



HAL
open science

Belief Reconfiguration

Sébastien Konieczny, Elise Perrotin, Ramón Pino Pérez

► **To cite this version:**

Sébastien Konieczny, Elise Perrotin, Ramón Pino Pérez. Belief Reconfiguration. European Conference on Logics in Artificial Intelligence (JELIA 2023), Sep 2023, Dresden, Germany. pp.446-461, 10.1007/978-3-031-43619-2_31 . hal-04494547

HAL Id: hal-04494547

<https://hal.science/hal-04494547>

Submitted on 7 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Belief Reconfiguration

Sébastien Konieczny^[0000-0002-2590-1222], Elise Perrotin^[0000-0003-4188-3789],
and Ramón Pino Pérez^[0000-0002-2912-263X]

CRIIL - CNRS - Université d'Artois - France

Abstract. We study a generalisation of iterated belief revision in a setting where we keep track not only of the received information (in the form of messages) but also of the source of each message. We suppose that we have a special source, the oracle, which never fails. That is, all of the information provided by the oracle is assumed to be correct. We then evaluate the reliability of each source by confronting its messages with the facts given by the oracle. In this case it is natural to give higher priority to messages coming from more reliable sources. We therefore re-order (reconfigure) the messages with respect to the reliability of the sources before performing iterated belief revision. We study how to compute this reliability, and the properties of the corresponding reconfiguration operators.

1 Introduction

In this work our aim is to provide a more realistic account of iterated belief revision [5, 2, 13]. A requirement in standard iterated belief revision is that every new evidence acquired by the agent is more plausible than the previous one. This assumption is usually not explicitly stated, but it is enforced by the postulates characterizing these operators. This is usually called “Primacy of Update”.

However, if this assumption is plausible in some scenarios, for instance when one wants to model the evolution of scientific theories, it makes relatively little sense in everyday-life scenarios: we usually obtain pieces of information at different points in time, which we consider sufficiently reliable to be taken into account, but they are not magically ordered from least to most plausible over the course of our life. We therefore need to adapt this “ideal” framework of iterated belief revision so that we can represent real, practical applications.

One way of doing this is to weaken the postulates in order to remove primacy of update altogether. This leads for instance to improvement operators [15, 14], which make ‘softer’ changes than revision operators. With improvement operators it is possible to completely get rid of primacy of update [19]. Another way to do this is by considering credibility limited revision operators [11, 8, 3, 4, 9].

Rather than weakening all of the revision steps, in this work we wish to base priority given to the information on the reliability of the source behind it, while remaining as close as possible to the standard iterated revision framework. To this end, we need to explicitly introduce a way of measuring reliability of these sources of information. In order to do this, we define a more general framework in which we attach to each received piece of evidence the source that provides this

information. We suppose that we have a special source, the oracle, which never fails (i.e. it only provides truthful information). Then, by comparing the claims of the different sources with the truth that we obtain from the oracle, we can have an estimation of their reliability. Once we have this reliability estimation, in order to work with the iterated revision operators, we reconsider the sequence of received information with respect to this reliability by reordering the messages, putting the messages of the more reliable sources after those of the less reliable ones. Hence the name of “reconfiguration” for these operators.

This reordering does not affect the relative order of messages of individual sources (or of sources of the same reliability). This is also expected, as we can expect more recent messages of a given source to be more reliable: this source may learn new things, and correct some of her initial mistakes.

The proposed setting is very natural in numerous scenarios. Suppose for instance that you receive information from different friends, whom you consider reliable enough to be listened to. Here the oracle is directly observed evidence, i.e. what you can experiment about the world. Your direct observations of the world will sometimes contradict previous information you have received from friends; in this case you can reevaluate the reliability of those friends.

In the next section we give the definitions of our reconfiguration operators, as well as their associated iterated belief revision operators and reliability functions. In Section 3 we focus on the reliability functions, propose expected properties for these functions, enumerate several possible instantiations, and check which properties they satisfy. Then in Section 4 we study the properties of the corresponding reconfiguration operators. In Section 5 we provide an example illustrating the behavior of these operators. We then conclude in Section 6 with a discussion of related and future work.

2 Reconfiguration

In this section we formally describe the reconfiguration process. We consider a propositional language \mathcal{L}_P built from a finite set of propositional variables P and the usual connectives. Lowercase letters of the Greek alphabet denote formulas. An interpretation ω is a total function from P to $\{0, 1\}$. The set of all interpretations is denoted by \mathcal{W} . An interpretation ω is a model of a formula $\varphi \in \mathcal{L}_P$ if and only if it makes it true in the usual truth functional way. The set of models of a formula φ is denoted by $\llbracket \varphi \rrbracket$. The set of consistent formulae is denoted by \mathcal{L}_P^c . \perp (resp. \top) is the Boolean constant always false (resp. true).

2.1 The ingredients

We consider an *epistemic space* $\mathcal{E} = \langle E, B \rangle$, where E is a set of epistemic states and B is a mapping $B : E \rightarrow \mathcal{L}_P$ characterizing beliefs in each epistemic state. We suppose the existence of an epistemic state Ψ_\top such that $B(\Psi_\top) = \top$ (for a systematic treatment of these structures and revision operators defined on them see [20]).

We also consider a DP revision operator (i.e. an operator satisfying all postulates of [5]) $\circ : E \times \mathcal{L}_P^c \rightarrow E$. If $\varphi_1 \dots \varphi_k$ is a sequence of formulas we define $\circledast(\varphi_1 \dots \varphi_k)$ as $\circledast(\varphi_1 \dots \varphi_k) = (\dots ((\Psi_{\top} \circ \varphi_1) \circ \varphi_2) \dots) \circ \varphi_n$.

In our framework we consider sequences of messages sent by a variety of sources. To formalize this, let \mathcal{S} be a finite set of sources of information, and o be an additional special source, called the oracle, which only provides correct information. We define $\mathcal{S}^* = \mathcal{S} \cup \{o\}$.

A *message* m is a couple (s, φ) where $s \in \mathcal{S}^*$ is the source of the message and $\varphi \in \mathcal{L}_P$ is the information given by the message. We denote by \mathcal{M} the set of messages. For a message $m_i = (s_i, \varphi_i)$, we denote by m_i^s the source of the message, that is s_i , and we denote by m_i^φ the information of m , that is φ_i . We consider that individual messages are always consistent, i.e. for any message m , $m^\varphi \not\vdash \perp$.

Given a finite sequence of messages $\sigma = m_1 \dots m_n$, we define $src(\sigma)$, the set of sources of σ , as $src(\sigma) = \{m_k^s : k = 1, \dots, n\}$. For any source $s \in \mathcal{S}^*$, we denote by σ_s the sequence of formulas in the messages from source s in σ . For example, if $\sigma = (s_1, p)(s_2, q)(s_1, r)$ then $\sigma_{s_1} = \langle p, r \rangle$. We furthermore define $\varphi_s^\sigma = \bigwedge_{\varphi \in \sigma_s} \varphi$ the conjunction of all information given by s in σ . In particular, if σ_s is the empty sequence (i.e. there are no messages from s in the sequence σ) then $\varphi_s^\sigma = \top$. As the messages of s in σ may be inconsistent with each other, resulting in $\varphi_s^\sigma \equiv \perp$, we also define the *opinion* of s as $\mathcal{O}_s^\sigma = B(\circledast\sigma_s)$.

As the oracle never fails, we suppose that $\varphi_o^\sigma \not\vdash \perp$ in all considered sequences σ . We call *Seq* the set of finite sequences σ of messages such that $\varphi_o^\sigma \not\vdash \perp$. We use \cdot as a concatenation symbol: if $\sigma = m_1 \dots m_n$ then $\sigma \cdot m = m_1 \dots m_n m$.

We assume that we have a function δ assigning a degree of reliability to a source in \mathcal{S} appearing in a finite sequence of messages as well as to the oracle. The mapping δ is a partial function having domain $Seq \times \mathcal{S}^*$ and co-domain \mathbb{R}^+ . If $(\sigma, s) \in Seq \times \mathcal{S}$ and $s \in src(\sigma)$ then $\delta(\sigma, s)$ is defined. The value $\delta(\sigma, s)$ represents the reliability degree of the source s given the sequence σ . $\delta(\sigma, o)$ is always defined, and required to be maximal, that is, $\delta(\sigma, o) > \delta(\sigma, s)$ for any source $s \in \mathcal{S}$. We adopt the notation $\delta_\sigma(s)$ instead of $\delta(\sigma, s)$.

2.2 The framework

In order to reorganize messages of a sequence $\sigma = m_1 \dots m_n$ from least to most reliable, we define a permutation function r which is a bijection from $\{1, \dots, n\}$ to $\{1, \dots, n\}$ as follows:

$$r(i) < r(j) \text{ iff } \begin{cases} \delta_\sigma(m_i^s) < \delta_\sigma(m_j^s), \text{ or} \\ \delta_\sigma(m_i^s) = \delta_\sigma(m_j^s) \text{ and } i < j \end{cases}$$

Intuitively, $r(i)$ is the relative reliability of the message m_i in σ . That is, a message m_i is considered less reliable than a message m_j if either the source of m_i is less reliable than that of m_j , or the sources of both messages have the same reliability and m_i was announced before m_j .

Finally, a sequence of messages $\sigma \in Seq$ induces an epistemic state $\Psi_\sigma \in E$ in which all messages have been taken into account relative to their respective reliability. Let $\sigma = m_1 \dots m_n$ be a sequence of messages. We define $r_i = r^{-1}(i)$

for $i \leq n$. Because we assume the oracle to be the (strictly) most reliable source, we know that if at least one message in σ is not from the oracle then there is a $k \leq n$ such that $m_{r_k}^s \neq o$ and $m_{r_j}^s = o$ for all $j > k$. We define Ψ_σ as follows:

$$\Psi_\sigma = (\otimes(m_{r_1}^\varphi \dots m_{r_k}^\varphi)) \circ \varphi_\sigma^\sigma.$$

In order to study the mechanisms of this revision process when receiving new messages, we define a new epistemic space as follows.

Definition 1. *The epistemic space of sequences \mathcal{E}_{Seq} associated to a DP operator \circ and a reliability function δ is defined by putting $\mathcal{E}_{Seq} = \langle Seq, B_{Seq} \rangle$, where the epistemic states are sequences of messages and B_{Seq} is the mapping $B_{Seq} : Seq \rightarrow \mathcal{L}_P$ defined by $B_{Seq}(\sigma) = B(\Psi_\sigma)$.*

3 Reliability functions

A key element of the reconfiguration framework is the function δ , which evaluates the reliability of sources by comparing their messages with those of the oracle. There are many possible definitions for this function. In this section we give a few general properties that such a function should satisfy. We then give some natural examples of δ functions, and check these functions against the stated properties.

3.1 Desirable properties

We give some natural properties which we expect any “good” reliability function δ to satisfy. We call these properties “general properties”. We then provide some additional optional properties which make sense in some contexts, and can be satisfied depending on the desired behavior of δ .

General properties

1. **(Source independence)** If $\sigma_s = \sigma'_s$ and $\sigma_o = \sigma'_o$ then $\delta_\sigma(s) = \delta_{\sigma'}(s')$.
i.e.: A source’s evaluation is independent of other sources. It depends solely on what the source and the oracle have announced.
2. **(Syntax independence)** If $\sigma = m_1 \dots m_n$ and $m_i^\varphi \equiv \psi$ for some i then $\delta_{\sigma[(m_i^s, \psi)/m_i]} = \delta_\sigma$, where $\sigma[(m_i^s, \psi)/m_i]$ is the sequence σ in which m_i has been replaced by the message (m_i^s, ψ) .
i.e.: Two logically equivalent messages have exactly the same effect.
3. **(Oracle)** For any $s \in src(\sigma) \setminus \{o\}$, we have $\delta_\sigma(s) < \delta_\sigma(o)$.
i.e.: The oracle is the (strictly) most reliable source.
4. **(Maximality)** For any $s, s' \in src(\sigma) \setminus \{o\}$, if $\varphi_s^\sigma \wedge \varphi_o^\sigma \not\vdash \perp$ then $\delta_\sigma(s') \leq \delta_\sigma(s)$.
i.e.: All consistent sources which have never contradicted the oracle have the same reliability, which is the maximal reliability among sources other than the oracle.
5. **(Non-maximality)** For any $s, s' \in src(\sigma) \setminus \{o\}$ if $\varphi_s^\sigma \wedge \varphi_o^\sigma \not\vdash \perp$, $\varphi_{s'}^\sigma \not\vdash \perp$ and $\varphi_{s'}^\sigma \wedge \varphi_o^\sigma \vdash \perp$ then $\delta_\sigma(s') < \delta_\sigma(s)$.
i.e.: A source who has contradicted the oracle will always be strictly less reliable than one who has made no mistakes.

Optional properties

6. If $\varphi_s^\sigma \equiv \varphi_s^{\sigma'}$ and $\varphi_s^\sigma \not\vdash \perp$ then $\delta_\sigma(s) = \delta_{\sigma'}(s)$.
i.e.: The reliability function does not depend on the exact messages provided by the agent, but only on their conjunction. In other words, providing any number of messages or just one message with their conjunction leads to the same reliability.
7. If $s \in \text{src}(\sigma) \setminus \{o\}$ and $\psi \wedge \varphi_o^\sigma \not\vdash \perp$ then $\delta_{\sigma \cdot (s, \psi)}(s) \geq \delta_\sigma(s)$.
i.e.: Not contradicting the oracle cannot decrease reliability of a source.
8. If $s \in \text{src}(\sigma) \setminus \{o\}$ and $\psi \wedge \varphi_o^\sigma \vdash \perp$ then $\delta_{\sigma \cdot (s, \psi)}(s) \leq \delta_\sigma(s)$.
i.e.: Contradicting the oracle cannot increase reliability of a source.
9. If $s, s' \in \text{src}(\sigma) \setminus \{o\}$, $\delta_\sigma(s) \leq \delta_\sigma(s')$, $\psi \wedge \varphi_o^\sigma \vdash \perp$ and $\psi' \wedge \varphi_o^\sigma \not\vdash \perp$ then $\delta_{\sigma \cdot (s, \psi)}(s) < \delta_{\sigma \cdot (s', \psi')}(s')$.
i.e.: A message directly contradicting the oracle is strictly worse (for reliability) than a message that does not contradict it.
10. If $\delta_\sigma(s) \leq \delta_\sigma(s')$ then $\delta_{\sigma \cdot (s, \psi) \cdot (s', \psi)}(s) \leq \delta_{\sigma \cdot (s, \psi) \cdot (s', \psi)}(s')$.
i.e.: If two sources give the same information then their relative reliability remains unchanged.
11. If $\sigma = m_1 \dots m_n$ and m is a message, call $\sigma^{+i, m}$ the sequence σ in which m is inserted after m_i : $\sigma^{+i, m} = m_1 \dots m_i m m_{i+1} \dots m_n$. Consider ψ such that $\psi \wedge \varphi_o^\sigma \vdash \perp$, and suppose that $\psi \wedge \varphi_s^\sigma \not\vdash \perp$ and for any α in σ_s , $\alpha \wedge \varphi_o^\sigma \not\vdash \perp$. Then $\delta_{\sigma^{+i, m}}(s) \leq \delta_{\sigma^{+j, m}}(s)$ if $i \geq j$.
i.e.: This is a temporality property. Contradicting the oracle is more problematic the more recently it has been done. This implies that we consider as more reliable a source has made a mistake a long time ago (and has had the time to correct it) than one that has made a mistake more recently.
12. If $\psi \wedge \varphi_s^\sigma \vdash \perp$, $\psi' \wedge \varphi_s^\sigma \not\vdash \perp$, $\psi \wedge \varphi_o^\sigma \not\vdash \perp$ and $\psi' \wedge \varphi_o^\sigma \not\vdash \perp$ then $\delta_{\sigma \cdot (s, \psi)}(s) < \delta_{\sigma \cdot (s, \psi')}(s)$.
i.e.: All the other properties focus on comparing messages from a given source to the messages of the oracle. Here we add a more local estimation of reliability, only confronting messages from a same source, and “punishing” sources that contradict themselves.

3.2 Some options for δ

We now give some examples of definitions for the reliability function δ . For all the considered functions we put $\delta_\sigma(o) = \infty$, and only give definitions of $\delta_\sigma(s)$ for $s \neq o$.

We wish to assess sources' reliability based on the consistency of their messages with the information from the oracle. We do this by using some inconsistency measure, that is, a function $d : \mathcal{L}_P^c \times \mathcal{L}_P^c \rightarrow \mathbb{R}^+$ which is intended to measure the disagreement between two consistent formulas [12, 10, 23, 1, 24]. We suppose that d is congruent with respect to logical equivalence and symmetric. We also suppose that for any formulas φ and ψ , if $\varphi \wedge \psi \not\vdash \perp$ then $d(\varphi, \psi) = 0$, and if $\varphi \wedge \psi \vdash \perp$ then $d(\varphi, \psi) > 0$. One example of such a function is the drastic measure d_D , defined by $d_D(\varphi, \psi) = 0$ if $\varphi \wedge \psi$ is consistent, otherwise $d_D(\varphi, \psi) = 1$.

We define $M_d = \max\{d(\varphi, \psi) : \varphi, \psi \in \mathcal{L}_P^c\}$. We extend d to $\mathcal{L}_P \times \mathcal{L}_P$ by putting $d(\varphi, \psi) = M_d + 1$ if φ or ψ is inconsistent.

A first naive definition for a reliability function is as follows:

$$\delta_\sigma^1(s) = M_d - d(\varphi_s^\sigma, \varphi_o^\sigma).$$

Note that if a source s is contradictory, that is, if $\varphi_s^\sigma \equiv \perp$, then $\delta_\sigma^1(s) = -1$. More generally, the messages from one source may become inconsistent with each other over time, and we wish to give sources the opportunity to correct past mistakes. We consider that when a source contradicts its past messages, its current opinion is that conveyed by its later messages, and that is what its reliability should be assessed from. There are several ways to implement this. The first is to consider the source's opinion as defined in Section 2.1:

$$\delta_\sigma^2(s) = M_d - d(\mathcal{O}_s^\sigma, \varphi_o^\sigma).$$

However, as the properties of a source's opinion \mathcal{O}_s^σ are difficult to characterize, we might want a simpler way to take into account the combination of its messages even when it contradicts itself. One way to achieve this is, when $\sigma_s = \alpha_1 \dots \alpha_k$, to define $c_s^\sigma = \alpha_{j_{\text{cons}}(\sigma_s)} \wedge \dots \wedge \alpha_k$, where $j_{\text{cons}}(\sigma_s) = \min\{j \mid \alpha_j \wedge \dots \wedge \alpha_k \not\vdash \perp\}$. For example, if $\sigma_s = \langle p, q, r, \neg q \rangle$ then $j_{\text{cons}}(\sigma_s) = 3$ and $c_s^\sigma = r \wedge \neg q$. We then define

$$\delta_\sigma^3(s) = M_d - d(c_s^\sigma, \varphi_o^\sigma).$$

Rather than considering a 'global opinion' for each source, we might want to take into account separately each individual message. This allows us, in particular to put weights on the evaluation of messages, so that older messages 'count less' when assessing a source's reliability. We may then also consider not only how much the source contradicts the oracle, but also how much the source contradicts itself from message to message. A general formula for computing the reliability of a source $s \neq o$ after a sequence σ in which $\sigma_s = \alpha_1 \dots \alpha_k$ could be:

$$\delta_\sigma^*(s) = \frac{1}{W_k} \sum_{i=1}^k w_{i,k} (A \cdot \text{OC}(\alpha_i, \varphi_o^\sigma) + B \cdot \text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1}))$$

where W_k is a normalization factor, $w_{i,k}$ is a weight function favoring more recent messages, $\text{OC}(\alpha_i, \varphi_o^\sigma)$ is a measure of how much the oracle is contradicted by α_i , $\text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1})$ is a measure of how much α_i contradicts its own previous messages, and A and B are weights representing the importance given to consistency with the oracle and with own previous messages respectively.

We put the following constraints on the elements of this definition: first, $A > 0$ and $B \geq 0$. For weights, we require for all k and i :

$$W_k, w_{i,k} > 0; \quad w_{k,k} = 1; \quad w_{i,k} \leq w_{i+1,k}; \quad w_{i,k} \leq w_{i,k+1}; \quad W_k \leq W_{k+1}.$$

As for the contradiction factors $\text{OC}(\alpha_i, \varphi_o^\sigma)$ and $\text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1})$, we suppose that both $\text{OC}(\alpha_i, \varphi_o^\sigma)$ and $\text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1})$ have a maximum and a minimum possible value, denoted maxOC , minOC , maxSC and minSC respectively. We then require $\text{minOC}, \text{minSC} \leq 0 \leq \text{maxOC}, \text{maxSC}$ and:

$$\begin{aligned} \text{OC}(\alpha_i, \varphi_o^\sigma) &= \text{maxOC} \text{ iff } \alpha_i \wedge \varphi_o^\sigma \not\vdash \perp; \\ &\text{if } \alpha_1 \wedge \dots \wedge \alpha_i \not\vdash \perp \text{ then } \text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1}) = \text{maxSC}; \\ &\text{if } \alpha_{i-1} \wedge \alpha_i \vdash \perp \text{ then } \text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1}) < \text{maxSC}; \\ &\text{if } \alpha_1 \wedge \dots \wedge \alpha_{i-1} \not\vdash \perp \text{ and } \alpha_1 \wedge \dots \wedge \alpha_i \vdash \perp \text{ then } \text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1}) < \text{maxSC}. \end{aligned}$$

Here are some examples of instantiations of these elements:

- $W_k = 1$ (no normalization) or $W_k = \sum_{i=1}^k w_{i,k}$;
- $w_{i,k} = 1$ (no weighting) or $w_{i,k} = (1 - \varepsilon)^{k-i}$ for some $\varepsilon < 1$;
- $\text{OC}(\alpha_i, \varphi_o^\sigma) = M_d - d(\alpha_i, \varphi_o^\sigma)$ (here $\text{maxOC} = M_d$) or $\text{OC}(\alpha_i, \varphi_o^\sigma) = -d(\alpha_i, \varphi_o^\sigma)$ (here $\text{maxOC} = 0$);
- for $i > 1$, $\text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1}) = -\frac{j_{\text{cons}}(\alpha_1 \dots \alpha_i) - 1}{i-1}$.

We now give a few instantiations of this definition:

$$\begin{aligned} \delta_\sigma^4(s) &= \frac{1}{k} \sum_{i=1}^k (M_d - d(\alpha_i, \varphi_o^\sigma)); \\ \delta_\sigma^5(s) &= \frac{1}{W_k} \sum_{i=1}^k w_{i,k} (M_d - d(\alpha_i, \varphi_o^\sigma)); \\ \delta_\sigma^6(s) &= \sum_{i=1}^k w_{i,k} (M_d - d(\alpha_i, \varphi_o^\sigma)); \\ \delta_\sigma^7(s) &= \frac{1}{W_k} \sum_{i=1}^k w_{i,k} (-d(\alpha_i, \varphi_o^\sigma) + \text{ctr}(\alpha_i)); \\ \delta_\sigma^8(s) &= -\sum_{i=1}^k w_{i,k} d(\alpha_i, \varphi_o^\sigma); \end{aligned}$$

where $w_{i,k}$ follow the non-trivial definition given above, $W_k = \sum_{i=1}^k w_{i,k}$, and $\text{ctr}(\alpha_i)$ follows the definition for $\text{SC}(\alpha_i, \alpha_1 \dots \alpha_{i-1})$ given above. The function δ^4 is normalized, but all messages have the same weight. It is a special case of δ^5 (for $\varepsilon = 0$), which features normalization and increasing weights for each message. The function δ^6 has increasing weights, but no normalization. Finally, the functions δ^7 and δ^8 consider negative reliability evaluations, with δ^7 also taking into account whether the source contradicts itself, and δ^8 not being normalized.

Here we have proposed two approaches to computing a source's reliability: either aggregating its messages into a 'global opinion' to compare to the oracle's announcements, or considering each of its messages separately. A core difference in these two approaches can be seen as follows: suppose that the oracle has announced $\neg(p \wedge q)$, and that a source announces p , then q . With the first approach we consider this to be equivalent to the source announcing $p \wedge q$, and being completely incorrect. With the second approach, we allow an interpretation in which the announcement of q is a correction of previous statements, that is, the source might have updated its opinion from $p \wedge \neg q$ to $q \wedge \neg p$.

3.3 Discussion of the proposed functions and properties

We now evaluate the properties and δ functions proposed above against each other, to confirm whether they indeed make sense¹.

General properties We first check whether the proposed functions satisfy the required general properties.

Proposition 1. *The functions δ^1 , δ^2 , δ^3 , and δ^* satisfy properties 1, 2, and 3, that is, source and syntax independence and oracle maximality.*

Proof. This follows from the definition of the functions and the fact that the inconsistency measure d is syntax-independent.

¹ Because of space constraints we do not put the proofs in the paper and only give some intuitions behind the results. The full proofs can be found in the supplementary material.

Proposition 2. *The function δ^1 satisfies properties 4 (Maximality) and 5 (Non-maximality). The functions δ^2 and δ^3 satisfy property 4 but not property 5. The function δ^* satisfies properties 4 and 5 if either $\frac{\sum_{i=1}^k w_{i,k}}{W_k}$ is constant or $\max\text{OC} = B\max\text{SC} = 0$; otherwise it satisfies neither property.*

The function δ^1 therefore satisfies all general properties, whereas δ^2 and δ^3 fail to satisfy property 5, so they should not be considered as satisfying reliability functions. For the general δ^* function, properties 1, 2 and 3 are always satisfied, and we need some mild additional condition to satisfy also properties 4 and 5.

Optional properties We now turn to the optional properties and check which of them are satisfied by the different proposed reliability functions.

Following Proposition 2, from now on we consider for δ^* only the cases where either $\frac{\sum_{i=1}^k w_{i,k}}{W_k}$ is constant in k (we can consider w.l.o.g. that this constant is 1), or $\max\text{OC} = B\max\text{SC} = 0$. This, in particular, rules out the function δ^6 .

Proposition 3. *The functions δ^1 , δ^2 and δ^3 satisfy property 6; the function δ^* does not.*

The intuition here is that with δ^* a source can increase its reliability by repeating tautologies. We now study properties 7, 8 and 9.

Proposition 4. *The functions δ^1 , δ^2 and δ^3 do not satisfy property 7 or property 9; they satisfy property 8 when the inconsistency measure d is the drastic measure d_D .*

Intuitively, these functions compute ‘how wrong’ the combination of a source’s messages is; a message might be correct but result in a mistake when combined with previous messages, or it might be incorrect but result in a lesser mistake when combined with previous messages. In particular, even if a source only makes mistakes, their evaluation can evolve from being ‘very wrong in general’ to being ‘almost correct in general’. Requiring d to be a 0/1 function removes this possibility of being ‘almost correct’.

The interplay between the different elements of δ^* are more complex, and we give some cases in which δ^* satisfies the different properties rather than giving a general criteria. We consider three additional properties in particular. The first is $B = 0$, so that not contradicting oneself cannot compensate for contradicting the oracle, or the other way around. The second is for there to be an $\varepsilon < 1$ such that $w_{i,k+1} = (1 - \varepsilon)w_{i,k}$ for all k and $i \leq k$, so that we can better characterize the evolution of a source’s reliability when it provides a new message. The third is for OC to be a 0/1 function. This means that for all formulas φ and ψ either $\text{OC}(\psi, \varphi) = \min\text{OC}$ or $\text{OC}(\psi, \varphi) = \max\text{OC}$. As noted above, this prevents sources’ reliability from increasing when making a mistake by removing the possibility of being ‘almost correct’.

We start by considering the case where $\sum_{i=1}^k w_{i,k} = W_k$.

Proposition 5. *δ^* satisfies property 7 when $\frac{\sum_{i=1}^k w_{i,k}}{W_k} = 1$, $B = 0$ and there exists some $\varepsilon \leq 1$ such that $w_{i,k+1} = (1 - \varepsilon)w_{i,k}$ for all i and k . In this case we have $\delta_{\sigma \cdot (s, \psi)}^*(s) = \delta_{\sigma}^*(s)$ iff $\delta_{\sigma}^*(s) = A\max\text{OC}$ and $\psi \wedge \varphi_{\sigma}^{\circ} \not\vdash \perp$.*

	1	2	3	4	5	6	7	8	9	10	11	12
δ^1	✓	✓	✓	✓	✓	✓	×	0/1	×	×	✓	✓
δ^2	✓	✓	✓	✓	×	✓	×	0/1	×	×	×	×
δ^3	✓	✓	✓	✓	×	✓	×	0/1	×	×	×	×
δ^4	✓	✓	✓	✓	✓	×	✓	0/1	0/1	×	✓	×
δ^5	✓	✓	✓	✓	✓	×	✓	0/1	0/1	×	✓	×
δ^6	✓	✓	✓	×	×							
δ^7	✓	✓	✓	✓	✓	×	×	$\varepsilon=0$ and 0/1	$\psi \wedge \varphi \vdash \perp \Rightarrow$ $\text{OC}(\psi, \varphi) \leq -1$	×	✓	✓
δ^8	✓	✓	✓	✓	✓	×	✓	0/1	✓	✓	✓	×

Table 1. Properties satisfied by the proposed reliability functions. Conditions that are both necessary and sufficient are in dark blue; merely sufficient conditions are in cyan.

Proposition 6. When $\frac{\sum_{i=1}^k w_{i,k}}{W_k} = 1$ for all k , $B = 0$, the function δ^* satisfies property 8 iff OC is a 0/1 function. In this case we have $\delta_{\sigma \cdot (s, \psi)}^*(s) = \delta_\sigma^*(s)$ iff $\delta_\sigma^*(s) = \text{AminOC}$ and $\psi \wedge \varphi_\sigma^\sigma \vdash \perp$.

Proposition 7. The function δ^* satisfies property 9 when $\frac{\sum_{i=1}^k w_{i,k}}{W_k} = 1$ for all k , there exists some $\varepsilon \leq 1$ such that $w_{i,k+1} = (1 - \varepsilon)w_{i,k}$ for all i and k , $B = 0$, and OC is a 0/1 function.

In particular the functions δ^4 and δ^5 satisfy property 7, and they satisfy properties 8 and 9 when d is the drastic measure d_D .

We now consider the case where $\frac{\sum_{i=1}^k w_{i,k}}{W_k}$ is not constant in k . Then in particular $\text{maxOC} = B\text{maxSC} = 0$.

Recall that W_k is a normalization factor. There are essentially two meaningful options in terms of normalization: either having $\frac{\sum_{i=1}^k w_{i,k}}{W_k}$ be constant, or having no normalization. We now focus on the latter case and require W_k to be constant in k . We once again consider the particular case where $\frac{w_{i,k+1}}{w_{i,k}}$ is constant in i and k .

Proposition 8. Suppose that $\text{maxOC} = B\text{maxSC} = 0$ and that there exist some W and ε such that $W_k = W$ and $w_{i,k+1} = (1 - \varepsilon)w_{i,k}$ for all k and $i \leq k$. Then the function δ^* satisfies property 7 iff $B = 0$; it satisfies property 8 iff either $\varepsilon = 0$ or $B = 0$ and OC is a 0/1 function; it satisfies property 9 iff for all ψ and φ such that $\psi \wedge \varphi \vdash \perp$ we have $\text{OC}(\psi, \varphi) \leq \frac{B}{A}\text{minSC}$.

In particular δ^7 does not satisfy property 7; it satisfies property 8 iff $\varepsilon = 0$ and d is the drastic measure d_D ; and it satisfies property 9 iff $d(\psi, \varphi) \geq 1$ for any φ and ψ such that $\varphi \wedge \psi \vdash \perp$. On the other hand, δ^8 satisfies properties 7 and 9, and it satisfies property 8 when d is the drastic measure d_D .

Proposition 9. The functions δ^1 , δ^2 and δ^3 do not satisfy property 10. The function δ^* does not satisfy it if W_k is not constant in k for $k > 1$ or $B \neq 0$. The function δ^* does satisfy property 10 when W_k is constant in k , $B = 0$ and $\frac{w_{i,k+1}}{w_{i,k}}$ is constant in i and k .

Intuitively, property 10 requires the impact of announcing a formula on a source's reliability not to depend on the rest of the source's messages. In particular, the function δ^8 satisfies this property, while the functions δ^4 , δ^5 and δ^7 do not.

Proposition 10. *The functions δ^1 and δ^* satisfy property 11; δ^2 and δ^3 do not.*

Proposition 11. *The function δ^1 satisfies property 12; the functions δ^2 and δ^3 do not. The function δ^* satisfies it iff $B \neq 0$.*

Table 1 sums up the properties satisfied by the reliability functions δ^1 – δ^8 . The 0/1 symbol represents the condition that d is the drastic measure d_D .

4 Reconfiguration operators

Let us now study the properties of the corresponding reconfiguration operators, i.e. the operators that we obtain when we use the reliability function to re-order (reconfigure) the sequence of messages.

Recall that using a DP operator \circ defined on an epistemic space $\mathcal{E} = \langle E, B \rangle$ and a reliability function δ , we build a new epistemic space $\mathcal{E}_{Seq} = \langle Seq, B_{Seq} \rangle$ where the elements of Seq (the new epistemic states) are sequences of messages and B_{Seq} is as in Definition 1. In this epistemic space we define a new operator \bullet as follows:

Definition 2. *The function $\bullet : Seq \times \mathcal{M} \rightarrow Seq$, called a reconfiguration operator, is defined in the following way: if $\sigma = m_1 \dots m_n$, then $\sigma \bullet m = m_1 \dots m_n m$.*

Let us give the translation² of the standard DP postulates [5] in this framework:

- (**r-R*1**) $B_{Seq}(\sigma \bullet m) \vdash m^\varphi$
- (**r-R*2**) If $B_{Seq}(\sigma) \wedge m^\varphi \not\vdash \perp$ then $B_{Seq}(\sigma \bullet m) \equiv B_{Seq}(\sigma) \wedge m^\varphi$
- (**r-R*3**) $B_{Seq}(\sigma \bullet m) \not\vdash \perp$
- (**r-R*4**) If $m^\varphi \equiv m'^\varphi$ and $m^s = m'^s$ then $B_{Seq}(\sigma \bullet m) \equiv B_{Seq}(\sigma \bullet m')$
- (**r-R*5**) Let μ be a formula and m_1, m_2 messages such that $m_1^s = m_2^s$ and $m_2^\varphi = m_1^\varphi \wedge \mu$; then $B_{Seq}(\sigma \bullet m_1) \wedge \mu \vdash B_{Seq}(\sigma \bullet m_2)$
- (**r-R*6**) Let μ be a formula and m_1, m_2 messages such that $m_1^s = m_2^s$ and $m_2^\varphi = m_1^\varphi \wedge \mu$; then, if $B_{Seq}(\sigma \bullet m_1) \wedge \mu \not\vdash \perp$ then $\vdash B_{Seq}(\sigma \bullet m_2) \vdash B_{Seq}(\sigma \bullet m_1) \wedge \mu$
- (**r-C1**) If $m_2^\varphi \vdash m_1^\varphi$ then $B_{Seq}((\sigma \bullet m_1) \bullet m_2) \equiv B_{Seq}(\sigma \bullet m_2)$
- (**r-C2**) If $m_2^\varphi \vdash \neg m_1^\varphi$ then $B_{Seq}((\sigma \bullet m_1) \bullet m_2) \equiv B_{Seq}(\sigma \bullet m_2)$
- (**r-C3**) If $B_{Seq}(\sigma \bullet m_2) \vdash m_1^\varphi$ then $B_{Seq}((\sigma \bullet m_1) \bullet m_2) \vdash m_1^\varphi$
- (**r-C4**) If $B_{Seq}(\sigma \bullet m_2) \not\vdash \neg m_1^\varphi$ then $B_{Seq}((\sigma \bullet m_1) \bullet m_2) \not\vdash \neg m_1^\varphi$

Please see [5] for a more complete description of these postulates. Briefly, (**r-R*1**) means that the last information of the sequence should be believed after the change. (**r-R*2**) means that when the new piece of information is consistent with the current beliefs of the agent, then the result should be the conjunction. (**r-R*3**) is a bit stronger than the original DP postulate, requiring coherence unconditionally (since we suppose that each message is consistent). (**r-R*4**) is the Independence of syntax postulate. (**r-R*5**) and (**r-R*6**) relates the change by a conjunction with the change by an element of the conjunction.

² We put the prefix r- (for reconfiguration) before the translated postulate.

(**r-C1**) says that if a message m_2^φ is logically stronger than m_1^φ (and provided by the same source), then we obtain the same result if we make the change by m_1^φ and then by m_2^φ and if we make the change directly by m_2^φ . (**r-C2**) says that if m_2^φ contradicts m_1^φ (and if they are provided by the same source), then we obtain the same result if we make the change by m_1^φ and then by m_2^φ and if we make the change directly by m_2^φ . (**r-C3**) says that if a change by m_2^φ implies m_1^φ , then making the change by m_1^φ before the one by m_2^φ should not hurt m_1^φ (so the result still implies m_1^φ). (**r-C4**) says that if a change by m_2^φ does not imply $\neg m_1^\varphi$, then making the change by m_1^φ before the one by m_2^φ should not help $\neg m_1^\varphi$ (so the result still not implies $\neg m_1^\varphi$).

It is easy to see that these postulates do not hold in general for reconfiguration, due to the reordering during the process, but they hold under certain restrictions on the reliability of the new information, showing that we keep the DP behavior when possible:

Definition 3. *We say that the operator \bullet has a DP behavior with respect to the triple σ, m_1, m_2 (a sequence of messages and two messages respectively) if the postulates (r-R*1–r-R*6) and (r-C1–r-C4) are satisfied.*

Proposition 12. *Let σ, m_1, m_2 be a sequence of messages and two messages respectively. Then*

1. *If m_1^s and m_2^s are two sources with highest reliability in the sequences $\sigma \cdot m_1$ and $\sigma \cdot m_2$, then \bullet has a DP behavior with respect to the triple σ, m_1, m_2 .*
2. *If $m_1^s, m_2^s \neq o$ and m_1^s (resp. m_2^s) is the source with highest reliability among the sources different from the oracle in the sequence $\sigma \cdot m_1$ (resp. $\sigma \cdot m_2$) and $\varphi_o^\sigma \equiv \top$, then \bullet has a DP behavior with respect to the triple σ, m_1, m_2 .*

Despite the fact that reconfiguration operators are not designed to be DP iterated belief revision operators, as the reconfiguration (re-ordering) has an important impact on how the last message is treated, these results illustrate the fact that we keep the DP iteration flavor, and maintain the DP iteration behavior in particular when receiving messages from a most reliable source.

5 Example

We now provide an example illustrating the impact of recomputing reliability following announcements from the oracle. For simplicity we use the epistemic space of total preorders over interpretations, where for a total preorder \preceq , $\llbracket B(\preceq) \rrbracket = \min(\preceq)$ and the underlying belief revision operator is Nayak's lexicographic revision operator \circ_N [18] defined as follows: $\preceq \circ_N \alpha = \preceq'$ where $\omega \preceq' \omega'$ iff $\omega \in \llbracket \alpha \rrbracket$ or $\omega' \notin \llbracket \alpha \rrbracket$. We consider the reliability function δ^{4d_D} , which uses the drastic measure d_D and computes the proportion of messages from a given source which contradicts the oracle. We denote by $\bullet_N^{d_D}$ the reconfiguration operator defined from \circ_N and the reliability function δ^{4d_D} .

Example 1. We consider three sources, and the following sequence of messages: $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \bullet_N^{d_D} (s_3, \neg a \wedge \neg c) \bullet_N^{d_D} (o, a) \bullet_N^{d_D} (s_3, a \wedge b) \bullet_N^{d_D}$

$(s_2, c \wedge \neg b) \bullet_N^{d_D} (\mathbf{o}, \neg c)$. Let us see and comment what happens at each iteration. In order to simplify the notations, we will write $\Psi \equiv \alpha$ instead of $B_{Seq}(\Psi) \equiv \alpha$.

1. $(s_1, a \wedge \neg c) \equiv a \wedge \neg c$. There is only one message for the moment, so there is no reason to reject it.

2. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \equiv a \wedge c$. Source s_2 contradicts source s_1 , but as the oracle has not yet given any information we keep the messages in order of reception, and accept the message from s_2 . Taking this temporal order into account (instead of finding some kind of consensus or compromise with operators such as belief merging [16]) can be justified by the fact that s_2 has potentially benefited from more time than s_1 to check this piece of information.

3. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \equiv a \wedge b \wedge c$. Source s_1 sends a new message about b . As b had not been mentioned up to this point we can accept it in addition to the previous message of s_2 .

4. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \bullet_N^{d_D} (s_3, \neg a \wedge \neg c) \equiv \neg a \wedge b \wedge \neg c$. Source s_3 sends a message that contradicts both s_1 and s_2 , but as it is the most recent message we accept it.

5. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \bullet_N^{d_D} (s_3, \neg a \wedge \neg c) \bullet_N^{d_D} (\mathbf{o}, \mathbf{a}) \equiv a \wedge b \wedge c$. We receive our first message from the oracle, which makes us realize that s_3 is the least reliable source; we still cannot distinguish between s_1 and s_2 . The reconfiguration gives the following sequence: $(\Psi_{\top \circ_N} \neg a \wedge \neg c \circ_N a \wedge \neg c \circ_N a \wedge c \circ_N b \circ_N \mathbf{a})$.

6. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \bullet_N^{d_D} (s_3, \neg a \wedge \neg c) \bullet_N^{d_D} (\mathbf{o}, \mathbf{a}) \bullet_N^{d_D} (s_3, a \wedge b) \equiv a \wedge b \wedge c$. We receive a new message from s_3 , which does not contradict the oracle, but contradicts the previous message from s_3 . This can mean that source s_3 has realized that it was wrong, revised its beliefs, and now sends a message it believes to be correct. Depending on the reliability function used, this can increase or decrease its reliability (since on the one hand this last message was consistent with the oracle, but on the other hand s_3 has contradicted itself). With the reliability function we have chosen, s_3 remains less reliable than the other sources. The corresponding reconfiguration gives the sequence $(\top \circ_N \neg a \wedge \neg c \circ_N a \wedge b \circ_N a \wedge \neg c \circ_N a \wedge c \circ_N b \circ_N \mathbf{a})$.

7. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \bullet_N^{d_D} (s_3, \neg a \wedge \neg c) \bullet_N^{d_D} (\mathbf{o}, \mathbf{a}) \bullet_N^{d_D} (s_3, a \wedge b) \bullet_N^{d_D} (s_2, c \wedge \neg b) \equiv a \wedge \neg b \wedge c$. We receive a new message from s_2 , which is one of the most trustworthy sources. We therefore accept this message. The corresponding reconfiguration sequence is $(\top \circ_N \neg a \wedge \neg c \circ_N a \wedge b \circ_N a \wedge \neg c \circ_N a \wedge c \circ_N b \circ_N c \wedge \neg b \circ_N \mathbf{a})$.

8. $(s_1, a \wedge \neg c) \bullet_N^{d_D} (s_2, a \wedge c) \bullet_N^{d_D} (s_1, b) \bullet_N^{d_D} (s_3, \neg a \wedge \neg c) \bullet_N^{d_D} (\mathbf{o}, \mathbf{a}) \bullet_N^{d_D} (s_3, a \wedge b) \bullet_N^{d_D} (s_2, c \wedge \neg b) \bullet_N^{d_D} (\mathbf{o}, \neg c) \equiv a \wedge b \wedge \neg c$. We receive a new message from the oracle, which contradicts the two messages of s_2 . Hence s_2 become less reliable than s_3 , as only half of s_3 's two messages contradict the oracle. The source s_1 which has never contradicted the oracle is now the single most reliable source. The corresponding reconfiguration sequence is $(\Psi_{\top \circ_N} a \wedge c \circ_N c \wedge \neg b \circ_N \neg a \wedge \neg c \circ_N a \wedge b \circ_N a \wedge \neg c \circ_N b \circ_N \mathbf{a} \wedge \neg c)$.

6 Discussion and conclusion

Reconfiguration operators are a very large family of operators as they have many parameters. It could be interesting to focus on particular subclasses, or to consider variations in the definitions we have given in this paper.

For instance, we have defined the epitemic state Ψ_σ as $(\otimes(m_{r_1}^\varphi \dots m_{r_k}^\varphi)) \circ \varphi_\sigma^\sigma$, i.e. we place the conjunction of all the messages of the oracle at the end of the sequence. Some interesting variations could be for instance $\Psi_\sigma^1 = (\otimes(m_{r_1}^\varphi \dots m_{r_k}^\varphi)) \circ (\otimes\sigma_o)$, in which the sequence of the oracle's messages is considered rather than merely their conjunction; or $\Psi_\sigma^2 = \otimes((m_{r_1}^\varphi \wedge \varphi_o^\sigma) \dots (m_{r_k}^\varphi \wedge \varphi_o^\sigma))$, in which every message from the sources is filtered by the information from the oracle. The latter approach would lead to ignoring all messages that contradict the oracle, as we know those messages to be incorrect. It could however be argued that this is too strong: for instance, if the message of a source is $a \wedge b \wedge c \wedge \dots \wedge z$ and the message of the oracle is $\neg a$, should the entire conjunction be ignored because of the conflict on a ? Studying the properties of these (and other) alternative definitions seems interesting.

With most of the reliability functions we have given, a source is more reliable the more correct messages it provides. This can be justified, since this evaluates how many “proofs” of reliability have been provided. However all of these correct messages can be of very little use if they are not very informative (if a source sends the message “the sky is blue” 50 times, does this make it a reliable source?). Moreover, this feature makes the evaluation weak to certain strategies: a manipulative source could provide many correct, but not very informative, messages, in order to raise its reliability, before sending a deliberately false (but not yet proven incorrect by oracle's messages) piece of information which it wants you to believe. Note that the importance of this issue must be balanced by the fact that such a strategy could be very difficult to carry out, as the malicious agent cannot predict the messages of the oracle, which could, at any moment before or after the planned false message of the source, provide a message that contradicts it. Nevertheless, taking into account the quantity of information a message carries could help avoid this problem.

Another choice we have made, once the reliability function is computed, is to use a standard iterated revision operator. One could instead use some weighted merging operator [17, 6, 7] to aggregate all messages. We claim that, even if reliability should be the most important point, between equally reliable agents, it still makes sense to take recency into account: agents can evolve and correct their beliefs, so we can still expect that, for a single agent or for several agents of the same reliability, more recent information is more correct : sources may learn new things, and correct some of their initial mistakes. Note that this recency could be encoded by adding a second step, after the reliability computation, in order to modify the obtained weights to add information about recency, and then use a weighted merging operator. But this is more naturally taken into account by iterated belief revision operators. Another advantage of iterated revision over merging is that, even though the reliability of the sources is computed numerically, what matters when it comes to the revision process in

our framework is only the order between the sources, in contrast with the more arbitrary numerical weights used in weighted merging.

There are two other closely related works. The first one is [22], which starts from a very similar motivation to ours, but presents several important differences: in [22] the credibility relation is a partial pre-order that is given as input, while we compute our (total) relation from the sequences of messages in a dynamic way. Moreover, in [22] they use the framework of multiple (belief base) revision to take into account messages of same credibility, whereas, as explained previously, iterated revision allows us to also take into account recency of the messages. Similarly to our framework, in [21] reliability of different sources is evaluated based on their announcements and on a special source which is known to be reliable. However their setting is different from ours: they consider a 0/1 notion of expertise of agents on formulas (e.g. “having disease X”), which is evaluated through the agents’ reports across different cases (e.g. patients).

In the future we plan to extend this work in two directions. The first is to consider that the oracle is not perfect, but almost perfect (it makes mistakes much less often than standard sources), and/or that we have several oracles, which may contradict each other. The second direction is to consider this reconfiguration framework, but with no oracle at all. The reliability of each source will then be computed by confronting its messages not with those of the oracle, but with those of the other sources.

Acknowledgements

This work has benefited from the support of the AI Chair BE4musIA of the French National Research Agency (ANR-20-CHIA-0028).

References

1. Besnard, P.: Revisiting postulates for inconsistency measures. In: Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA’14). Lecture Notes in Computer Science, vol. 8761, pp. 383–396. Springer (2014)
2. Booth, R., Meyer, T.A.: Admissible and restrained revision. *Journal of Artificial Intelligence Research* **26**, 127–151 (2006)
3. Booth, R., Fermé, E., Konieczny, S., Pino Pérez, R.: Credibility-limited revision operators in propositional logic. In: Proceedings of the 13th International Conference on the Principles of Knowledge Representation and Reasoning (2012)
4. Booth, R., Fermé, E.L., Konieczny, S., Pino Pérez, R.: Credibility-limited improvement operators. In: Proceedings of the 21st European Conference on Artificial Intelligence (ECAI’14), vol. 263, pp. 123–128 (2014)
5. Darwiche, A., Pearl, J.: On the logic of iterated belief revision. *Artificial Intelligence* **89**(1-2), 1–29 (1997)
6. Delgrande, J.P., Dubois, D., Lang, J.: Iterated revision as prioritized merging. In: Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning (KR’06). pp. 210–220 (2006)
7. Everaere, P., Fellah, C., Konieczny, S., Pérez, R.P.: Weighted merging of propositional belief bases. In: Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning (KR’23) (2023)

8. Fermé, E.L., Mikalef, J., Taboada, J.: Credibility-limited functions for belief bases. *Journal of Logic and Computation* **13**(1), 99–110 (2003)
9. Garapa, M., Fermé, E., Reis, M.D.L.: Credibility-limited base revision: New classes and their characterizations. *Journal of Artificial Intelligence Research* **69**, 1023–1075 (2020)
10. Grant, J., Martinez, M.V.: *Measuring Inconsistency in Information*. College Publications, London (2018)
11. Hansson, S.O., Fermé, E.L., Cantwell, J., Falappa, M.A.: Credibility limited revision. *Journal of Symbolic Logic* **66**(4), 1581–1596 (2001)
12. Hunter, A., Konieczny, S.: Approaches to measuring inconsistent information. In: *Inconsistency tolerance*. Springer LNCS 3300. pp. 189–234 (2005)
13. Jin, Y., Thielscher, M.: Iterated belief revision, revised. *Artificial Intelligence* **171**(1), 1–18 (2007)
14. Konieczny, S., Medina Grespan, M., Pino Pérez, R.: Taxonomy of improvement operators and the problem of minimal change. In: *Proceedings of the 12th International Conference on Principles of Knowledge Representation and Reasoning (KR'10)*. pp. 161–170 (2010)
15. Konieczny, S., Pino Pérez, R.: Improvement operators. In: *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*. pp. 177–187 (2008)
16. Konieczny, S., Pino Pérez, R.: Merging information under constraints: a logical framework. *Journal of Logic and Computation* **12**(5), 773–808 (2002)
17. Lin, J.: Integration of weighted knowledge bases. *Artificial Intelligence* **83**(2), 363–378 (1996)
18. Nayak, A.: Iterated belief change based on epistemic entrenchment. *Erkenntnis* **41**, 353–390 (1994)
19. Schwind, N., Konieczny, S.: Non-prioritized iterated revision: Improvement via incremental belief merging. In: *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR'20)*. pp. 738–747 (2020)
20. Schwind, N., Konieczny, S., Pino Pérez, R.: On the representation of Darwiche and Pearl's epistemic states for iterated belief revision. In: *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning (KR'22)* (2022)
21. Singleton, J., Booth, R.: Who's the expert? on multi-source belief change. In: *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning (KR'22)* (2022)
22. Tamargo, L.H., Deagustini, C.A., García, A.J., Falappa, M.A., Simari, G.R.: Multi-source multiple change on belief bases. *International Journal of Approximate Reasoning* **110**, 145–163 (2019)
23. Thimm, M.: Inconsistency measurement. In: *Proceedings of the 13th International Conference on Scalable Uncertainty Management (SUM'19)*. *Lecture Notes in Computer Science*, vol. 11940, pp. 9–23. Springer (2019)
24. Thimm, M., Wallner, J.P.: On the complexity of inconsistency measurement. *Artificial Intelligence* **275**, 411–456 (2019)