},\Sigma}$; example 4 provides a counterexample.

Example 4

- $K_1 = \{w_1, w_2, w_3\}$
- $K_2 = \{w_5\}$
- $E = \{K_1, K_2\}$
- $\Delta^{\text{avg},\Sigma}(E) = \{w_1\}$

As can be seen, although $K_1 \wedge K_2$ is not consistent, K_1 is favoured: $\Delta^{\text{avg},\Sigma}(\{K_1, K_2\}) \vdash K_1$.

□

(A5) and (A6) taken together mean that if the results of merging two separate profiles agree on a set of models, then the result of merging their union will be exactly that set of models.

Theorem 2 *$\Delta^{avg,\Sigma}$ satisfies (A5) and (A6).*

Proof See Appendix, Theorem 2.

[7] $\sqcup$ denotes the union of two profiles. Thus $\sqcup$ is the multiset equivalent of $\cup$ and can give results with multiple instances of an element.

The failure of $\Delta^{\text{avg},\Sigma}$ to satisfy A4 is not a problem given the purpose of this operator; we have an epistemic focus and are interested in merging for the purpose of approaching the truth, not in resolving conflict via majority or arbitration. As Grégoire and Konieczny [7, p. 12] note, A4 "is a strong impartiality requirement and may appear too strong in some cases"; this is one of those cases.

The failure of A2 is another issue; $\Delta^{\text{avg},\Sigma}$ does not necessarily respect agreement among agents in that even if all of the belief bases have at least one model in common (i.e. they are mutually consistent), the result of this merging operator will not necessarily consist of models from this common set. Whilst this peculiarity might strike some as particularly problematic, by demonstrating the truthlikeness conduciveness of $\Delta^{\text{avg},\Sigma}$ in the following sections, its general behaviour and failure to satisfy axioms such as A2 will be given justification according to its purpose. Let us for now discuss the matter a bit more and get a feel for how $\Delta^{\text{avg},\Sigma}$ behaves in relation to A2.

To begin with, it must be emphasised that often $\Delta^{\text{avg},\Sigma}$ will act in accordance with A2. Firstly, given Theorem 1, in cases where E is based on c-statements, A2 will be satisfied. Secondly, given sufficient reliable agreement amongst sources, the result of $\Delta^{\text{avg},\Sigma}$ will be in accordance with A2. The following serves as a good example:

Example 5

- $K_1 = \{w_1, w_2\}$
- $K_2 = \{w_1, w_3\}$
- $K_3 = \{w_1, w_5\}$
- $E = \{K_1, K_2, K_3\}$
- $\Delta^{\text{avg},\Sigma}(E) = \{w_1\}$

□

Furthermore, sometimes $\Delta^{\text{avg},\Sigma}$ will not strictly act in accordance with A2, but will respect agreement amongst belief bases, as shown with the following example, where the result is a subset of $\wedge E$:

Example 6

- $K_1 = \{w_1, w_2, w_3\}$
- $K_2 = \{w_1, w_2, w_4\}$
- $K_3 = \{w_1, w_2, w_5\}$
- $E = \{K_1, K_2, K_3\}$
- $\wedge E = \{w_1, w_2\}$
- $\Delta^{\text{avg},\Sigma}(E) = \{w_1\}$[8]

□

Sometimes the result of $\Delta^{\text{avg},\Sigma}$ will keep elements common to all belief bases with the addition of extra states that also reflect the average:

[8] As will become clearer once our formal truthlikeness measure is introduced, the narrowing down here of the result to w_1 rather than $\{w_1, w_2\}$ can be conducive to truthlikeness.

Example 7

- $K_1 = \{w_1, w_4\}$
- $K_2 = \{w_1, w_4\}$
- $K_3 = \{w_1, w_6\}$
- $E = \{K_1, K_2, K_3\}$
- $\Delta^{\text{avg},\Sigma}(E) = \{w_1, w_2\}$

□

Finally, there will be cases where no element common to all belief bases is kept. The following is but one example, along with a possible scenario to which it would be suited.

Example 8 Suppose, as represented in the following, that three investigators conduct independent tests into whether p_1 is the case, whether p_2 is the case and whether p_3 is the case.

Each investigator tests using their own methods, repeats their test three times and delivers three results. Furthermore, none of the testing methods are completely reliable. The results are summed up in the following list. So K_1 for example, observed p_1, p_2 and p_3 on their first trial, p_1, p_2 and $\neg p_3$ on their second trial and p_1, $\neg p_2$ and p_3 on their third trial.

- $K_1 = \{w_1, w_2, w_3\}$
- $K_2 = \{w_1, w_2, w_5\}$
- $K_3 = \{w_2, w_3, w_5\}$
- $E = \{K_1, K_2, K_3\}$
- $\Delta^{\text{avg},\Sigma}(E) = \{w_1\}$

□

Despite w_2 being common to all bases, the result of w_1 reflects the fact that taking all tests into consideration, p_1 occurred more as a result than $\neg p_1$, p_2 occurred more as a result than $\neg p_2$ and p_3 occurred more as a result than $\neg p_3$.

We continue by establishing a new property that I shall term Singleton Identity:

$$\text{(Sing Id) } \forall K \Delta(\{K\}) = K$$

This property, which follows from A2, is not satisfied by $\Delta^{\text{avg},\Sigma}$. For example:

- $K = \{w_2, w_3, w_5\}$
- $\Delta^{\text{avg},\Sigma}(\{K\}) = \{w_1\} \neq K$

Whilst this might be seen as strange for a merging operator, it will be explicated in terms of truthlikeness maximisation in Section 4.

Finally, a distinction is made in the literature between *majority* merging operators and *arbitration* merging operators [8]. A merging operator (that satisfies A1 - A6) is a majority operator if it satisfies the following:[9]

[9] $K^x = \underbrace{K \sqcup ... \sqcup K}_{x}$

$$(\text{M7})\ \forall K \exists x \Delta(E \sqcup K^x) \vdash K$$

This postulate expresses the fact that for any belief base, there is a sufficient number of times x the base can occur in a belief profile, at which point the base will prevail in the result of merging the profile.

As shown in [8,9], $\Delta^{\min,\Sigma}$ is a majority merging operator. Whilst $\Delta^{\text{avg},\Sigma}$ is not a merging operator in the sense of satisfying all of postulates (A1 - A6), as example 9 demonstrates, it does not satisfy (M7) either.

Example 9 It is clear that when $K = \{w_2, w_3, w_5\}$ for example and $E = \{\{w_1\}\}$, any value of n will be such that $\Delta(E \sqcup K^n) = \{w_1\} \nvdash K$.
□

A basic way to characterise arbitration is the following strong postulate:

$$(\text{A7'})\ \forall K \forall n \Delta(E \sqcup K^n) = \Delta(E \sqcup K)$$

Konieczny and Pérez [8] show that there can be no merging operator satisfying both (A1 - A6) and (A7'). Whilst $\Delta^{\text{avg},\Sigma}$ does not satisfy all of postulates (A1 - A6), it too does not satisfy the strong (A7').

In order to characterise majority insensitivity arbitration in a way that is consistent with (A1 - A6), they propose the following:

$$(\text{A7})\ \forall K' \exists K K' \nvdash K \forall n \Delta(K' \sqcup K^n) = \Delta(K' \sqcup K)$$

According to this postulate, "the result of the arbitration is [to a large extent] independent from the frequency of the different views" [8]. Konieczny and Pérez show that $\Delta^{\min,\text{Max}}$ satisfies (A7) and that $\Delta^{\min,G\text{Max}}$ satisfies (A7) if and only if the number of atoms n in the logical space is such that $n > 1$.

In fact, $\Delta^{\text{avg},\Sigma}$ also satisfies (A7) given this restriction.

Theorem 3 *The $\Delta^{avg,\Sigma}$ operator satisfies postulate (A7) if and only if $n > 1$*

Proof See Appendix, Theorem 3.

Thus interestingly, $\Delta^{\text{avg},\Sigma}$ and $\Delta^{\min,\Sigma}$ both use the sum aggregation function but $\Delta^{\min,\Sigma}$ satisfies the majority postulate whereas $\Delta^{\text{avg},\Sigma}$ does not. Instead, like the merging operators that use the max and leximax aggregation functions, $\Delta^{\text{avg},\Sigma}$ satisfies a form of arbitration.

3 Truthlikeness and estimated truthlikeness

Truthlikeness or verisimilitude [17] is a concept from the philosophy of science that concerns measuring the extent to which a theory approaches the truth. Popper was the first philosopher to seriously investigate this notion as part of his falsificationist program. In what sense can progress be made if fallibilism about science is adopted? Well, two theories might both be false, but one might be closer to the truth than the other and if one false theory is replaced by another false theory that is more truthlike then this constitutes progress.

More generally, the concept of truthlikeness can be applied to any collection of data/beliefs. A simple, mundane example will suffice to elucidate this notion and demonstrate its direct practical import. Take a situation in which a catering service is organising for an event. The truth is that at this event there will be exactly 64 guests. Now take the following statements:

(A) There will be 64 guests
(B) There will be somewhere between 50 and 90 guests
(C) There will be 65 guests
(D) There will be 30 guests

With regards to the number of guests, statement (A) is precisely true and thus is closest to the truth of the matter. Statements (C) and (D) are both false, but (C) gives a number closer to the actual number and is thus more truthlike than (D). The comparison between statements (B) and (C) is an illuminating one. Statement (B) is true and statement (C) is false, yet statement (C) seems to be more accurate, as it gives a number just above 64 as opposed to an unspecific yet true disjunctive range from 50 to 90. Thus whilst (B) is true and (C) is false, (C) is deemed more truthlike than (B). From this ordering we can see that in such cases the more truthlike a statement is the more useful it is and hence the value in devising methods that conduce truthlikeness.

There is a variety of approaches to constructing formal truthlikeness measures in logical space.[10] The Tichy/Oddie proposal to calculate the truthlikeness of a dataset/proposition makes it inversely related to the average distance between the models of the dataset and the interpretation that is the true state. The setup to calculate distances involves assigning a weight to each atomic element in the logical space such that these weights sum up to 1. The distance between two states is a summation of these weights for each atomic evaluation difference and is thus a derivation of the Hamming distance measured in terms of these weights. Thus variation in weights can be used to reflect differences in the truthlikeness significance an atom has. Here we shall simply assign a uniform weight of $\frac{1}{n}$ to each atomic element, where n is the number of atoms in the logical space. Thus in the case of our example this weight is $\frac{1}{3}$.

The truthlikeness of belief base K relative to some state w thus starts with the sum of these differences divided by the number of models for K and is given by the following function with range $[0, 1]$:

$$\mathrm{Tr}(K, w) = 1 - d_{\mathrm{avg}}(w, K)$$

With w_t standing for the actual state of the world, the actual truthlikeness of K is given by $\mathrm{Tr}(K, w_t)$, which we will shorten to $\mathrm{Tr}(K)$.[11] Throughout this paper the actual state will be set to w_1. As an example of this measure take the belief base $K = \{w_1, w_2\}$. Then $d_{\mathrm{avg}}(w_1, K) = \frac{1}{2} \times \frac{1}{3} = \frac{1}{6}$ and thus $\mathrm{Tr}(K) = 1 - \frac{1}{6} = \frac{5}{6}$.

[10] See [16,13] for some surveys.
[11] This is known as the Tichy-Oddie truthlikeness measure in the literature.

Besides the semantic problem of truthlikeness, which concerns what we mean when we claim that the proposition A is closer to the truth than the proposition B, there is also the epistemic problem, which concerns on what evidence we are to judge that A is closer to the truth than B [22, p. 121]. Regarding the latter, given some probability distribution the standard formula for expected utility in decision theory can be used to calculate the estimated truthlikeness of a dataset K given evidence Evd [15, p. 180]:

$$\mathrm{Tr}_{\mathrm{est}}(K|Evd) = \sum_{i=1}^{s} \mathrm{Tr}(K, w_i)\mathrm{Pr}(w_i|Evd)$$

s stands for the number of states in the logical space and $\mathrm{Pr}(w_i|Evd)$ is the conditional probability of w_i given Evd.

Example 10 Referring back to Table 1 and starting off with a uniform probability of $\frac{1}{8}$ for each possible state, conditionalization upon evidence $Evd = p_1 \vee \neg p_2 \vee \neg p_3$ gives the following probability distribution:

- $\mathrm{Pr}(w_5) = 0$
- $\mathrm{Pr}(w_1) = \mathrm{Pr}(w_2) = \mathrm{Pr}(w_3) = \mathrm{Pr}(w_4) = \mathrm{Pr}(w_6) = \mathrm{Pr}(w_7) = \mathrm{Pr}(w_8) = \frac{1}{7}$

With this distribution:

- $\mathrm{Tr}_{\mathrm{est}}(p_1 \wedge p_2 \wedge p_3|Evd) = 0.48$
- $\mathrm{Tr}_{\mathrm{est}}(\neg p_1 \wedge \neg p_2 \wedge \neg p_3|Evd) = 0.52$

As can be seen, given a certain piece of true evidence (since w_1 is the actual state Evd is true), there can be a marked difference between actual truthlikeness and estimated truthlikeness: $\mathrm{Tr}(p_1 \wedge p_2 \wedge p_3) = 1 > \mathrm{Tr}(\neg p_1 \wedge \neg p_2 \wedge \neg p_3) = 0$ whilst $\mathrm{Tr}_{\mathrm{est}}(p_1 \wedge p_2 \wedge p_3|Evd) < \mathrm{Tr}_{\mathrm{est}}(\neg p_1 \wedge \neg p_2 \wedge \neg p_3|Evd)$. □

Hopefully, this brief overview of truthlikeness suffices to give an adequate picture of its significance. For more discussion on the idea of using estimated truthlikeness as an epistemic utility see [15, 12].

4 Belief merging as truthlikeness maximisation

Combining a process of belief change with a truthlikeness measure opens up the possibility of exploring their relationship. Niiniluoto [14] and others [2, 20] for example have investigated general properties concerning how the truthlikeness measure of a belief base changes in relation to the process of belief revision.

Whilst there has been a modicum of papers investigating the combination of belief revision and truthlikeness, the combination of belief merging and truthlikeness has received close to no coverage. Regarding this combination, at least two distinct general topics can be investigated:

1. Given a merging method and truthlikeness measure, when/how does the truthlikeness of $\Delta(E)$ increase relative to the members of E. Are there ways to characterise the set of E such that the truthlikeness of $\Delta(E)$ is greater than the truthlikeness of the members of E?
2. Given a truthlikeness measure and one or more merging methods, which of the methods is most conducive to obtaining truthlike results.

To my knowledge Cevolani [1] provides the only published paper that investigates the combination of truthlikeness and belief merging. Using his basic feature approach to truthlikeness, he tackles the first of these topics, going through some simple test cases and exploring how the truthlikeness of $\Delta^{\min,\Sigma}$ results increase/decrease relative to the structure and truthlikeness of input belief bases.

In this paper we are interested in the second of these topics. We have introduced the $\Delta^{\text{avg},\Sigma}$ operator with the motivation of truthlikeness. How does it fare though in terms of truthlikeness conduciveness? In this section it will be demonstrated that $\Delta^{\text{avg},\Sigma}$ is actually a special case of a procedure that combines probabilities and maximises estimated truthlikeness.

Given some probability distribution across the possible states in a logical space, one approach to selecting an optimal set of states is to select those states which have the highest estimated truthlikeness. Niiniluoto [14] discusses this approach to using truthlikeness as a guiding factor in developing belief revision constructions. This form of belief revision by estimating truthlikeness involves preferring theories which have a larger estimated truthlikeness value than their rivals. Thus we get a rule of acceptance such that given evidence *Evd*, the resulting theory to be selected is simply that which has maximal estimated truthlikeness conditional upon *Evd*.

In cases where there is a set of such probability distributions from various sources, this approach would involve firstly combining the probability distributions into one and then using the result for the estimation calculations. There is a variety of approaches to combining probability distributions.[12] For our purposes we shall use the simple *linear opinion pool* approach to probability aggregation, which is a weighted linear combination of the source probabilities given by the formula:

$$\Pr(A) = \sum_{i=1}^{n} \mathrm{w}_i \Pr_i(A)$$

where n is the number of probability distributions, $\Pr_i(A)$ represents distribution i's probability for A and the weights w_i are non-negative and sum to one [4, p. 189]. These weightings can be adjusted to reflect the relative importance of or trust in a belief base source. For example, if two equally distinguished experts each contribute a probability distribution on some matter then the weights would be set to $\frac{1}{2}$. Alternatively, if the first contributor has a better record of prediction than the second, then the weights might be $\frac{3}{4}$ and

[12] See [4] for an overview.

$\frac{1}{4}$ respectively. Obviously probability pooling can be seen as a form of basic numerical merging. As we will now see, it actually forms one half of a procedure which combines probability aggregation and truthlikeness estimation and that, as it turns out, is equivalent to $\Delta^{\text{avg},\Sigma}$.

Given a belief profile E, we determine a set of probabilities by going through each belief base K in E and assigning to each of its models w a uniform probability $\Pr_K(w) = \frac{1}{|K|}$. These probabilities are then combined and the resulting probability distribution is used for estimated truthlikeness calculations.

Example 11 Take the following belief bases that were used in Example 2:

- $K_1 = \{w_1, w_6\}$
- $K_2 = \{w_2, w_3, w_5\}$
- $K_3 = \{w_2, w_5\}$
- $K_4 = \{w_5, w_6\}$
- $E = \{K_1, K_2, K_3, K_4\}$

Given this belief profile, the probability determination method just described results in:

- $\Pr_{K_1}(w_1) = \Pr_{K_1}(w_6) = \frac{1}{2}$
- $\Pr_{K_2}(w_2) = \Pr_{K_2}(w_3) = \Pr_{K_2}(w_5) = \frac{1}{3}$
- $\Pr_{K_3}(w_2) = \Pr_{K_3}(w_5) = \frac{1}{2}$
- $\Pr_{K_4}(w_5) = \Pr_{K_4}(w_6) = \frac{1}{2}$

Combining all of these probabilities using the linear opinion pool formula with a uniform weighting of $\frac{1}{4}$ for each of the four probability distribution sources, we get:

- $\Pr(w_1) = \frac{1}{2} \times \frac{1}{4} = \frac{1}{8}$
- $\Pr(w_2) = (\frac{1}{3} + \frac{1}{2}) \times \frac{1}{4} = \frac{5}{24}$
- $\Pr(w_3) = \frac{1}{3} \times \frac{1}{4} = \frac{1}{12}$
- $\Pr(w_5) = (\frac{1}{3} + \frac{1}{2} + \frac{1}{2}) \times \frac{1}{4} = \frac{1}{3}$
- $\Pr(w_6) = (\frac{1}{2} + \frac{1}{2}) \times \frac{1}{4} = \frac{1}{4}$

With this probability distribution, the estimated truthlikeness calculations for each state are:

- $\text{Tr}_{\text{est}}(w_1) = 0.625$
- $\text{Tr}_{\text{est}}(w_2) = 0.597$
- $\text{Tr}_{\text{est}}(w_3) = 0.347$
- $\text{Tr}_{\text{est}}(w_4) = 0.319$
- $\text{Tr}_{\text{est}}(w_5) = \mathbf{0.681}$
- $\text{Tr}_{\text{est}}(w_6) = 0.653$
- $\text{Tr}_{\text{est}}(w_7) = 0.403$
- $\text{Tr}_{\text{est}}(w_8) = 0.375$

The selected result based on highest estimated truthlikeness consists of the model w_5, the same result as Example 2.[13]
□

Let $F_{PT}(E)$ stand for this function that combines the probabilities of all belief bases in E with linear pooling and returns the set of models with highest estimated truthlikeness based on the combination distribution. As can be seen, if the $d(w, E)$ results from Example 2 are ordered from lowest to highest and the $\mathrm{Tr}_{\mathrm{est}}()$ results from Example 11 are ordered from highest to lowest then the orderings are equivalent. In fact, we have the following general result:

Theorem 4 *For any belief profile E, if each K in E adopts a uniform probability distribution on the set of its models, then the model ordering results of $\Delta^{avg,\Sigma}(E)$ ordered from lowest to highest and the model ordering results of $F_{PT}(E)$ ordered from highest to lowest are ordinally equivalent.*

Proof See Appendix, Theorem 4.

From Theorem 4 it simply follows that:

Corollary 1 *For any belief profile E, if each K in E adopts a uniform probability distribution on the set of its models, then $\Delta^{avg,\Sigma}(E) = F_{PT}(E)$.*

Using this equivalence as a starting point, this method of combining probability distributions and selecting states based on highest estimated truthlikeness can be seen as a more general way to perform the type of belief merging performed by the operator $\Delta^{\mathrm{avg},\Sigma}$, a way that can incorporate probability distributions into the calculations.

Example 12 Take a situation like that of Example 11 with the belief profile $E = \{K_1, K_2, K_3, K_4\}$. This time though, suppose that the probability distribution associated with each of the sources is non-uniform:

- $\mathrm{Pr}_{K_1}(w_1) = \frac{3}{4}, \mathrm{Pr}_{K_1}(w_6) = \frac{1}{4}$
- $\mathrm{Pr}_{K_2}(w_2) = \frac{1}{2}, \mathrm{Pr}_{K_2}(w_3) = \frac{3}{8}, \mathrm{Pr}_{K_2}(w_5) = \frac{1}{8}$
- $\mathrm{Pr}_{K_3}(w_2) = \frac{3}{4}, \mathrm{Pr}_{K_3}(w_5) = \frac{1}{4}$
- $\mathrm{Pr}_{K_4}(w_5) = \mathrm{Pr}_{K_4}(w_6) = \frac{1}{2}$

So for example, in this case the K_1 source believes that either w_1 or w_6 is the actual state, but w_1 is more likely. Combining all of these probabilities using the linear opinion pool formula with a uniform weighting of $\frac{1}{4}$ for each of the four probability distribution sources, we get:

- $\mathrm{Pr}(w_1) = \frac{3}{4} \times \frac{1}{4} = \frac{3}{16}$

[13] Using the Tichy-Oddie truthlikeness measure, one or more individual states will always be amongst the datasets with highest estimated truthlikeness. In cases where there is more than one such individual state, any datasets formed as a collection of these individual states will also have this highest truthlikeness value since the truthlikeness of a collection of states is simply their average truthlikeness.

- $\Pr(w_2) = (\frac{1}{2} + \frac{3}{4}) \times \frac{1}{4} = \frac{5}{16}$
- $\Pr(w_3) = \frac{3}{8} \times \frac{1}{4} = \frac{3}{32}$
- $\Pr(w_5) = (\frac{1}{8} + \frac{1}{4} + \frac{1}{2}) \times \frac{1}{4} = \frac{7}{32}$
- $\Pr(w_6) = (\frac{1}{4} + \frac{1}{2}) \times \frac{1}{4} = \frac{3}{16}$

With this probability distribution, the estimated truthlikeness calculations for each state are:

- $\mathrm{Tr}_{\mathrm{est}}(w_1) = \mathbf{0.667}$
- $\mathrm{Tr}_{\mathrm{est}}(w_2) = \mathbf{0.667}$
- $\mathrm{Tr}_{\mathrm{est}}(w_3) = 0.396$
- $\mathrm{Tr}_{\mathrm{est}}(w_4) = 0.396$
- $\mathrm{Tr}_{\mathrm{est}}(w_5) = 0.604$
- $\mathrm{Tr}_{\mathrm{est}}(w_6) = 0.604$
- $\mathrm{Tr}_{\mathrm{est}}(w_7) = 0.333$
- $\mathrm{Tr}_{\mathrm{est}}(w_8) = 0.333$

Thus although the belief profile in this example is the same as that of Example 11, with this probability distribution the result becomes $\{w_1, w_2\}$.
□

Given that the focus of this paper is on the merging of factual data, ending up with a form of merging that can incorporate probability distributions is a welcome extension. Unlike non-probabilistic voting preferences, the states comprising factual data from a source might very well have non-uniform probabilities associated with them to reflect the varying degrees of credence the source has.

It is worth mentioning that this equivalence elucidates the point from Section 2 that $\Delta^{\mathrm{avg},\Sigma}$ fails Singleton Identity and can produce a change when applied to a singleton belief profile; that is, where $E = \{K\}$, it is possible for $\Delta^{\mathrm{avg},\Sigma}(E) \neq K$. This is simply explained by the fact that a singleton belief profile corresponds to a single probability distribution and that selecting states with maximum estimated truthlikeness based on that probability distribution can give a result that differs from the initial states.

This method of combining probability distributions followed by maximising truthlikeness based on the resulting probability distribution was simply motivated by standard estimated truthlikeness involving a single distribution. Since estimated truthlikeness is based on a single probability distribution, it naturally followed in the case of multiple distributions to combine them into one and use the result to estimate truthlikeness. The equivalence between this method and $\Delta^{\mathrm{avg},\Sigma}$ was subsequently realised, thus providing mutual support. However, as pointed out by a referee, the 'merge + maximise' sequence could be reversed to get a 'maximise + merge' procedure. That is, another approach to dealing with multiple belief bases is to have agents maximise truthlikeness based on their own distributions first and then merge the results. Whilst it is beyond the scope of this paper to undertake an investigation into this approach, some remarks are in order. Firstly, this approach does not seem to

have support in the way of a correspondence such as that between $\Delta^{\text{avg},\Sigma}$ and 'merge + maximise'. Secondly, 'maximise + merge' using $\Delta^{\text{avg},\Sigma}$ would involve two levels of estimated truthlikeness maximisation. The first concerns the initial maximising that is performed for each base. The second concerns the truthlikeness maximisation that, as has been shown, is implicit in the $\Delta^{\text{avg},\Sigma}$ process. Is such a doubling in the one procedure legitimate in terms of expected utility theory? Finally, Jon Williamson [21] argues that judgements are best aggregated by merging the evidence on which they are based, rather than by directly merging the judgements themselves:

> one should not apply belief merging directly to agents' judgement sets. Instead, one should apply belief merging to the agents' evidence bases and consider a single hypothetical agent with this merged evidence base. One can determine appropriate degrees of belief for this hypothetical agent using objective Bayesian theory and then determine appropriate judgements using decision theory. These judgements can be viewed as the aggregate of the original agents' judgement sets. [21][p. 468].

Whilst the technical details of Williamson's argument are somewhat different from the formal apparatus of this paper, a general application of his idea suggests that it would be better to merge the probabilities (evidence) first and then maximise truthlikeness (decision theoretic judgement), rather than *vice versa*.

5 Concluding remarks

It was established at the outset of this paper that the process of merging applies to two different types of problem:

- aggregating preferences/opinions (PM)
- aggregating factual beliefs/data (BM)

Whilst it has not been the intention of this paper to provide a complete division and formal characterisation of operators for these two types of problems, a few concluding points are in order. Unlike the min-based merging operators brought in from the literature in Section 2, the operator $\Delta^{\text{avg},\Sigma}$ is not suited to PM. For one thing, this is apparent in its failure to satisfy the A2 property; if a group of agents all have a common preference, then that preference should be selected. With regards to BM things are different. Whilst this paper offers no definitive account of BM operators and does not rule out the usage of operators that satisfy properties A1 - A6 (such as the min-based merging operators), it has conclusively shown that in cases where the factual beliefs/data are aggregated with an interest in maximising the truthlikeness of the result, then an operator such as $\Delta^{\text{avg},\Sigma}$ is a most suitable choice. Despite the fact that this method violates certain postulates established in the belief merging literature, these violations are reasonable given its success with

regards to truthlikeness conduciveness. At the very least, we have established that not all BM operators need satisfy A1 - A6.

With regards to future work on the topic of belief merging and truthlikeness, one place to start is the equivalence demonstrated in Section 5 of $\Delta^{\text{avg},\Sigma}$ and the combine probabilities + maximise estimated truthlikeness procedure. Since this general combination approach involves parameters that can be adjusted and components that can be replaced with alternatives there is room to tailor and experiment with it. Here are three possibilities:

- If a different truthlikeness measure is adopted then it can straightforwardly replace the method used in this paper.
- Using methods of probability aggregation other than the linear opinion pool.
- Alternatives to expected utility decision theory as a method of estimating truthlikeness.[14]

As just mentioned, there is also the possibility of investigating other untested combination approaches. One alternative is to first calculate the results of maximum estimated truthlikeness for each belief base individually and then combine these results somehow. For example, if one were to first calculate the results of maximum estimated truthlikeness for each belief base and then select the resulting states with the highest minimum estimated truthlikeness values across all belief bases, then one would get the counterpart of the operator $\Delta^{\text{avg,Max}}$. Such possibilities and other methods involving belief merging and truthlikeness await further investigation.

Appendix

Theorem 1. If E consists only of belief bases corresponding to a c-statement then $\Delta^{\text{avg},\Sigma}(E)$ and $\Delta^{\text{min},\Sigma}(E)$ are equivalent.

Proof It suffices to show that for any given belief base K that corresponds to a c-statement and interpretations w_x and w_y:

$$d_{\min}(w_x, K) \gtreqless d_{\min}(w_y, K) \text{ if and only if } d_{\text{avg}}(w_x, K) \gtreqless d_{\text{avg}}(w_y, K)$$

We will show it using $<$, but the same applies to the other two relations. For any K, $|K| = 2^{n-c}$ models, where c is the number of atoms in the corresponding c-statement and n is the number of atoms in the logical space.

Now, suppose that $d_{\min}(w_x, K) = x < d_{\min}(w_y, K) = y$. The value of $d_{\text{avg}}(w_x, K)$ will equate to $\frac{|K|x+z}{|K|}$ and the value of $d_{\text{avg}}(w_y, K)$ will equate to $\frac{|K|y+z}{|K|}$.

We get $|K|x$ from the fact that whatever state w_k in K gives the min distance x, the other $|K| - 1$ states will agree with the atom valuations of

[14] See [11,19] for a survey.

w_k that give that distance and hence add $|K| - 1$ to the total sum difference. Same for $|K|y$.

Here, z is equal to the sum distance between any given model w of K and K: $z = d_{\text{sum}}(w \in K, K) = \sum_{v \in K} d(w, w')$. For any two models w_a, $w_b \in K$ such that K corresponds to a c-statement, $d_{\text{sum}}(w_a, K) = d_{\text{sum}}(w_b, K)$, which is why z is the same for both $d_{\text{avg}}(w_x, K)$ and $d_{\text{avg}}(w_y, K)$. Specifically, it can be seen that each of $d_{\text{sum}}(w_a, K)$ and $d_{\text{sum}}(w_b, K)$ will equal $\sum_{i=0}^{n-c} \binom{n-c}{i} i$.

Given this, it straightforwardly follows that $d_{\text{min}}(w_x, K) = x < d_{\text{min}}(w_y, K) = y$ if and only if $d_{\text{avg}}(w_x, K) = \frac{|K|x+z}{|K|} < d_{\text{avg}}(w_y, K) = \frac{|K|y+z}{|K|}$.

For example, let $K = \{w_1, w_2, w_3, w_4\}$ (p_1). $d_{\text{min}}(w_1, K) = 0 < d_{\text{min}}(w_7, K) = 1$. $d_{\text{sum}}(w_1, K) = d_{\text{sum}}(w_3, K) = 4$. $d_{\text{avg}}(w_1, K) = \frac{(4\times 0)+4}{4} = 1 < d_{\text{avg}}(w_7, K) = \frac{(4\times 1)+4}{4} = 2$.

Theorem 2. $\Delta^{\text{avg},\Sigma}$ satisfies (A5) and (A6).

Proof When using the sum aggregation function, $d(w, E_1)+d(w, E_2) = d(w, E_1 \sqcup E_2)$. Let d_1 stand for the minimum distance for $\Delta(E_1)$ and d_2 stand for the minimum distance for $\Delta(E_2)$. For each $w_x \in \Delta(E_1) \wedge \Delta(E_2)$, $d(w_x, E_1 \sqcup E_2) = d_1 + d_2$. Since $d_1 + d_2$ is the minimum distance for $d(w, E_1) + d(w, E_2)$, it follows that it is the minimum distance for $d(w, E_1 \sqcup E_2)$ and hence the result of $\Delta(E_1) \wedge \Delta(E_2)$ will equate to the result of $\Delta(E_1 \sqcup E_2)$.

Theorem 3. The $\Delta^{\text{avg},\Sigma}$ operator satisfies postulate (A7) if and only if $n > 1$

Proof Preliminary observation suggests that for any K' it will often be easy to find some such K. The following selection procedure guarantees one way to go about finding some K for any K' such that $K' \nvdash K \forall n \Delta(K' \sqcup K^n) = \Delta(K' \sqcup K)$.

First, divide K' into two types of cases: (1) those where K' consists of two or more models and (2) those where K consists of one model.

When K' consists of two or more models, perform $\Delta(\{K'\})$ and select one of the resulting models w_x such that $w_x \in \Delta(\{K'\})$. This w_x will by definition have a minimum distance out of all $d(w, K')$. Next, set $K = \{w_x\}$. Since $|K'| > 1 = |K|$, $K' \nvdash K$. Now, since $w_x \in \Delta(\{K'\})$, $\Delta(K' \sqcup K) = \{w_x\}$, since K is adding a minimum distance of 0 to only $d(w_x, E)$. But then, any repetition of $K = \{w_x\}$ added to the input profile will only contribute 0 to $d(w_x, E)$ and so $\Delta(K' \sqcup K) = \{w_x\} = \Delta(K' \sqcup K^n)$.

In the case where K' consists of one model, this strategy is not possible. What we can do is the following. Let $K' = \{w_y\}$. Let K consist of each of the n states that differs from w_y by one atom. Given this, $K' \nvdash K$. Next, we show that generally it is a fact that $\Delta(\{K\}) = \{w_y\}$. Given this, it follows that $\Delta(K' \sqcup K^n) = \Delta(K' \sqcup K)$.

The average distance between K and w_y will be $\frac{n}{n}$ (there are n states in K that differ from w_y by 1), so if $\Delta(\{K\}) \neq \{w_y\}$, we would need to show that

there is some w_z such that the average distance between K and w_z is less than $\frac{n}{n}$. Since n is the common denominator, this amounts to showing that there is some w_z with total sum difference less than n. But this is not possible:

- If w_z is not in K, then the total sum distance is m such that $m > n$, as none of the distances between w_z and members of K can be 0 and at least one will have to be greater than 1, as there is only one state (i.e. w_y) such that each member of K will differ by 1.
- If w_z is in K, then, apart from w_z, which has a distance of 0 to itself, it can be seen that all other members of K will have a difference of 2 from w_z, so the total sum distance will be $\frac{0+2(n-1)}{n} = 2n-2$.
  For all values $n > 2$, $n < 2n-2$, thus $\Delta(\{K\}) = \{w_y\}$. When $n = 2$, $n = 2n-2$, in which case K has an average distance of 1 to all of the four possible states and thus $\Delta(\{w_y\} \sqcup K)$ is still equal to w_y.

We finish off with a counterexample for when $n = 1$. Let w_1 be the states in which the one atom p is true and w_2 be the state in which it is false. $K' = \{w_1\}$ and $K = \{w_2\}$ is the only possible type of setup such that $K^{'} \nvdash K$. But $\Delta(K^{'} \sqcup K) = \{w_1, w_2\} \neq \Delta(K^{'} \sqcup K^2) = \{w_2\}$.

Theorem 4. For any belief profile E, if each K in E adopts a uniform probability distribution on the set of its models, then the model ordering results of $\Delta^{\mathrm{avg},\Sigma}(E)$ ordered from lowest to highest and the model ordering results of $F_{PT}(E)$ ordered from highest to lowest are ordinally equivalent.

Proof Noting that $\mathrm{Tr}_{\mathrm{est}}(K|Evd) = \sum_{i=1}^{s} \mathrm{Tr}(K, w_i)\mathrm{Pr}(w_i|Evd) = 1 - [\sum_{i=1}^{s} d_{\mathrm{avg}}(w_i, K)\mathrm{Pr}(w_i|Evd)]$, let $F^*_{PT}(E)$ stand for a modified $F_{PT}(E)$, where rather than calculating $\sum_{i=1}^{s} \mathrm{Tr}(K, w_i)\mathrm{Pr}(w_i|Evd)$ for each interpretation, $\sum_{i=1}^{s} d_{\mathrm{avg}}(w_i, K)\mathrm{Pr}(w_i|Evd)$ is used instead.

So as to simplify things, we will show that $\Delta^{\mathrm{avg},\Sigma}(E)$ ordered from lowest to highest is ordinally equivalent to $F^*_{PT}(E)$ ordered from lowest to highest. We do this by showing that for each pair of interpretations (w_x, w_y):

$$d(w_x, E) \gtreqless d(w_y, E) \text{ if and only if}$$
$$[\sum_{i=1}^{s} d(w_i, w_x)\mathrm{Pr}(w_i|Evd)] \gtreqless [\sum_{i=1}^{s} d(w_i, w_y)\mathrm{Pr}(w_i|Evd)]$$

Let $\mathrm{In}(w_z, K)$ be a function such that:

$$\mathrm{In}(w_z, K) = \begin{cases} 1 & \text{if } w_z \in K \\ 0 & \text{if } w_z \notin K \end{cases}$$

Beginning with $F^*_{PT}(E)$, we first establish that given the belief profile $E = \{K_1, K_2, ..., K_k\}$ the uniform probability of some interpretation w_z is given by:

$$\Pr(w_z) = \left[\sum_{i=1}^{k} \frac{\mathrm{In}(w_z, K_i)}{|K_i|}\right]^{\frac{1}{k}}$$

Next we proceed with the following formulation of the relation $[\sum_{i=1}^{s} d(w_i, w_x)\Pr(w_i|Evd)] \gtreqless [\sum_{i=1}^{s} d(w_i, w_y)\Pr(w_i|Evd)]$:

$$\sum_{i=1}^{s} d(w_i, w_x)\Pr(w_i|Evd) \gtreqless \sum_{i=1}^{s} d(w_i, w_y)\Pr(w_i|Evd)$$

if and only if

$$\sum_{i=1}^{s} \Pr(w_i|Evd)(d(w_i, w_x) - d(w_i, w_y)) \gtreqless 0$$

if and only if

$$\sum_{i=1}^{s}\left[\sum_{j=1}^{k} \frac{(d(w_i, w_x) - d(w_i, w_y))\mathrm{In}(w_i, K_j)}{|K_j|k}\right] \gtreqless 0$$

if and only if

$$\sum_{j=1}^{k}\left[\frac{\sum_{i=1}^{s}[(d(w_i, w_x) - d(w_i, w_y))\mathrm{In}(w_i, K_j)]}{|K_j|}\right] \gtreqless 0$$

Next we formulate the relation between $d(w_x, E)$ and $d(w_y, E)$.

$$\sum_{j=1}^{k}\left[\frac{\sum_{i=1}^{s} d(w_i, w_x)\mathrm{In}(w_i, K_j)}{|K_j|}\right] \gtreqless \sum_{j=1}^{k}\left[\frac{\sum_{i=1}^{s} d(w_i, w_y)\mathrm{In}(w_i, K_j)}{|K_j|}\right]$$

if and only if

$$\sum_{j=1}^{k}\left[\frac{\sum_{i=1}^{s}[d(w_i, w_x)\mathrm{In}(w_i, K_j)] - \sum_{i=1}^{s}[d(w_i, w_y)\mathrm{In}(w_i, K_j)]}{|K_j|}\right] \gtreqless 0$$

if and only if

$$\sum_{j=1}^{k}\left[\frac{\sum_{i=1}^{s}[(d(w_i,w_x)-d(w_i,w_y))\mathrm{In}(w_i,K_j)]}{|K_j|}\right]\lessgtr 0$$

At this stage we note that the formulation for $F^*_{PT}(E)$ is equivalent to the formulation for the relation between $d(w_x, E)$ and $d(w_y, E)$ and thus the proof is complete.

References

1. Cevolani, G.: Truth approximation, belief merging, and peer disagreement. Synthese **191**(11), 2383–2401 (2014)
2. Cevolani, G., Calandra, F.: Approaching the truth via belief change in propositional languages. In: EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association, pp. 47–62 (2010)
3. Cevolani, G., Crupi, V., Festa, R.: Verisimilitude and belief change for conjunctive theories. Erkenntnis **75**(2), 183–202 (2011)
4. Clemen, R.T., Winkler, R.L.: Combining probability distributions from experts in risk analysis. Risk Analysis **19**(2), 187–203 (1999)
5. Eckert, D., Pigozzi, G.: Belief merging, judgment aggregation, and some links with social choice theory. In: In Belief Change in Rational Agents: Perspectives from Artificial Intelligence, Philosophy, and Economics, Dagstuhl Seminar Proceedings 05321 (2005)
6. Everaere, P., Konieczny, S., Marquis, P.: The epistemic view of belief merging: Can we track the truth? In: H. Coelho, R. Studer, M. Wooldridge (eds.) ECAI 2010 - 19th European Conference on Artificial Intelligence, Lisbon, Portugal, August 16-20, 2010, Proceedings, vol. 215, pp. 621–626. IOS Press (2010)
7. Grégoire, E., Konieczny, S.: Logic-based approaches to information fusion. Information Fusion **7**, 4–18 (2006)
8. Konieczny, S., Pérez, R.P.: On the logic of merging. pp. 488–498 (1998)
9. Konieczny, S., Pérez, R.P.: Merging information under constraints: A logical framework. Journal of Logic and Computation **12**(5), 773–808 (2002)
10. Konieczny, S., Pérez, R.P.: Logic based merging. Journal of Philosophical Logic **40**(2), 239–270 (2011)
11. Machina, M.J.: Non-expected utility theory. In: S.N. Durlauf, L.E. Blume (eds.) The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke (2008)
12. Niiniluoto, I.: Truthlikeness. D. Reidel, Dordrecht, Holland (1987)
13. Niiniluoto, I.: Verisimilitude: The third period. British Journal for the Philosophy of Science **49**(1), 1–29 (1998)
14. Niiniluoto, I.: Revising beliefs towards the truth. Erkenntnis **75**(2), 165–181 (2011)
15. Oddie, G.: Likeness to Truth. D. Reidel, Dordrecht, Holland (1986)
16. Oddie, G.: The content, consequence and likeness approaches to verisimilitude: Compatibility, trivialization, and underdetermination. Synthese **190**(9), 1647–1687 (2013)
17. Oddie, G.: Truthlikeness. In: E.N. Zalta (ed.) The Stanford Encyclopedia of Philosophy, summer 2014 edn. Stanford University (2014). URL = <`http://plato.stanford.edu/archives/sum2014/entries/truthlikeness/`>
18. Pigozzi, G., Hartmann, S.: Judgment aggregation and the problem of truth-tracking. In: Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge, TARK '07, pp. 248–252. ACM, New York, NY, USA (2007)
19. Schmidt, U.: Alternatives to expected utility: Formal theories. In: S. Barber, P. Hammond, C. Seidl (eds.) Handbook of Utility Theory, pp. 757–837. Springer US (2004)

20. Schurz, G.: Verisimilitude and belief revision. with a focus on the relevant element account. Erkenntnis **75**(2), 203–221 (2011)
21. Williamson, J.: Aggregating judgements by merging evidence. Journal of Logic and Computation **19**(3), 461–473 (2009)
22. Zwart, S.: Refined Verisimilitude. Kluwer, Dordrecht, Holland (2001)