



On Paraconsistent Belief Revision in LP

Nicolas Schwind, Sébastien Konieczny, Ramón Pino Pérez

► To cite this version:

Nicolas Schwind, Sébastien Konieczny, Ramón Pino Pérez. On Paraconsistent Belief Revision in LP. The Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI) 2022, Feb 2022, Virtual conference, United States. pp.5879-5887, 10.1609/aaai.v36i5.20532 . hal-03855356

HAL Id: hal-03855356

<https://hal.science/hal-03855356>

Submitted on 16 Nov 2022

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.

On Paraconsistent Belief Revision in LP

Nicolas Schwind¹, Sébastien Konieczny², Ramón Pino Pérez²

¹ National Institute of Advanced Industrial Science and Technology, Tokyo, Japan

² CRIL - CNRS, Université d'Artois, Lens, France

nicolas-schwind@aist.go.jp, {konieczny,pinoperez}@cril.fr

Abstract

Belief revision aims at incorporating, in a rational way, a new piece of information into the beliefs of an agent. Most works in belief revision suppose a classical logic setting, where the beliefs of the agent are consistent. Moreover, the consistency postulate states that the result of the revision should be consistent if the new piece of information is consistent. But in real applications it may easily happen that (some parts of) the beliefs of the agent are not consistent. In this case then it seems reasonable to use paraconsistent logics to derive sensible conclusions from these inconsistent beliefs. However, in this context, the standard belief revision postulates trivialize the revision process. In this work we discuss how to adapt these postulates when the underlying logic is Priest's LP logic, in order to model a rational change, while being a conservative extension of AGM/KM belief revision. This implies, in particular, to adequately adapt the notion of expansion. We provide a representation theorem and some examples of belief revision operators in this setting.

Introduction

Belief revision aims at incorporating, in a rational way, a new piece of information into the beliefs of an agent. The core of belief change theory (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988; Katsuno and Mendelzon 1991; Hansson 1999; Fermé and Hansson 2011) is well-established now, and the numerous representation theorems, for instance (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988; Katsuno and Mendelzon 1991; Alchourrón and Makinson 1985) as well as the results showing the closeness between belief change and non-monotonic inference (Gärdenfors 1990; Kraus, Lehmann, and Magidor 1990; Lehmann and Magidor 1992) and possibilistic logic (Dubois and Prade 1991; Dubois, Lang, and Prade 1994) confirm that the AGM framework (for Alchourrón, Gärdenfors and Makinson 1985) correctly models this fundamental process.

Nonetheless, some adaptations are required when one leaves the standard classical setting. In particular, one fundamental assumption of the AGM framework is that one works in extensions of classical logic (AGM postulates are stated for any logic satisfying some basic requirements, one

of which is that the consequence relation must contain all classical consequences). If one does not want this to occur then one quickly enters into unknown territories.

To motivate this work, let us recall that the AGM revision postulates aim to formalize three intuitive principles:

- **primacy of update.** The new piece of information must be believed after the change;
- **consistency.** The result of the change has to be a consistent belief base whenever the new piece of information is consistent;
- **minimal change.** We want the result to be as close as possible to the previous beliefs: we do not want to add unnecessary new beliefs and we want to give up only the beliefs that prevent the first two principles from holding.

Minimal change is really at the heart of belief change, as the revised base should be as pertinent as possible. So it is not a principle that can be relaxed.

Primacy of update is a very natural requirement. Nonetheless, in some applications it may be sensible to expect a different behaviour. Sometimes we may want to be given the choice on whether to accept only a part of the new piece of information (Hansson 1998; Makinson 1998; Hansson et al. 2001; Booth et al. 2012; Falappa et al. 2012; Booth et al. 2014; Garapa, Fermé, and Reis 2020). Or we may not want to give such a high priority on the new piece of information with respect to the current beliefs, which could lead us to promotion (Schwind, Konieczny, and Marquis 2018) or improvement (Konieczny, Medina Grespan, and Pino Pérez 2010; Konieczny and Pino Pérez 2008), or, when we want to give the same weight to both, to merging operators (Revesz 1993; Konieczny and Pino Pérez 2002, 2011; Schwind and Konieczny 2020). So relaxations of this principle have been highly investigated.

Contrastingly, the relaxations of the last principle, consistency, have rarely been investigated. We can see several reasons for that. First because, as explained above, the AGM framework requires working with a logic that is an extension of classical logic. So it prevents us from working with paraconsistent logics, that do not contain all classical consequences (in particular when the current belief base is not classically consistent). Another reason could be that adding a new piece of information into the beliefs of the agent when working in a paraconsistent logic could appear simple: as

these logics do not fear logical conflicts, simply performing the conjunction seems to be enough. This is a debatable point, since it forgets the fact that belief revision is about rational change; and the minimal change principle can be interpreted as a requirement to reject the conjunction in this case. This is what this paper is about.

A number of AGM postulates are based on consistency conditions. For instance, in line with the previous point, one of the AGM postulates says that if the conjunction of the belief base and the new piece of information is consistent, then the result of the revision must be exactly this conjunction. This is very sensible in classical logic: if there is no logical conflict caused by the new piece of information we have nothing else to do than adding this piece of information. A direct counterpart of this postulate in paraconsistent logics would require to trivialize the change: we would always have to use the conjunction. So this postulate, and other ones, have to be adapted to be able to cope with paraconsistent logics.

We want to stress that studying belief revision in a paraconsistent logic setting is more than a technical exercise or a purely theoretical question. The AGM setting assumes that the beliefs of the agent are consistent. If it is not the case, then any revision (by a consistent formula) will restore consistency. This is very sensible theoretically. But we want to stress that in real applications this will certainly be the exception rather than the general case. If an agent gathers its beliefs from a lot of different sources (think for instance of an ontology built from gathering information from internet sources), it may very easily happen that some parts of these beliefs are not consistent. In this case, either we can try to repair the inconsistencies, but this will certainly require arbitrary choices and does not really represent the current beliefs of the agent, or we can live with the inconsistency (Gabbay and Hunter 1991; Priest 2002; Béziau, Carnielli, and Gabbay 2007) until we explicitly obtain information allowing to solve these inconsistencies, while being able to derive sensible conclusions from the beliefs of the agent that are not related to these inconsistencies.

As an example, consider that the current beliefs of the agents are $p \wedge \neg p \wedge q \wedge r$ and that the new piece of information is $\neg q$. If someone uses a classical AGM/KM operator¹ then the result will have to be consistent, so learning $\neg q$ will solve the conflict on p . But, without any additional information, it is arguable to consider that these variables can be independent, and that the expected result could be $p \wedge \neg p \wedge \neg q \wedge r$, i.e., that we take into account the new piece of information, but that it does not change anything about the conflict on p . There is no existing operator that allows that, and this is the kind of operators that we want to introduce in this paper.

To do so, we will work in the LP logic setting. LP logic (for Logic of Paradox) (Priest 1979, 1991) is a 3-valued logic, with the third value meaning “inconsistent” (“both true and false”), that allows to isolate inconsistencies in the concerned propositional variables. For instance in a large base (or ontology), we can have several topics (identified by

sets of variables) with inconsistencies on some of these topics, but we want to be able to have non-trivial consequences on topics with no inconsistencies. LP logic allows to do that.

We discuss how to adapt the AGM/KM postulates when the underlying logic is Priest’s LP logic, in order to model a rational change, while being a conservative extension of AGM/KM belief revision. This requires in particular to adequately adapt the definition of expansion, since its direct translation is not adequate for non classical settings. We provide a representation theorem for this class of revision operators in terms of plausibility preorders (faithful assignments) on interpretations. And we define a whole family of distance-based operators, that generalize Dalal revision in this setting. For space reasons the proofs are omitted, but an extended version containing all the proofs is available from <http://www.cril.fr/~konieczny/AAAI22-SKP.pdf>.

Formal Preliminaries

Let \mathcal{L}_{PS} be a propositional language built up from a finite set of propositional variables $PS = \{x_1, \dots, x_n\}$ and the usual connectives. Given $X \subseteq PS$, \bar{X} denotes the set $PS \setminus X$. The symbol \perp_X denotes the formula $\bigwedge_{x_i \in X} x_i \wedge \neg x_i$. Given a formula α , $Var(\alpha)$ denotes the set of propositional variables appearing in α .

An LP world ω is a mapping from PS to $\{0, 1, B\}$ (the value “B” intuitively means “both true and false”). These three truth values are ordered as $0 <_t B <_t 1$. An LP world ω is said to be a *classical world* if for each $x_i \in PS$, $\omega(x_i) \in \{0, 1\}$. The set Ω_* denotes the set of all LP worlds, and the set Ω denotes the set of all classical worlds. The LP semantics of a formula in an LP world ω are defined inductively as follows: $\omega(\neg\alpha) = B$ if $\omega(\alpha) = B$, $\omega(\neg\alpha) = 1$ if $\omega(\alpha) = 0$, $\omega(\neg\alpha) = 0$ if $\omega(\alpha) = 1$, $\omega(\alpha \wedge \beta) = \min(\{\omega(\alpha), \omega(\beta)\}, \leq_t)$, $\omega(\alpha \vee \beta) = \max(\{\omega(\alpha), \omega(\beta)\}, \leq_t)$.

An LP world ω is an *LP model* of a formula α iff $\omega(\alpha) \in \{1, B\}$, and it is said to be a *classical model* of α when ω is a classical world. A formula is said to be *LP-consistent* if it has an LP model; it is said to be *consistent* if it has a classical model. Note that the notion of LP-consistency is a trivial one, since the LP world ω_\perp defined as $\omega_\perp(x_i) = B$ for each $x_i \in PS$ is an LP model of every formula. The set of LP models of a formula α is denoted by $\llbracket \alpha \rrbracket$, and its set of classical models is denoted by $[\alpha]$.

An LP world will be written as a sequence of symbols from $\{0, 1, B\}$, assuming an implicit ordering on PS . For instance, for $PS = \{p, q, r\}$ the LP world ω defined as $\omega(p) = 0$, $\omega(q) = B$ and $\omega(r) = 1$ will simply be denoted by $\omega = 0B1$.

The LP consequence relation, denoted by \models_{LP} , is defined as $\alpha \models_{LP} \beta$ iff $\llbracket \alpha \rrbracket \subseteq \llbracket \beta \rrbracket$. Two formulae α, β are said to be equivalent, denoted by $\alpha \equiv_{LP} \beta$, iff $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket$. Likewise, the classical consequence relation \models is defined as $\alpha \models \beta$ iff $[\alpha] \subseteq [\beta]$, and two formulae are said to be *classically* equivalent, denoted by $\alpha \equiv \beta$, iff $[\alpha] = [\beta]$.

¹The KM framework is a particular case of AGM in the finite propositional case (Katsuno and Mendelzon 1991).

Representative LP Models & LP-DNF

Let us now consider the partial ordering² \leq_{LP} on the set of all LP worlds Ω_* , defined for all $\omega, \omega' \in \Omega_*$ as $\omega \leq_{LP} \omega'$ if and only if for each $x_i \in PS$, $\omega'(x_i) \in \{\omega(x_i), B\}$. Intuitively, $\omega \leq_{LP} \omega'$ could be read as ω' is “less classical” than ω . In particular, it can be seen that the minimal elements of Ω_* with respect to \leq_{LP} form the set of all classical worlds Ω . Now, let us stress that if an LP world ω is an LP model of a formula α , then all of the LP worlds that are “less classical” than ω are also LP models of α . This is formalized in the following lemma:

Lemma 1. *If $\omega \in \llbracket \alpha \rrbracket$ and $\omega \leq_{LP} \omega'$, then $\omega' \in \llbracket \alpha \rrbracket$.*

Given an LP world ω , let us define the formula ψ_ω as $\psi_\omega = \bigwedge_{x_i \in PS} l_i$, where $l_i = x_i$ if $\omega(x_i) = 1$, $l_i = \neg x_i$ if $\omega(x_i) = 0$, otherwise $l_i = x_i \wedge \neg x_i$.

One extends the above definition to a set of LP worlds $S \neq \emptyset$, i.e., ψ_S is defined as $\psi_S = \bigvee_{\omega \in S} \psi_\omega$. The set $\llbracket \psi_S \rrbracket$ is called the *LP closure* of S .

The LP closure of S only contains (i) the LP worlds from S , and (ii) all LP worlds that are “less classical” than some LP world from S (which is a necessary condition according to Lemma 1). This is formalized in the following lemma:

Lemma 2. 1. $\omega \in \llbracket \psi_S \rrbracket$
 2. If $\omega \in \llbracket \psi_S \rrbracket$, then $\exists \omega' \in S$ s.t. $\omega' \leq_{LP} \omega$
 3. If $S \subseteq \llbracket \alpha \rrbracket$, then $\psi_S \models_{LP} \alpha$

The LP closure also satisfies the standard closure properties of extensivity, monotonicity, and idempotence:

Lemma 3. 1. $S \subseteq \llbracket \psi_S \rrbracket$
 2. If $S \subseteq T$, then $\psi_S \models_{LP} \psi_T$
 3. $\psi_S \equiv_{LP} \psi_{\llbracket \psi_S \rrbracket}$

Now, with every formula one can associate a representative subset of its set of LP models that will be of particular importance in the rest of this paper.

Definition 1. *Given a formula α , the representative set of LP models of α , denoted by $\llbracket \alpha \rrbracket_*$, is defined as $\llbracket \alpha \rrbracket_* = \min(\llbracket \alpha \rrbracket, \leq_{LP})$.*

Let us introduce a few remarkable properties on representative LP models:

Lemma 4. 1. $\alpha \equiv_{LP} \psi_{\llbracket \alpha \rrbracket_*}$
 2. If $S \subsetneq \llbracket \alpha \rrbracket_*$, then $\alpha \not\models_{LP} \psi_S$

The first point states that the representative set of LP models of α is *sufficient* to retrieve all LP models of α . And the second point says that each LP model from the representative set of α is *necessary* to retrieve all LP models of α .

These two points together mean that each formula can be characterized (up to equivalence) by its representative set. Given that the LP worlds from any representative set are all pairwise incomparable w.r.t. \leq_{LP} (see Definition 1), this characterization result can be formally expressed as follows: Let $(2^{\Omega_*})_\sim$ denote the set of all subsets of LP worlds that are pairwise incomparable w.r.t. \leq_{LP} , i.e., $(2^{\Omega_*})_\sim = \{S \subseteq \Omega_* \mid \forall \omega, \omega' \in S, \omega \leq_{LP} \omega' \implies \omega = \omega'\}$:

²Given a partial ordering \leq , the symbols $<$ and \simeq denote respectively the corresponding strict and equivalence relations. When the ordering is denoted by \preceq , its strict counterpart is denoted by \prec .

Corollary 1. *For any formula α , there is a unique set $S_\alpha \in (2^{\Omega_*})_\sim$ such that $\alpha \equiv_{LP} \psi_{S_\alpha}$. Moreover, $S_\alpha = \llbracket \alpha \rrbracket_*$.*

Corollary 1 says that the representative set of a formula α is the *only* set from $(2^{\Omega_*})_\sim$ which can characterize α (up to equivalence). The other way around, one can also see that every set of LP worlds $S \in (2^{\Omega_*})_\sim$ is the representative set of some formula:

Corollary 2. *For each $S \in (2^{\Omega_*})_\sim$, there exists a formula α such that $\llbracket \alpha \rrbracket_* = S$.*

So Corollaries 1 and 2 make clear the one-to-one correspondence between the set of all formulae (up to equivalence) and their representative sets from $(2^{\Omega_*})_\sim$. Since any formula α can be represented equivalently in a canonical form characterized by its representative set, in what follows, given any formula α , the formula $\psi_{\llbracket \alpha \rrbracket_*} \equiv_{LP} \alpha$ will be called the *LP Disjunctive Normal Form (LP-DNF)* of α . And $lpdnf(\alpha)$ will denote the set $\{\psi_\omega \mid \omega \in \llbracket \alpha \rrbracket_*\}$. Accordingly, we have that $\alpha \equiv_{LP} \bigvee_{\psi_\omega \in lpdnf(\alpha)} \psi_\omega$.

Issues with KM Postulates in LP Revision

A revision operator \circ associates all formulae φ, μ with a formula $\varphi \circ \mu$. A set of standard properties are expected in the classical case. Let us recall the standard KM postulates:

Definition 2. (Katsuno and Mendelzon 1991) *A revision operator \circ is said to be a KM revision operator if it satisfies the following postulates:*

- (R1) $\varphi \circ \mu \models \mu$
- (R2) If $\varphi \wedge \mu$ is consistent, then $\varphi \wedge \mu \equiv \varphi \circ \mu$
- (R3) If μ is consistent, then $\varphi \circ \mu$ is consistent
- (R4) If $\varphi \equiv \varphi'$ and $\mu \equiv \mu'$, then $\varphi \circ \mu \equiv \varphi' \circ \mu'$
- (R5) $(\varphi \circ \mu) \wedge \mu' \models \varphi \circ (\mu \wedge \mu')$
- (R6) If $(\varphi \circ \mu) \wedge \mu'$ is consistent, then $\varphi \circ (\mu \wedge \mu') \models (\varphi \circ \mu) \wedge \mu'$

Let us recall that the notion of consistency referred to in these postulates corresponds to the notion of *classical* consistency, i.e., it refers to formulae which admit at least one classical model.

KM revision operators can be characterized in terms of faithful assignments:

Definition 3. *A faithful assignment (denoted by $\varphi \mapsto \preceq_\varphi$) is a mapping associating every formula φ with a preorder \preceq_φ over classical worlds, such that:*

1. $\forall \omega, \omega' \in \Omega$, if $\omega, \omega' \in [\varphi]$ then $\omega \simeq_\varphi \omega'$
2. $\forall \omega, \omega' \in \Omega$, if $\omega \in [\varphi]$ and $\omega' \notin [\varphi]$, then $\omega \prec_\varphi \omega'$
3. If $[\varphi] = [\varphi']$, then $\preceq_\varphi = \preceq_{\varphi'}$

Proposition 1. (Katsuno and Mendelzon 1991) *A revision operator \circ is a KM revision operator if and only if there exists a faithful assignment $\varphi \mapsto \preceq_\varphi$ associating every formula φ with a total preorder \preceq_φ over classical worlds, such that for all formulae φ, μ , $[\varphi \circ \mu] = \min([\mu], \preceq_\varphi)$.*

Our goal is to define interesting revision operators in the LP setting, and in order to do so, our first step is to look at the KM postulates with the LP semantics in mind. Yet in the paraconsistent case the notion of LP-consistency is a trivial one, i.e., each formula is LP-consistent since the world ω_\perp defined as $\omega(x_i) = B$ for each $x_i \in PS$ is an LP model

of every formula. So, if one interprets the KM postulates in the LP setting, postulate **(R3)** becomes trivially true and the remaining KM postulates can be rephrased as follows:

- (R1-LP)** $\varphi \circ \mu \models_{LP} \mu$
- (R2-LP)** $\varphi \wedge \mu \equiv_{LP} \varphi \circ \mu$
- (R4-LP)** If $\varphi \equiv_{LP} \varphi'$ and $\mu \equiv_{LP} \mu'$, then $\varphi \circ \mu \equiv_{LP} \varphi' \circ \mu'$
- (R5-LP)** $(\varphi \circ \mu) \wedge \mu' \models_{LP} \varphi \circ (\mu \wedge \mu')$
- (R6-LP)** $\varphi \circ (\mu \wedge \mu') \models_{LP} (\varphi \circ \mu) \wedge \mu'$

Postulates **(R1-LP)** and **(R4-LP)** express requirements that are similar to the classical case and can be seen as natural extensions of **(R1)** and **(R4)** to the LP setting. As to the remaining postulates, let us first discuss **(R2-LP)**. This postulate is clearly too strong since it forces the revision to trivialize the conjunction. So one must consider a weakening of **(R2-LP)**. One of the least demanding weakenings is to ask that the agent's current beliefs should not change when they are revised by a tautology (see, e.g., (Benferhat, Lagrue, and Papini 2005)):

- (R2-LP')** $\varphi \circ \top \equiv_{LP} \varphi$

Requiring **(R2-LP')** is very natural even in the paraconsistent case: there is no reason for changing some agent's beliefs φ when the new information is tautological, whether φ is consistent or not. But then, it turns out that the postulates **(R5-LP)** and **(R6-LP)** again trivialize the revision to be the conjunction in presence of **(R2-LP')**:

Proposition 2. *If \circ satisfies **(R2-LP')**, **(R5-LP)** and **(R6-LP)**, then it satisfies **(R2-LP)**.*

What is Expansion?

The previous results illustrate the trivialization caused by a naive interpretation of KM revision postulates in LP. Most of these problems are caused by the use of LP conjunction. In particular, as LP conjunction is always LP-consistent, then **(R2-LP)** trivializes the revision to a simple conjunction in LP. As explained in the introduction, we want to go further than that, and we motivated the fact that revision is not just about consistent change, but about rational change.

To solve this trivialization issue, one has to come back to the AGM definition of expansion, revision and contraction (Alchourrón, Gärdenfors, and Makinson 1985), and to remember that the addition of a formula into a belief set has to be done using expansion, and that revision is defined using expansion. It turns out that AGM showed that, in the classical setting, the only expansion operator is logical conjunction. But this is not necessarily the case in other, non-classical settings, like ours. Expansion basically aims at acquiring new information when a new evidence does not raise conflicts in the beliefs of the agent. Interpreted in terms of possible worlds, this is done by selecting the possible worlds of the beliefs of the agent that satisfy the new evidence.

This makes perfect sense in the AGM “coherentist” setting, where all the beliefs, so all the possible worlds, have the same status. But in LP, the representative set of models of a formula are more important than the others. For instance, given $\varphi = p \wedge q$, we get that $\llbracket \varphi \rrbracket = \{11, 1B, B1, BB\}$ and $\llbracket \varphi \rrbracket_\star = \{11\}$, so the LP model 11 is more informative than the others which are only there because of 11.

So, expansion in this setting should take this dimension into account, and make a selection among the LP worlds from the representative set. Let us put it formally. Recall that the LP-DNF formula which is equivalent to a formula α is $\psi_{\llbracket \alpha \rrbracket_\star}$, and that $lpdnf(\alpha) = \{\psi_\omega \mid \omega \in \llbracket \alpha \rrbracket_\star\}$. Then given two formulae α and β , let us denote by $lpdnf(\alpha \mid \beta)$ the subset of $lpdnf(\alpha)$ defined as $lpdnf(\alpha \mid \beta) = \{\psi_\omega \in lpdnf(\alpha) \mid \psi_\omega \models_{LP} \beta\}$. Let us simply denote by $\bigvee lpdnf(\alpha \mid \beta)$ the formula $\bigvee \{\psi_\omega \mid \psi_\omega \in lpdnf(\alpha \mid \beta)\}$. From the semantical point of view, the representative set of the formula $\bigvee lpdnf(\alpha \mid \beta)$ (when defined) precisely corresponds to the LP models in the representative set of α which are LP models of β .

Proposition 3. *Assume that $lpdnf(\alpha \mid \beta) \neq \emptyset$. Then $\llbracket \bigvee lpdnf(\alpha \mid \beta) \rrbracket_\star = \llbracket \alpha \rrbracket_\star \cap \llbracket \beta \rrbracket_\star$.*

And as a direct consequence of the definition of $lpdnf(\alpha \mid \beta)$ and Proposition 3, we get:

Corollary 3. *$lpdnf(\alpha \mid \beta) \neq \emptyset$ iff $\llbracket \alpha \rrbracket_\star \cap \llbracket \beta \rrbracket_\star \neq \emptyset$.*

We are now ready to define our expansion operator, which we call LP+expansion:

Definition 4. *The LP+expansion of α by β , denoted by $\alpha +_{LP} \beta$ is defined as*

$$\alpha +_{LP} \beta = \begin{cases} \bigvee lpdnf(\alpha \mid \beta), & \text{if } lpdnf(\alpha \mid \beta) \neq \emptyset, \\ \perp_{PS}, & \text{otherwise.} \end{cases}$$

When $lpdnf(\alpha \mid \beta) \neq \emptyset$ we say that $\alpha +_{LP} \beta$ is conclusive.

Example 1. *Let us illustrate this difference in behaviour between conjunction and LP+expansion with a simple example. Consider a formula α such that $\llbracket \alpha \rrbracket_\star = \{01BB, 1111\}$. Now $\llbracket \alpha +_{LP} \psi_{\{0111\}} \rrbracket = \llbracket \psi_{\{01BB\}} \rrbracket$, whereas $\llbracket \alpha \wedge \psi_{\{0111\}} \rrbracket = \llbracket \psi_{\{01BB, B111\}} \rrbracket$. And $\llbracket \alpha +_{LP} \psi_{\{0BBB\}} \rrbracket = \llbracket \perp_{PS} \rrbracket$, whereas $\llbracket \alpha \wedge \psi_{\{0BBB\}} \rrbracket = \llbracket \psi_{\{0BBB\}} \rrbracket$.*

As we will have to deal with representative worlds, we have to define a (strong) inference relation between formulae that takes into account these worlds only.

Definition 5. *The strong LP inference relation, denoted by $\models_{LP\star}$, is the relation on $\mathcal{L}_{PS} \times \mathcal{L}_{PS}$ defined for all formulae α, β as: $\alpha \models_{LP\star} \beta$ iff $lpdnf(\alpha) \subseteq lpdnf(\beta)$.*

Obviously enough, this inference relation $\models_{LP\star}$ can be semantically characterized in terms of inclusion between representative sets:

Remark 1. *$\alpha \models_{LP\star} \beta$ iff $\llbracket \alpha \rrbracket_\star \subseteq \llbracket \beta \rrbracket_\star$.*

We are now ready to give a translation of the AGM expansion postulates in this setting. In the following postulates we will use strong LP inference to compare the belief bases of the agent, in order to be able to focus on the conservation of representative worlds. Whereas we use (standard) LP inference to look at (all) the consequences of these bases.

Definition 6. *An operator $+$: $\mathcal{L}_{PS} \times \mathcal{L}_{PS} \rightarrow \mathcal{L}_{PS}$ is an LP-expansion operator if it satisfies the following properties:*

- (K+2)** $\varphi + \alpha \models_{LP} \alpha$
- (K+3)** $\varphi + \alpha \models_{LP\star} \varphi$
- (K+4)** If $\varphi \models_{LP} \alpha$, then $\varphi + \alpha \equiv_{LP} \varphi$

(K+5) If $\varphi \models_{LP^*} \psi$, then $\varphi + \alpha \models_{LP^*} \psi + \alpha$
(K+6) If $+$ satisfies (K+2)-(K+5), then $\varphi +' \alpha \models_{LP} \varphi + \alpha$

These postulates are an adaptation of the AGM expansion postulates for this setting where some LP models are more important than others.

As in the standard AGM framework there is a unique operator satisfying these postulates:

Proposition 4. *The LP -expansion $+_{LP}$ is the only LP -expansion operator.*

Reasonable LP Revision

Based on the notion of LP -expansion, which we denote simply by $+$ in the following (since it is unique), we propose the following new set of postulates for LP revision:

Definition 7. *A revision operator \circ is an LP revision operator if it satisfies the following postulates:*

- (LP1) $\varphi \circ \mu \models_{LP} \mu$
- (LP2) If $\varphi + \mu$ is conclusive, then $\varphi \circ \mu \equiv_{LP} \varphi + \mu$
- (LP4) If $\varphi \equiv_{LP} \varphi'$ and $\mu \equiv_{LP} \mu'$, then $\varphi \circ \mu \equiv_{LP} \varphi' \circ \mu'$
- (LP5) $(\varphi \circ \mu) + \mu' \models_{LP} \varphi \circ (\mu \wedge \mu')$
- (LP6) If $(\varphi \circ \mu) + \mu'$ is conclusive, then $\varphi \circ (\mu \wedge \mu') \models_{LP} (\varphi \circ \mu) + \mu'$

These postulates are similar to the original KM ones, except that we use the correct notion of expansion in this setting. We did not add a translation to (R3) since as we previously explained, its direct translation is trivially true in LP . But we will discuss an adaptation of (R3) in the next section. Let us now give a representation theorem for these operators.

Definition 8. *An LP faithful assignment is a mapping $\varphi \mapsto \preceq_\varphi$ which associates every formula φ with a preorder over LP worlds and which satisfies the following conditions:*

- (L1) If $\omega, \omega' \in \llbracket \varphi \rrbracket_*$, then $\omega \simeq_\varphi \omega'$
- (L2) If $\omega \in \llbracket \varphi \rrbracket_*$ and $\omega' \notin \llbracket \varphi \rrbracket_*$, then $\omega \prec_\varphi \omega'$
- (L3) If $\varphi \equiv_{LP} \varphi'$ then $\preceq_\varphi = \preceq_{\varphi'}$

Proposition 5. *An operator \circ is an LP revision operator if and only if there is an LP faithful assignment $\varphi \mapsto \preceq_\varphi$ associating every formula with a total preorder over LP worlds such that for all formulae $\varphi, \mu, \varphi \circ \mu \equiv_{LP} \psi_{\min(\llbracket \mu \rrbracket, \preceq_\varphi)}$.*

So, as in the classical case, there is a whole family of sensible revision operators, which can be defined as a selection of the most plausible (LP) possible worlds.

Adaptation of (R3)

(R3) requires that the result of the revision has to be consistent whenever the new piece of information is consistent. As all formulae are LP -consistent, a direct translation of (R3) into our LP setting would always be trivially satisfied. Nonetheless, one could expect the fact that if the new piece of information is consistent, then the result should be consistent. In fact, such a property would be both too strong and too weak. So we will look for a more adequate version of it.

To see that it is too strong, one has to realize that two formulae that are equivalent in classical logic are not always so in LP . In the following example, we show that revising a consistent base by a consistent new information does not lead to a consistent result.

Example 2. Let $\alpha_1 = p \wedge q$ and $\alpha_2 = p \wedge (\neg p \vee q)$. Remark that in classical logic these two formulae are equivalent. This is no longer the case in LP , since $\llbracket \alpha_1 \rrbracket_* = \{11\}$, whereas $\llbracket \alpha_2 \rrbracket_* = \{11, B0\}$. Indeed, in LP logic there is a possible LP world in α_2 where the variable p is “inconsistent” and q is false (i.e., the LP world $B0$), which is not an LP model of α_1 . Then it can be perfectly sensible to expect the revision by $\neg q$ to give different results in these two cases, for instance obtaining $\llbracket \alpha_1 \circ \neg q \rrbracket_* = \{10\}$ and $\llbracket \alpha_2 \circ \neg q \rrbracket_* = \{B0\}$. Remark that $\alpha_2 \circ \neg q$ is not consistent, whereas both α_2 and $\neg q$ are.

On the other hand, (R3) focuses on (fully) consistent cases (i.e., where the models are classical ones). But consider for instance a (almost consistent) case with ten propositional variables, with only one variable having a B value in all (representative) models of the base and the new information, all the other ones having classical truth values. Then we could expect an (R3)-like postulate to force those other variables to remain classical. It is this property that we intend to formalize now.

Given an LP world $\omega \in \Omega_*$, we denote by $\omega!$ the set of “inconsistent” variables in ω , i.e., $\omega! = \{x_i \in PS \mid \omega(x_i) = B\}$. Given $X \subseteq PS$, a formula α is said to be strongly X -consistent whenever $\omega! \subseteq X$ for all ω such that $\omega \in \llbracket \alpha \rrbracket_*$.

Obviously enough, strong PS -consistency implies (classical) consistency. The converse statement does not hold, e.g., the formula $\alpha = p \wedge (\neg p \vee q)$ is consistent, but $\llbracket \alpha \rrbracket_* = \{11, B0\}$, thus α is not strongly PS -consistent.

Definition 9. *An LP revision operator \circ is said to be strong if it satisfies the following additional postulate:*

- (LP3) If φ is strongly X -consistent and μ is strongly X -consistent, then $\varphi \circ \mu$ is strongly X -consistent

This postulate ensures that if some propositional variables have consistent truth values in the (representative set of the) beliefs and in the (representative set of the) new information, then they will have consistent truth values in the (representative set of the) revised beliefs.

Definition 10. *An LP faithful assignment $\varphi \mapsto \preceq_\varphi$ is said to be strong if it satisfies the following additional condition:*

- (S) If φ is strongly X -consistent, $\omega! \subseteq X$, $\omega' \cap X \neq \emptyset$ and $\omega \prec_{LP} \omega'$, then there exists ω'' s.t. $\omega''! \subseteq X$, $\omega \leq_{LP} \omega''$, $\omega'' \preceq_\varphi \omega'$, and $(\omega'' \not\leq_{LP} \omega' \implies \omega'' \prec_\varphi \omega')$

This condition states that if the beliefs are strongly X -consistent, then any possible world that does not comply with X (i.e., $\exists x \in X$ s.t. $\omega(x) = B$) is worse than some more classical possible world that complies with X .

We can now state the representation theorem for strong LP revision operators:

Proposition 6. *An operator \circ is a strong LP revision operator if and only if there is a strong LP faithful assignment $\varphi \mapsto \preceq_\varphi$ associating every formula with a total preorder over LP worlds such that for all formulae $\varphi, \mu, \varphi \circ \mu \equiv_{LP} \psi_{\min(\llbracket \mu \rrbracket, \preceq_\varphi)}$.*

Condition (S) of a strong LP faithful assignment characterizes (LP3) for LP revision operators. But as it may be difficult to check, it can be useful to consider a simpler (stronger) version of this condition:

(C) If φ is strongly X -consistent, $\omega! \subseteq \bar{X}$ and $\omega'! \cap X \neq \emptyset$, then $\omega \preceq_{\varphi} \omega'$

This condition states that if the beliefs are strongly X -consistent, then all possible worlds that do not comply with X are worse than all worlds that comply with X .

Proposition 7. If $\varphi \mapsto \preceq_{\varphi}$ is an LP faithful assignment which satisfies (C), then it satisfies (S).

Let us show that strong LP revision operators can be seen as an extension of classical AGM/KM revision operators.

First, let us recall that strong PS -consistency implies (classical) consistency, but not conversely. However, the DNF of every consistent formula α is a formula which is classically equivalent to α and is strongly PS -consistent³:

Proposition 8. If α is consistent, then $\text{dnf}(\alpha)$ is strongly PS -consistent.

Notice that for classical revision operators we have that $\varphi \circ \alpha \equiv \text{dnf}(\varphi) \circ \alpha$. This is a consequence of (R4).

Now we can show that:

Proposition 9. For every KM revision operator \circ , there exists a strong LP revision operator \otimes such that for every consistent formula φ and every formula μ , we have that $\varphi \circ \mu \equiv \text{dnf}(\varphi) \otimes \mu$.

So, the class of strong LP revision operators can be seen as a “safe” extension of classical AGM/KM revision operators. Some strong LP revision operators behave exactly as classical AGM/KM revision on consistent formulae, while giving non-trivial results on inconsistent formulae.

Distance-Based LP Revision Operators

Let us now investigate how to generalize Dalal revision (Dalal 1988), and more generally, distance-based revision operators into this LP logic setting.

Definition 11. Let d be a mapping $\Omega_{\star} \times \Omega_{\star} \mapsto \mathbb{R}^+$ such that $d(\omega, \omega') = 0$ if and only if $\omega = \omega'$. The distance-based LP revision operator \circ^{d_b} is defined for all formulae φ, μ by $\varphi \circ^{d_b} \mu = \psi_{\min(\|\mu\|, \preceq_{\varphi}^d)}$, where \preceq_{φ}^d is a total preorder over LP worlds induced by φ and defined by:

- $\omega \preceq_{\varphi}^d \omega'$ if and only if $d(\omega, \varphi) \leq d(\omega', \varphi)$
- $d(\omega, \varphi) = \min_{\omega' \in \|\varphi\|_{\star}} d(\omega, \omega')$

Note that the distance is computed from the representative models of the base. Moreover, d is not required to satisfy the properties of symmetry and triangular inequality of standard metrics. For convenience, such a mapping d will be called a distance in the following, since it generalizes well-known distance-based operators.

Proposition 10. Every distance-based LP revision operator is an LP revision operator, i.e., it satisfies (LP1), (LP2), (LP4), (LP5), and (LP6).

One of the most natural ways to define a distance d is to take advantage of existing distances in the classical case, like the Hamming distance d^H , defined as

$d^H(\omega, \omega') = \sum_{x \in PS} d_b^H(\omega(x), \omega'(x))$, where d_b^H is defined by $d_b^H(\omega(x), \omega'(x)) = 0$ if $\omega(x) = \omega'(x)$, otherwise $d_b^H(\omega(x), \omega'(x)) = 1$. In the classical case, this distance induces a revision operator which corresponds to the Dalal revision operator (Dalal 1988)⁴. We will study similar distance definitions in our LP setting, so let us put it formally:

Definition 12. A decomposable distance between LP worlds is a mapping $\Omega_{\star} \times \Omega_{\star} \mapsto \mathbb{R}^+$ defined for all LP worlds ω, ω' as $d(\omega, \omega') = \sum_{x \in PS} d_b(\omega(x), \omega'(x))$, where d_b is a mapping $\{0, 1, B\}^2 \mapsto \mathbb{R}^+$ satisfying $d_b(\omega(x), \omega'(x)) = 0$ if and only if $\omega(x) = \omega'(x)$.

Since a decomposable distance d is fully characterized by d_b , one denotes by \circ^{d_b} the induced LP distance-based revision operator. It can be easily seen that every operator \circ^{d_b} is indeed a distance-based LP operator according to Definition 11. Let us see what kind of properties could make sense for the d_b mapping.

First, let us remark that symmetry is not a natural requirement in our setting. For instance, one may want in some cases to favor classical truth values (true and false) over the inconsistent one, e.g., to ask that $d_b(0, B) \neq d_b(B, 0)$.

But a sensible requirement is to be neutral with respect to the classical values true and false, i.e., to decomposable distances that are “0-1 symmetric”, i.e., such that $d_b(0, B) = d_b(1, B)$, $d_b(B, 0) = d_b(B, 1)$, and $d_b(0, 1) = d_b(1, 0)$.

Now, since in the classical case 2-valued worlds are considered, to define an operator based on a 0-1 symmetric decomposable distance the only reasonable choice for d_b is to define $d_b(0, 1) = d_b(1, 0) = 1$, which defines the Dalal operator (taking any positive value other than 1 defines the same operator). But in our setting, with 3-valued interpretations, there is more room for non-trivial choices. Still, by setting the distance between the “classical” values to the reference value 1 (i.e., $d_b(0, 1) = d_b(1, 0) = 1$) and since $d_b(\omega(x), \omega(x)) = 0$ by definition, we already require that:

- $d_b(0, 0) = d_b(1, 1) = d_b(B, B) = 0$
- $d_b(0, 1) = d_b(1, 0) = 1$

So one is left with a choice on the remaining values:

- $d_b(0, B) = d_b(1, B) = d_{01 \leftarrow B}$
- $d_b(B, 0) = d_b(B, 1) = d_{B \leftarrow 01}$

But it turns out that the choice for $d_{01 \leftarrow B}$ does not matter:

Proposition 11. Let d_b, d'_b be such that $d_{B \leftarrow 01} = d'_{B \leftarrow 01}$. Then $\circ^{d_b} = \circ^{d'_b}$.

So an operator \circ^{d_b} based on a 0-1 symmetric decomposable distance is in fact characterized by a single value: the value $d_{B \leftarrow 01}$.

We know from Proposition 10 that whatever the choice of this value, the obtained operator will satisfy all required postulates. But the exact value will determine the exact behaviour of the operator.

In particular the choice of $d_{B \leftarrow 01} < 1$, so $d_{B \leftarrow 01} < d_b(0, 1) = d_b(1, 0)$, will give rise to operators that always

⁴In his paper (Dalal 1988), Dalal defines his operator by means of formula dilation (Bloch and Lang 2002), but it can also be defined using the Hamming distance between worlds.

³Note that for a consistent formula α , $\text{dnf}(\alpha) = \psi_{[\alpha]}$.

prefer to change the valuation of a propositional variable to B instead of the classical truth values 0 and 1.

Conversely the choice of $d_{B \leftarrow 01} > |PS|$ will force the best possible worlds to be the “most classical ones” (since the price to pay for a B value will be too high).

And choosing an intermediate value ($1 \leq d_{B \leftarrow 01} \leq |PS|$) gives rise to operators that more or less favor classical truth values with respect to the B value, so one can choose the cost of not being consistent.

Now let us show that there is only one operator with $d_{B \leftarrow 01} < 1$, and similarly there is a unique operator with $d_{B \leftarrow 01} > |PS|$:

Proposition 12.

1. Let d_b, d'_b be such that $d_{B \leftarrow 01} < 1$ and $d'_{B \leftarrow 01} < 1$. Then $\circ^{d_b} = \circ^{d'_b}$.
2. Let d_b, d'_b be such that $d_{B \leftarrow 01} > |PS|$ and $d'_{B \leftarrow 01} > |PS|$. Then $\circ^{d_b} = \circ^{d'_b}$.

Since values $d_{B \leftarrow 01} < 1$ define the same operator, let us denote it by $\circ^{d_{\inf}}$. Likewise, the values $d_{B \leftarrow 01} > |PS|$ define the operator denoted by $\circ^{d_{\sup}}$.

So, intuitively, when defining an operator, one would rather choose a lower value for $d_{B \leftarrow 01}$ when one is more reluctant to change, or stated equivalently, more tolerant to inconsistencies. In the most change-reluctant cases, in particular for $\circ^{d_{\inf}}$, one would expect the underlying operators not to be inclined to any change at all, giving as output a revised formula $\varphi \circ \mu$ which always entails φ . Note that whereas this is not a desired behaviour in the classical setting where consistency must be preserved in all cases (cf. (R3)), this property is perfectly acceptable (although not necessary) in our LP setting. Let us call this property “persistence”:

(Per) $\varphi \circ \mu \models_{LP} \varphi$

It turns out that, among the class of all operators based on a 0-1 symmetric decomposable distance, the operator $\circ^{d_{\inf}}$ is the only operator satisfying (Per):

Proposition 13. An operator \circ^{d_b} based on a 0-1 symmetric decomposable distance satisfies (Per) if and only if $\circ^{d_b} = \circ^{d_{\inf}}$.

But, as our work shows, this is the simplest of all possibilities. And in fact all the remaining operators \circ^{d_b} , i.e., when $d_{B \leftarrow 01} \geq 1$ are strong LP revision operators:

Proposition 14. An operator \circ^{d_b} based on a 0-1 symmetric decomposable distance satisfies (LP3) if and only if $d_{B \leftarrow 01} \geq 1$.

Actually, the operator $\circ^{d_{\sup}}$ can be characterized by an additional interesting property:

(LP3') If φ is strongly PS-consistent and μ is consistent, then $\varphi \circ \mu$ is strongly PS-consistent

Proposition 15. An operator \circ^{d_b} based on a 0-1 symmetric decomposable distance satisfies (LP3') if and only if $\circ^{d_b} = \circ^{d_{\sup}}$.

So the operator $\circ^{d_{\sup}}$ forces the result of the revision to be consistent whenever it is possible.

The postulates (LP3') and (LP3) are two adaptations of (R3) in our framework. Note that there can be other interesting ones.

Let us illustrate the (differences in) behaviour of these operators on the following example:

Example 3. Consider first the following worlds:

$\omega_\varphi = 000000000000$ $\omega_1 = 0000000BBBBB$
 $\omega_2 = 1000000BBBB0$ $\omega_3 = 1110000BB000$
 $\omega_4 = 111111100000$ $\omega_5 = B000000BBBB0$
 $\omega_6 = BBB0000BB000$

Now, let φ and μ be two formulae such that $\llbracket \varphi \rrbracket = \{\omega_\varphi\}$ and $\llbracket \mu \rrbracket = \{\omega_1, \omega_2, \omega_3, \omega_4\}$. Let $d_{B \leftarrow 01}^1 = 1$ and $d_{B \leftarrow 01}^2 = 2$. We get that $\varphi \circ^{d_{\inf}} \mu = \{\omega_1, \omega_5, \omega_6\}$, $\varphi \circ^{d_{B \leftarrow 01}^1} \mu = \{\omega_1, \omega_2, \omega_3\}$, $\varphi \circ^{d_{B \leftarrow 01}^2} \mu = \{\omega_3, \omega_4\}$, and $\varphi \circ^{d_{\sup}} \mu = \{\omega_4\}$.

Accordingly, the higher the value of $d_{B \leftarrow 01}$, the more classical the result of the revision of φ by μ .

Related Work

There is a few related work on revision on paraconsistent logics. The first approach we are aware of is (da Costa and Bueno 1998), which contains an extensive discussion on why belief revision in paraconsistent logics makes perfect sense. It uses da Costa C_n logics, but does not go further than discussing the standard AGM postulates in this setting. In (Mares 2002) the author works with a dedicated logic **R**, where an additional structure is required to guide the revision process, and no link is made with the standard approach. In (Priest 2001) the proposed operator used a complex process using different criteria to rank all bases considered as potential solutions. This paper also discussed which are the AGM postulates satisfied by this operator. In (Girard and Tanaka 2016) the authors use LP logic, like us, but they defined two particular operators, corresponding to Segerberg's irrevocable revision (Segerberg 1998) and to Nayak's lexicographic revision (Nayak 1994) on 3-valued interpretations. Their operators are defined using dynamic epistemic logic, and they do not study the links with AGM postulates. Note that none of these works provide a representation theorem, and none of them safely extend the classical AGM/KM framework. The only two works that we are aware of where the authors provide representation theorems are (Testa, Coniglio, and Ribeiro 2017) and (Testa et al. 2018). But in both cases they use a logic with an explicit consistency connector, and the theorems are set for the basic postulates only, not for the full set of AGM postulates.

Conclusion

In this work we discussed how to adapt the AGM/KM postulates when the underlying logic is Priest's LP logic, in order to model a rational change, as a conservative extension of AGM/KM belief revision. This implied in particular to adequately adapt the notion of expansion. We provided a representation theorem and some examples of belief revision operators in this setting.

We hope that this work will allow to rethink the definition of revision as rational change instead of as consistent change, and will be applied to frameworks where some beliefs are more important/fundamental than others.

Acknowledgements

This work has benefited from the support of the AI Chair BE4musIA of the French National Research Agency (ANR-20-CHIA-0028) and of the JSPS KAKENHI Grant Number JP20K11947. The third author has also been partially funded by the program PAUSE of Collège de France.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50: 510–530.
- Alchourrón, C. E.; and Makinson, D. 1985. On the logic of theory change: safe contraction. *Studia Logica*, 44(4): 405–422.
- Benferhat, S.; Lagrue, S.; and Papini, O. 2005. Revision of partially ordered information: axiomatization, semantics and iteration. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI'05)*, 376–381.
- Bloch, I.; and Lang, J. 2002. *Towards mathematical morpho-logics*, 367–380. Heidelberg: Physica-Verlag HD.
- Booth, R.; Fermé, E.; Konieczny, S.; and Pino Pérez, R. 2012. Credibility limited revision operators in propositional logic. In *Proceedings of the 13th International Conference on Principles of Knowledge Representation and Reasoning (KR'12)*, 116–125.
- Booth, R.; Fermé, E.; Konieczny, S.; and Pino Pérez, R. 2014. Credibility-limited improvement operators. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, 123–128.
- Béziau, J.-Y.; Carnielli, W. A.; and Gabbay, D. M. 2007. Handbook of paraconsistency. Studies in logic. College publications.
- da Costa, N. C.; and Bueno, O. 1998. Belief change and inconsistency. *Logique et Analyse*, 41(161/163): 31–56.
- Dalal, M. 1988. Investigations into a theory of knowledge base revision: preliminary report. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'88)*, 475–479.
- Dubois, D.; Lang, J.; and Prade, H. 1994. *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3 : Nonmonotonic Reasoning and Uncertain Reasoning, chapter Possibilistic Logic. Oxford Science Publications.
- Dubois, D.; and Prade, H. 1991. Epistemic entrenchment and possibilistic logic. *Artificial Intelligence*, 50: 223–239.
- Falappa, M. A.; Kern-Isberner, G.; Reis, M. D. L.; and Simari, G. R. 2012. Prioritized and non-prioritized multiple change on belief bases. *Journal of Philosophical Logic*, 41(1): 77–113.
- Fermé, E.; and Hansson, S. O. 2011. AGM 25 years: twenty-five years of research in belief change. *Journal of Philosophical Logic*, 40(2): 295–331.
- Gabbay, D. M.; and Hunter, A. 1991. Making inconsistency respectable: a logical framework for inconsistency in reasoning. In *International Workshop on Fundamentals of Artificial Intelligence Research (FAIR'91)*, 19–32.
- Garapa, M.; Fermé, E.; and Reis, M. D. L. 2020. Credibility-limited base revision: new classes and their characterizations. *Journal of Artificial Intelligence Research*, 69: 1023–1075.
- Gärdenfors, P. 1988. *Knowledge in flux*. MIT Press.
- Gärdenfors, P. 1990. Belief revision and nonmonotonic logic: two sides of the same coin? In *Proceedings of the 9th European Conference on Artificial Intelligence (ECAI'90)*, 768–773.
- Girard, P.; and Tanaka, K. 2016. Paraconsistent dynamics. *Synthese*, 193(1): 1–14.
- Hansson, S.; Fermé, E.; Cantwell, J.; and Falappa, M. 2001. Credibility limited revision. *Journal of Symbolic Logic*, 66(4): 1581–1596.
- Hansson, S. O. 1998. What's new isn't always best. *Theoria*, 1–13. Special issue on non-prioritized belief revision.
- Hansson, S. O. 1999. *A textbook of belief dynamics. Theory change and database updating*. Kluwer.
- Katsuno, H.; and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52: 263–294.
- Konieczny, S.; Medina Grespan, M.; and Pino Pérez, R. 2010. 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)*, 161–170.
- Konieczny, S.; and Pino Pérez, R. 2002. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5): 773–808.
- Konieczny, S.; and Pino Pérez, R. 2008. Improvement operators. In *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, 177–187.
- Konieczny, S.; and Pino Pérez, R. 2011. Logic based merging. *Journal of Philosophical Logic*, 40(2): 239–270.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44: 167–207.
- Lehmann, D.; and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1): 1–60.
- Makinson, D. 1998. Screened revision. *Theoria*, 14–23. Special issue on non-prioritized belief revision.
- Mares, E. D. 2002. A paraconsistent theory of belief revision. *Erkenntnis*, 56(2): 229–246.
- Nayak, A. C. 1994. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41(3): 353–390.
- Priest, G. 1979. The logic of paradox. *Journal of Philosophical Logic*, 8(1): 219–241.
- Priest, G. 1991. Minimally inconsistent LP. *Studia Logica*, 50: 321–331.
- Priest, G. 2001. Paraconsistent belief revision. *Theoria*, 67(3): 214–228.
- Priest, G. 2002. *Paraconsistent logic*, 287–393. Springer Netherlands.

- Revesz, P. Z. 1993. On the semantics of theory change: arbitration between old and new information. In *Proceedings of the 12th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Databases*, 71–92.
- Schwind, N.; and Konieczny, S. 2020. 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)*, 738–747.
- Schwind, N.; Konieczny, S.; and Marquis, P. 2018. On belief promotion. In *Proceedings of the 16th International Conference on the Principles of Knowledge Representation and Reasoning (KR'18)*, 297–307.
- Segerberg, K. 1998. Irrevocable belief revision in dynamic doxastic logic. *Notre Dame Journal of Formal Logic*, 39(3): 287–306.
- Testa, R.; Fermé, E.; Garapa, M.; and Reis, M. 2018. How to construct remainder sets for paraconsistent revisions: preliminary report. In *17th International Workshop on Non-Monotonic Reasoning (NMR'18)*.
- Testa, R. R.; Coniglio, M. E.; and Ribeiro, M. M. 2017. AGM-like paraconsistent belief change. *Logic Journal of the IGPL*, 25(4): 632–672.