



Reasoning under inconsistency: A forgetting-based approach [☆]

Jérôme Lang^a, Pierre Marquis^{b,*}

^a LAMSADE-CNRS / Université Paris-Dauphine, Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16, France

^b CRIL-CNRS / Université d'Artois, rue Jean Souvraz, S.P. 18, 62307 Lens Cedex, France

ARTICLE INFO

Article history:

Received 15 September 2009

Received in revised form 23 April 2010

Accepted 23 April 2010

Available online 29 April 2010

Keywords:

Knowledge representation

Reasoning under inconsistency

Forgetting

ABSTRACT

In this paper, a fairly general framework for reasoning from inconsistent propositional bases is defined. Variable forgetting is used as a basic operation for weakening pieces of information so as to restore consistency. The key notion is that of *recoveries*, which are sets of variables whose forgetting enables restoring consistency. Several criteria for defining preferred recoveries are proposed, depending on whether the focus is laid on the relative relevance of the atoms or the relative entrenchment of the pieces of information (or both). Our framework encompasses several previous approaches as specific cases, including reasoning from preferred consistent subsets, and some forms of information merging. Interestingly, the gain in flexibility and generality offered by our framework does not imply a complexity shift compared to these specific cases.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Reasoning from inconsistent pieces of information, represented as logical formulas, is an important issue in Artificial Intelligence. Thus, there are at least two very different contexts where inconsistent sets of formulas have to be dealt with. The first one is when the formulas express *beliefs* about the real world, that may stem from different sources. In this case, inconsistency means that some of the formulas are just wrong. The second one is when the input formulas express *preferences* (or goals, desires) expressed by different agents (or by a single agent according to different criteria). In this case, inconsistency does not mean that anything is incorrect, but that some preferences will not be able to be fulfilled. Even if the nature of the problems is very different whether we are in one context or the other one, many notions and techniques that can be employed to reason from inconsistent sets of formulas are similar.

Whatever the nature of the information represented, classical reasoning is inadequate to derive significant consequences from inconsistent formulas, since it trivializes in this situation (*ex falso quodlibet sequitur*). This calls for other inference relations which avoid the trivialization problem (namely, paraconsistent inference relations) but there is no general consensus about what such relations should be. Actually, both the complexity of the problem of reasoning under inconsistency and its significance are reflected by the number of approaches that have been developed for decades and can be found in the literature under various names, like paraconsistent logics, belief revision, argumentative inference, information merging, model fitting, arbitration, knowledge integration, knowledge purification, etc. (see [7,5] for surveys).

Corresponding to these approaches, many different mechanisms to avoid trivialization can be exploited. A first taxonomy allows for distinguishing between *active approaches*, where inconsistency is removed by identifying wrong pieces of belief through knowledge-gathering actions (see e.g. [34,36,28]) or by the group of agents agreeing on some goals to be given

[☆] This is an extended and revised version of a paper that appeared in the Proceedings of the 8th International Conference on Knowledge Representation and Reasoning (KR'02), 2002, pp. 239–250.

* Corresponding author.

E-mail addresses: lang@lamsade.dauphine.fr (J. Lang), marquis@cril.univ-artois.fr (P. Marquis).

up after a negotiation process, from *passive approaches*, where inconsistency is dealt with. In this latter case, trivialization is avoided by weakening the set of consequences that can be derived from the given base (a set of formulas). In the propositional case, this can be achieved by two means:

- (1) By *weakening the consequence relation of classical logic while keeping the base intact*. Such an approximation by below of classical entailment can be achieved, which typically leads to paraconsistent logics.
- (2) By *weakening the input base while keeping classical entailment*. The pieces of information from the initial base are weakened such that their conjunction becomes consistent. This technique is at work in so-called *coherence-based approaches* to paraconsistent inference (see e.g. [46,24,25,11,3,41,44,4] for some of the early references), where weakening the input base consists in *inhibiting* some of the pieces of information it contains (by removing them). It is also at work in belief merging (see e.g. [37,47,31,40] for some of the early references). Belief merging, and especially distance-based merging consists in weakening the pieces of information by *dilating* them: the piece of information ϕ , instead of expressing that the real world is for sure among the models of ϕ , now expresses that it is close to be a model of ϕ (the further a world ω from the models of ϕ , the less plausible it is that ω is the real world).

The above dichotomy between (1) and (2) is reminiscent of the dichotomy between *actual* and *potential* contradictions, as discussed in [7,5]. *Actual* contradictions tolerate inconsistency by reasoning with a set of inconsistent statements, whereas *potential* contradictions are prevented from arising by putting individually consistent yet jointly inconsistent information together.

In the rest of this paper we deal with potential inconsistencies and focus on the class of approaches, consisting in weakening the input base. While existing weakening-based approaches work well on some families of problems, there are typical examples that they fail to handle in a satisfactory way (see Section 6 for a detailed discussion), the reason being that while some of these approaches take account for the relative importance of pieces of information (or of the corresponding sources), they do not handle the relative importance of *atoms* in the problem at hand. This is problematic in many situations where some atoms are less central than others, especially when some atoms are meaningful only in the presence of others. For instance, it makes little sense to reason about whether John's car is grey if there is some strong conflict about whether John actually *has* a car. Or, in a preference merging context, suppose that a group of co-owners of a residence tries to agree about whether a tennis court or a swimming pool should be built: if there is no agreement about whether the swimming pool is to be constructed, any preference concerning its colour must be (at least temporarily) ignored (this would not be the case for its size, however, because it influences its price in a dramatic way).

More generally, it is sometimes the case that ignoring a small set of propositional atoms of the formulas from an inconsistent set renders it consistent. When reasoning from inconsistent *beliefs*, this allows for giving some useful information about the *other* atoms (those that are not forgotten); information about other atoms can be processed further (for instance through knowledge-gathering actions) if these atoms are important enough. When trying to reach a common decision from an inconsistent set of *preferences*, ignoring small sets of atoms allows for making a decision about all other atoms; the decision about the few remaining atoms can then take place after a negotiation process among agents.

In the following, we define a framework for reasoning from inconsistent propositional bases, using *forgetting* [10,39,33] as a basic operation for weakening formulas. Belief (or preference) bases are viewed as (finite) vectors of propositional formulas, conjunctively interpreted. Without loss of generality, each formula is assumed to be issued from a specific source of information (or a specific agent). Forgetting a set X of atoms in a formula consists in replacing it by its logically strongest consequence which is independent of X , in the sense that it is equivalent to a formula in which no atom from X occurs [33]. The key notion of our approach is that of *recoveries*, which are sets of atoms whose forgetting enables restoring consistency. The intuition of this simple principle is that if a collection of pieces of information is jointly inconsistent, weakening it by ignoring some atoms (for instance the least important ones) can help restoring consistency and derive reasonable conclusions. Several criteria for defining preferred recoveries are proposed, depending on whether the focus is laid on the relative relevance of the atoms or the relative entrenchment of the pieces of information (or both).

Our contributions are composed of the following models and results. We first define a general model for using variable forgetting in order to reason under inconsistency. We show that our model is general enough to encompass several classes of paraconsistent inference relations, including reasoning from preferred consistent subbases (Propositions 4.1 and 4.2) and various types of belief merging (Propositions 4.3, 4.4 and 4.5). Our framework does not only recover known approaches as specific cases (which would make its interest rather limited) but it allows for new families of paraconsistent inferences, including homogeneous inferences, where propositional variables have to be forgotten in a homogeneous way from the different sources, and abstraction-based inferences, where the most specific variables are forgotten. Our contribution with respect to the computational study of these forgetting-based inferences starts with a result stating that these inference relations are not more difficult to compute than some of the well-known specific cases (Proposition 5.1) and also include complexity results about homogeneous inferences (Proposition 5.2) and tractable fragments (Proposition 5.3).

Our framework offers several advantages compared with many existing approaches to paraconsistent reasoning. The main ones are three fold. First, the inference relations defined in our framework are typically less cautious than inference relations stemming from other approaches to paraconsistent reasoning (especially those based on the selection of consistent subsets, but also some forms of merging) because it is based on a weakening mechanism (namely, forgetting) that is more fine-grained than the weakening mechanisms at work in such approaches (inhibition and dilation of formulas). Accordingly,

more information can be derived from inconsistent bases. Second, it is quite general and flexible. As to generality, our framework encompasses several previous approaches as specific cases, including reasoning from preferred consistent subsets, and forms of merging (mainly because inhibition and dilation of formulas can be achieved via forgetting). As to flexibility, our framework enables us to model situations where some sources of information are considered more reliable than others in an absolute way, but also relatively to some topics of interest. Some form of equity between some sources of information can also be achieved by imposing to forget the same atoms in the formulas associated to the sources. Third, the gain in generality and flexibility offered by our framework does not imply a complexity shift compared with these specific cases.

The rest of this paper is organized as follows. Formal preliminaries including a presentation of variable forgetting are given in Section 2. Our general framework is presented in Section 3. In Section 4 we give several particular cases of our general approach, including well-known frameworks (namely, reasoning from consistent subsets, and some forms of merging); we also give specific cases that are not covered by any of the existing classes of approaches. In Section 5 we briefly discuss the computational aspects of our framework. Especially, we show that although it is quite general, our framework does not imply a complexity shift compared to the well-known specific classes of approaches that are recovered. In Section 6 we discuss in detail the benefits of our general approach compared to these specific cases. We mention some related work in Section 7 and we draw some conclusions in Section 8. Full proofs are given in Appendix A.

2. Formal preliminaries

2.1. Propositional logic

$PROP_{PS}$ denotes the propositional language built up from a finite set PS of symbols (also referred to as atoms or variables), the Boolean constants \top (true) and \perp (false), and the connectives $\neg, \vee, \wedge, \Rightarrow, \Leftrightarrow, \oplus$. $Var(\phi)$ denote the set of propositional variables occurring in the formula ϕ . $\phi_{x\leftarrow 0}$ (resp. $\phi_{x\leftarrow 1}$) denotes the formula obtained by replacing in ϕ every occurrence of symbol x by \perp (resp. \top).

An interpretation (or world) ω over PS is an assignment of a truth value to each variable of PS . $\Omega = 2^{PS}$ is the set of all interpretations. Formulas are interpreted in the classical way. Every finite set of formulas is interpreted conjunctively. $Mod(\phi)$ denotes the set of models of ϕ . Rather than denoting interpretations by sets of variables, we choose to denote them by words over the set of $\{x, \bar{x} \mid x \in PS\}$. For instance, if $PS = \{a, b, c, d\}$, $\omega = a\bar{b}\bar{c}d$ is the interpretation that assigns a and d to true and b and c to false. \models denotes logical entailment and \equiv denotes logical equivalence. Finally, ω and ω' being two interpretations, $Diff(\omega, \omega')$ is the set of propositional variables assigned different truth values by ω and ω' . For instance, $Diff(a\bar{b}\bar{c}d, a\bar{b}c\bar{d}) = \{b, d\}$.

If $X \subseteq PS$, an X -interpretation ω_X is a truth assignment to the variables in X , that is, an element of 2^X . While it is not the case that every formula of $PROP_{PS}$ can be given a truth value in an X -assignment, the notion of satisfaction of a formula ϕ from $PROP_X$ by an X -interpretation ω_X is well defined (and coincides with the standard notion of satisfaction). The *projection* of an interpretation ω on a subset of variables X of PS , denoted by $\omega^{\downarrow X}$, is the X -interpretation which is the restriction of ω to X .

2.2. Forgetting

Our approach to restore consistency is based on (*variable*) *forgetting*, also known as projection. Forgetting can be defined as follows (see [39,33] for more details):

Definition 2.1 (*Forgetting*). Let ϕ be a formula from $PROP_{PS}$ and $V \subseteq PS$. The *forgetting* of V in ϕ , noted $\exists V.\phi$, is a formula from $PROP_{PS}$ that is inductively defined up to logical equivalence as follows:

- $\exists \emptyset.\phi \equiv \phi$;
- $\exists \{x\}.\phi \equiv \phi_{x\leftarrow 0} \vee \phi_{x\leftarrow 1}$;
- $\exists (\{x\} \cup V).\phi \equiv \exists V.(\exists \{x\}.\phi)$.

For example, $\exists \{a\}.\neg a \wedge b \equiv b$ and $\exists \{a\}.(a \vee b) \equiv \top$.

When V is a singleton $\{x\}$, we typically write $\exists x.\phi$ instead of $\exists \{x\}.\phi$.

As the notation used suggests it, the forgetting of $V = \{v_1, \dots, v_k\}$ in ϕ is equivalent to the quantified Boolean formula $\exists v_1.(\dots(\exists v_k.\phi)\dots)$.

$\exists V.\phi$ represents the logically strongest consequence ψ (unique up to logical equivalence) of ϕ that is independent of V , where “ ψ is independent of X ” means that there exists a formula ψ' from $PROP_{PS}$ s.t. $\psi \equiv \psi'$ and $Var(\psi') \cap V = \emptyset$. Accordingly, forgetting a set of variables within a formula leads to *weaken* it. To be more precise, if $V \subseteq W$ holds, then $\exists V.\phi \equiv \exists W.\phi$ holds. Moreover, ϕ is consistent if and only if $\exists Var(\phi).\phi$ is valid (see [33]).

Many characterizations of forgetting, together with complexity results, are reported in [33]. Noticeably, for every $V \subseteq PS$ and every formula ϕ from $PROP_{PS}$, we have $\exists V.\phi \equiv \exists V_\phi.\phi$, where $V_\phi = V \cap Var(\phi)$ – which means that forgetting variables that do not appear in a formula does not have any effect.

The following characterization of forgetting [33] will be used further in the paper: for any $\phi \in \text{PROP}_{PS}$, $V \subseteq PS$ and $\omega_{\bar{V}} \in 2^{\bar{V}}$ where $\bar{V} = PS \setminus V$, we have $\omega_{\bar{V}} \models \exists V.\phi$ if and only if there exists an interpretation $\omega \in \Omega$ such that $\omega \models \phi$ and $\omega \downarrow^{\bar{V}} = \omega_{\bar{V}}$.

3. Reasoning from preferred recoveries

3.1. Bases and recoveries

We start by defining propositional bases.

Definition 3.1 (Base). A base B is a vector $\langle \phi_1, \dots, \phi_n \rangle$ of n formulas from PROP_{PS} , where n is a positive integer.

Each i ($1 \leq i \leq n$) identifies a source of information and ϕ_i denotes the piece of information conveyed by source i . Note that it can be the case that a formula occurs more than once in B , which can be used to model the situation where two different sources (or more) give the same information.

$B = \langle \phi_1, \dots, \phi_n \rangle$ is conjunctively interpreted, so that it is said to be *inconsistent* if and only if $\bigwedge_{i=1}^n \phi_i$ is inconsistent; in the remaining case, it is said to be *consistent*. $B = \langle \phi_1, \dots, \phi_n \rangle$ and $B' = \langle \phi'_1, \dots, \phi'_n \rangle$ are said to be *equivalent*, noted $B \equiv B'$, if and only if for each $i \in \{1, \dots, n\}$, we have $\phi_i \equiv \phi'_i$. Finally, we note $\text{Var}(\langle \phi_1, \dots, \phi_n \rangle) = \bigcup_{i=1}^n \text{Var}(\phi_i)$.

The key notions of our approach are *forgetting vectors* and *recoveries*. A forgetting vector consists of sets of variables to be forgotten in each formula from the base. These sets of variables need not be identical, but they should obey some constraints bearing on the forgetting process:

Definition 3.2 (Forgetting context). Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base. A *forgetting context* for B is a consistent propositional formula \mathcal{C} over $F_{PS,n} = \{\text{forget}(x, i) \mid x \in PS, i \in \{1, \dots, n\}\}$ (a set of propositional atoms). $\text{forget}(x, i)$ means that atom x may be forgotten in ϕ_i .

For instance, $\neg\text{forget}(x, i)$ means that forgetting variable x in ϕ_i is forbidden. This is helpful to model the situation where source i is fully reliable to what concerns x , i.e., what i says about x must be taken for sure. $\text{forget}(x, i) \Rightarrow \text{forget}(y, i)$ enables to express that for some sources (i), some atoms (y) can be significant only in presence of others (x) so that forgetting the latter imposes to forget the former. $\text{forget}(x, i) \Rightarrow \text{forget}(x, j)$ can be used to force some sources of information (here i and j) to be considered on equal terms w.r.t. weakening. For instance, if ϕ_i and ϕ_j are together inconsistent and consistency can be recovered by forgetting x in ϕ_i , then it could be expected that x should also be forgotten in ϕ_j .

While the notion of forgetting context has already been considered in [35,6], this way of representing them allows for going beyond previous representations based on binary relations; the generalization is both in terms of expressiveness and spatial efficiency. Thus, a binary relation \mathcal{R} over $PS \times \{1, \dots, n\}$ (see [6]), where $(x, i)\mathcal{R}(y, j)$ means that if x is forgotten in ϕ_i , then y must be forgotten in ϕ_j can be easily (i.e., in linear time) encoded as a formula $\bigwedge_{x,y \in PS} \bigwedge_{i,j \in \{1, \dots, n\}} (\text{forget}(x, i) \Rightarrow \text{forget}(y, j))$. Conversely, the forgetting context $\neg\text{forget}(x, i) \vee \neg\text{forget}(y, i)$ (stating that in source i at most one of the two atoms x and y may be forgotten) cannot be represented by such a binary relation \mathcal{R} . Representing forgetting contexts by propositional formulas also allows for taking advantage of existing approaches for managing propositional representations. Consider for instance the following update problem: if B becomes independent from an atom x (through an update) (resp. if a source i becomes unavailable), the forgetting constraint corresponding to the updated base can be characterized by the forgetting of $\{\text{forget}(x, i) \mid i \in \{1, \dots, n\}\}$ (resp. $\{\text{forget}(x, i) \mid x \in \text{Var}(B)\}$) in the forgetting context \mathcal{C} for B .

Whereas we do not put any restriction (other than consistency) on forgetting contexts, clearly some contexts are intuitively more satisfactory than others, and in many practical situations, representing contexts by definite, binary clauses as in [6] or as specific classes of such clauses as in [35], can prove to be sufficient. However, in the rest of the paper, such assumptions are unnecessary, which is why we prefer to allow every possible context. The only two (very weak) conditions on contexts that are often desirable and that we will need later in the paper are:

- *Downward normality*: $\{\neg\text{forget}(x, i) \mid x \in \text{Var}(B), i \in \{1, \dots, n\}\} \cup \{\mathcal{C}\}$ is consistent;
- *Upward normality*: $\{\text{forget}(x, i) \mid x \in \text{Var}(B), i \in \{1, \dots, n\}\} \cup \{\mathcal{C}\}$ is consistent.

Downward (resp. upward) normality means that the forgetting context does not exclude the possibility not to forget anything (resp. the possibility to forget everything).

Example 3.1. (Inspired from [31].) As a matter of illustration, let us consider the following preference merging scenario. Suppose that a group of four co-owners of a residence tries to agree about building a new tennis court t and/or a new swimming pool s . If it is constructed, the swimming pool can be either red (s_r) or blue (s_b). If both a tennis court and a pool – respectively, one of them, none of them – are (is) constructed, the induced cost is 2 money units (c_2) – respectively, 1 unit (c_1), nothing (c_0). The first co-owner would not like to spend more than 1 unit, and prefers a red swimming pool,

should it be constructed. The second co-owner would like to improve the quality of the residence by the construction of at least a tennis court or a swimming pool and would prefer a blue swimming pool, should it be constructed. The third co-owner just prefers a swimming pool to be built, whatever its colour. The fourth co-owner would like both a swimming pool (of whatever colour) and a tennis court to be built.

Of course, if there is no agreement about whether the swimming pool is to be constructed, any preference concerning its colour must be ignored. Furthermore, it is meaningless to forget about whether the pool should be blue without forgetting whether it should be red: either we forget about the colour of the swimming pool or not. Similarly, either we forget about the expenses, i.e., we forget $\{c_0, c_1, c_2\}$ or not (i.e., we forget none of them).

Clearly enough, the preferences of the group are jointly inconsistent. This scenario can be encoded in our framework using the following base $B = \langle \phi_1, \phi_2, \phi_3, \phi_4, \phi_5 \rangle$ and forgetting context \mathcal{C} :

- $\phi_1 = (s \Rightarrow (s_r \oplus s_b)) \wedge (\neg s \Rightarrow \neg s_r \wedge \neg s_b) \wedge (c_2 \Leftrightarrow (s \wedge t)) \wedge (c_1 \Leftrightarrow (s \oplus t)) \wedge (c_0 \Leftrightarrow (\neg s \wedge \neg t))$;
- $\phi_2 = (c_0 \vee c_1) \wedge (s \Rightarrow s_r)$;
- $\phi_3 = (s \vee t) \wedge (s \Rightarrow s_b)$;
- $\phi_4 = s$;
- $\phi_5 = s \wedge t$;
- \mathcal{C} is the conjunction of the following formulas:
 - $\bigwedge_{x \in \text{Var}(\phi_1)} \neg \text{forget}(x, 1)$;
 - $\bigwedge_{i=2}^5 (\text{forget}(s, i) \Rightarrow \text{forget}(s_r, i)) \wedge (\text{forget}(s_r, i) \Leftrightarrow \text{forget}(s_b, i)) \wedge (\text{forget}(c_0, i) \Leftrightarrow \text{forget}(c_1, i)) \wedge (\text{forget}(c_1, i) \Leftrightarrow \text{forget}(c_2, i))$.

Given that $\bigwedge_{x \in \text{Var}(\phi_1)} \neg \text{forget}(x, 1)$ is a logical consequence of \mathcal{C} , ϕ_1 expresses an *integrity constraint* which cannot be weakened under any circumstance (this integrity constraint expresses that the swimming pool must be either red or blue if it is constructed, as well as the logical definition of the induced price from the number of equipments built). ϕ_2, ϕ_3, ϕ_4 and ϕ_5 encode the preferences of the co-owners. The second conjunct of \mathcal{C} ensures that s_r and s_b are irrelevant if there is no agreement on s , that s_r is irrelevant if and only if s_b is, and that for any $j, k \in \{0, 1, 2\}$, c_j is irrelevant if and only if c_k is. In the situation where all co-owners must be considered on equal terms with respect to the set of variables to be forgotten, a third conjunct must be added to \mathcal{C} : $\bigwedge_{x \in \text{Var}(B)} \bigwedge_{i,j=2}^5 \text{forget}(x, i) \Rightarrow \text{forget}(x, j)$; indeed, if we forget any atom x in ϕ_i for some $i \geq 2$ then x has to be forgotten in every ϕ_j for $j \geq 2$.

We can now define the central notions of this paper: those of *forgetting vector*, *projection* and *recovery*.

Definition 3.3 (*Forgetting vector*). Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base and \mathcal{C} a forgetting context for B . For any vector $\vec{V} = \langle V_1, \dots, V_n \rangle$ of subsets of PS , $F_{\vec{V}}$ denotes the interpretation over $F_{PS,n}$ which assigns $\text{forget}(x, i)$ to true if and only if $x \in V_i$, for each $x \in PS$ and each $i \in \{1, \dots, n\}$. \vec{V} is said to be a *forgetting vector* for B given \mathcal{C} if and only if $F_{\vec{V}}$ is a model of \mathcal{C} . $\mathcal{F}_{\mathcal{C}}(B)$ denotes the set of all forgetting vectors for B given \mathcal{C} .

For any forgetting vector $\vec{V} = \langle V_1, \dots, V_n \rangle$ for B given \mathcal{C} , we define $\bigcup \vec{V} = V_1 \cup \dots \cup V_n$.

Observe that the forgetting vectors for a base B do not depend directly on the formulas ϕ_i in it, but on the number n of such formulas. Contrastingly, $\mathcal{F}_{\mathcal{C}}(B)$ heavily depends on \mathcal{C} .

Definition 3.4 (*Projection*). Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base, \mathcal{C} a forgetting context for B , and \vec{V} a forgetting vector for B given \mathcal{C} . The *projection* of \vec{V} on B is the set of formulas $B \mid \vec{V} = \{\exists V_i. \phi_i \mid i \in \{1, \dots, n\}\}$.

Definition 3.5 (*Recovery*). Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base, and \mathcal{C} a forgetting context for B . A vector $\vec{V} = \langle V_1, \dots, V_n \rangle$ of subsets of PS is a *recovery* for B given \mathcal{C} if and only if the following two conditions hold:

- (1) \vec{V} is a forgetting vector for B given \mathcal{C} .
- (2) $B \mid \vec{V}$ is consistent.

$\mathcal{R}_{\mathcal{C}}(B) \subseteq \mathcal{F}_{\mathcal{C}}(B)$ denotes the set of all recoveries for B given \mathcal{C} . B is said to be *recoverable* with respect to \mathcal{C} if and only if $\mathcal{R}_{\mathcal{C}}(B) \neq \emptyset$.

Example 3.2. Let B and \mathcal{C} be as in Example 3.1. All the vectors \vec{V}^1 to \vec{V}^9 considered in Table 1 are recoveries for B given \mathcal{C} , while the following ones are *no* recoveries for B given \mathcal{C} :

- $\langle \emptyset, \{t\}, \{t\}, \{t\}, \{t\} \rangle$, because $\exists V_1. \phi_1 \wedge \exists V_2. \phi_2 \wedge \exists V_3. \phi_3 \wedge \exists V_4. \phi_4 \wedge \exists V_5. \phi_5$ is inconsistent;
- $\langle \emptyset, \{s\}, \{s\}, \{s\}, \{s\} \rangle$, because it does not satisfy the constraints of \mathcal{C} : if we forget s we also have to forget s_b and s_r .

Table 1

Recovery	$\exists V_1.\phi_1$	$\exists V_2.\phi_2$	$\exists V_3.\phi_3$	$\exists V_4.\phi_4$	$\exists V_5.\phi_5$	$B \mid \vec{V}$
$\vec{V}^1 = (\emptyset, \{s, s_b, s_r\}, \{s, s_b, s_r\}, \{s, s_b, s_r\}, \{s, s_b, s_r\})$	ϕ_1	$c_0 \vee c_1$	\top	\top	t	$\phi_1 \wedge \neg s \wedge t \wedge c_1$
$\vec{V}^2 = (\emptyset, \{t, s_b, s_r\}, \{t, s_b, s_r\}, \{t, s_b, s_r\}, \{t, s_b, s_r\})$	ϕ_1	$c_0 \vee c_1$	\top	s	s	$\phi_1 \wedge s \wedge \neg t \wedge c_1$
$\vec{V}^3 = (\emptyset, \{c_0, c_1, c_2, s_b, s_r\}, \{c_0, c_1, c_2, s_b, s_r\}, \{c_0, c_1, c_2, s_b, s_r\}, \{c_0, c_1, c_2, s_b, s_r\})$	ϕ_1	\top	$s \vee t$	s	$s \wedge t$	$\phi_1 \wedge s \wedge t \wedge c_2$
$\vec{V}^4 = (\emptyset, \{s_b, s_r\}, \{s_b, s_r\}, \emptyset, \{t\})$	ϕ_1	$c_0 \vee c_1$	$s \vee t$	s	s	$\phi_1 \wedge s \wedge \neg t \wedge c_1$
$\vec{V}^5 = (\emptyset, \{c_0, c_1, c_2\}, \{s_b, s_r\}, \emptyset, \emptyset)$	ϕ_1	$s \Rightarrow s_r$	$s \vee t$	s	$s \wedge t$	$\phi_1 \wedge s \wedge t \wedge s_r \wedge c_2$
$\vec{V}^6 = (\emptyset, \{s_b, s_r\}, \emptyset, \emptyset, \{t\})$	ϕ_1	$c_0 \vee c_1$	$(s \vee t) \wedge (s \Rightarrow s_b)$	s	s	$\phi_1 \wedge s \wedge \neg t \wedge s_b \wedge c_1$
$\vec{V}^7 = (\emptyset, \emptyset, \emptyset, \{s, s_b, s_r\}, \{s, s_b, s_r\})$	ϕ_1	$(c_0 \vee c_1) \wedge (s \Rightarrow s_r)$	$(s \vee t) \wedge (s \Rightarrow s_b)$	\top	t	$\phi_1 \wedge \neg s \wedge t \wedge c_1$
$\vec{V}^8 = (\emptyset, \{c_0, c_1, c_2, s_b, s_r\}, \emptyset, \emptyset, \emptyset)$	ϕ_1	\top	$(s \vee t) \wedge (s \Rightarrow s_b)$	s	$s \wedge t$	$\phi_1 \wedge s \wedge t \wedge s_b \wedge c_2$
$\vec{V}^9 = (\emptyset, \emptyset, \{s_b, s_r\}, \emptyset, \{t\})$	ϕ_1	$(c_0 \vee c_1) \wedge (s \Rightarrow s_r)$	$s \vee t$	s	s	$\phi_1 \wedge s \wedge \neg t \wedge s_r \wedge c_1$

The other columns of the table (except the rightmost one) give formulas equivalent to the “local” projections $\exists V_i.\phi_i$, and the last column gives formulas equivalent to the projection of the corresponding recovery on B .

By construction, replacing the pieces of information ϕ_i of B by the projection on B of any of its recoveries is sufficient to restore consistency, provided that B is recoverable.

In the general case, it may happen that $\mathcal{R}_{\mathcal{C}}(B)$ is empty, because some atoms must not be forgotten. For instance, in Example 3.1, if \mathcal{C} is strengthened, so that now c_0 and c_1 are protected in ϕ_2 and s and t are protected in ϕ_5 , then the corresponding $\mathcal{R}_{\mathcal{C}}(B)$ is empty. Now, another important sufficient condition for a base not to be recoverable is when one of the ϕ_i is inconsistent. Indeed, if ϕ_i is inconsistent then even forgetting all atoms from PS will not help us restoring consistency, because $\exists PS.\phi_i \equiv \exists PS.\perp \equiv \perp$. Therefore, for any forgetting vector \vec{V} for B given \mathcal{C} , $(B \mid \vec{V})$ is inconsistent. Conversely, when each formula ϕ_i of B is consistent and the forgetting context \mathcal{C} is equivalent to \top , then B is recoverable (especially, the forgetting vector \vec{V} such that each V_i is equal to $Var(B)$ is a recovery for B given \mathcal{C}).

It is easy to show that for any base B , the logical strength of its forgetting context \mathcal{C} has a direct impact on the set of all forgetting vectors $\mathcal{F}_{\mathcal{C}}(B)$ for B given \mathcal{C} , hence on the set of all recoveries $\mathcal{R}_{\mathcal{C}}(B)$ for B given \mathcal{C} ; to be more precise, we have for any B :

$$\text{If } \mathcal{C} \models \mathcal{C}' \text{ then } \mathcal{F}_{\mathcal{C}}(B) \subseteq \mathcal{F}_{\mathcal{C}'}(B) \text{ and } \mathcal{R}_{\mathcal{C}}(B) \subseteq \mathcal{R}_{\mathcal{C}'}(B).$$

The number of forgetting vectors $\mathcal{F}_{\mathcal{C}}(B)$ for B given \mathcal{C} may vary between 1 (remind that \mathcal{C} is a consistent formula) and $2^{n|PS|}$, while the number of recoveries $\mathcal{R}_{\mathcal{C}}(B)$ for B given \mathcal{C} may vary between 0 and $2^{n|PS|}$.

3.2. Forgetting contexts

As expected, our definition of forgetting contexts is not very restrictive so that many different forgetting contexts can be considered. A first distinguished forgetting context, viewed as a baseline, is the *standard* forgetting context, denoted $\mathcal{C}_{\mathcal{S}}$, where no constraint bears on the atoms to be forgotten:

Definition 3.6 (Standard context). The *standard forgetting context* $\mathcal{C}_{\mathcal{S}}$ for $B = \langle \phi_1, \dots, \phi_n \rangle$ is defined by $\mathcal{C}_{\mathcal{S}} = \top$.

In this context, every atom can be forgotten in any piece of information, and atoms can always be forgotten in an independent way. Obviously enough, the standard forgetting context for a base is the logically weakest one (up to logical equivalence). Imposing stronger constraints on the forgetting vectors is a simple way to focus on expected recoveries.

Thus, in some situations, the atoms forgotten in each of the pieces of information must be identical so that all sources of information (or, more generally, a distinguished subset G among them) are weakened in the same way. This can be captured by considering *homogeneous* contexts:

Definition 3.7 (Homogeneous context). A forgetting context \mathcal{C} for $B = \langle \phi_1, \dots, \phi_n \rangle$ is *homogeneous* for $G \subseteq \{1, \dots, n\}$ (homogeneous for short, when $G = \{1, \dots, n\}$) if and only if $\bigwedge_{i,j \in G} \bigwedge_{x \in Var(B)} \text{forget}(x, i) \Rightarrow \text{forget}(x, j)$ is a logical consequence of \mathcal{C} .

A simple case when \mathcal{C} should not be homogeneous (unless B is consistent) is when some of the pieces of information ϕ_i must be left intact: for such a ϕ_i , $\bigwedge_{x \in \text{Var}(\phi_i)} \neg \text{forget}(x, i)$ is expected to be a logical consequence of \mathcal{C} . As explained before, this is useful to encode integrity constraints (or revision formulas), i.e., formulas which must be protected since they are required to be true (just like ϕ_1 in Example 3.1).

Another distinguished family of forgetting contexts consists of the *binary* ones, where, as soon as an atom is forgotten from any source of information then all other atoms from a distinguished subset $V \subseteq \text{Var}(B)$ must be forgotten from ϕ_i .

Definition 3.8 (*Binary context*). A forgetting context \mathcal{C} for $B = \langle \phi_1, \dots, \phi_n \rangle$ is *binary* for $V \subseteq \text{Var}(B)$ (binary for short, when $V = \text{Var}(B)$) if and only if $\bigwedge_{i=1}^n \bigwedge_{x, y \in V} \text{forget}(x, i) \Rightarrow \text{forget}(y, i)$ is a logical consequence of \mathcal{C} .

Binary contexts are useful for scenarios where one considers that a source of information is unreliable as soon as one of the pieces of information it gives is involved in a contradiction (and in such a case, no information from the source are preserved: every “useful” atom from V is forgotten from it).

3.3. Preferred recoveries

Generally, many recoveries for a base given a forgetting context are possible, but most of the time, some of the corresponding recoveries are more expected than others. What are usually available are selection policies (or preference criteria) which characterize such expected recoveries independently from the base, in an implicit way. While forgetting contexts only specify *hard constraints* on possible recoveries, preference criteria specify soft constraints, which have to be optimized, so that optimal recoveries are the preferred ones.

As expected, many preference criteria can be considered so as to capture in formal terms several intuitions about the way the pieces of information should be merged. For instance, in some situations, we may prefer recoveries $\vec{V} = \langle V_1, \dots, V_n \rangle$ for $B = \langle \phi_1, \dots, \phi_n \rangle$ in which the set V_i of atoms forgotten in some of the ϕ_i 's are as close to each other as possible, and ideally coincide (which always is the case whenever homogeneous contexts are considered). In this extreme case, i.e., when $V_1 = \dots = V_n$, $\vec{V} = \langle V_1, \dots, V_n \rangle$ is said to be *homogeneous*. Thus, if both $\vec{V} = \langle \{a\}, \{b, c\} \rangle$ and $\vec{V}' = \langle \{a\}, \{a, c\} \rangle$ are recoveries for B , \vec{V}' should (in some contexts) be preferred to \vec{V} because it is “more homogeneous” in some sense.

We may also prefer recoveries that lead to forget minimal sets of atoms, where minimality may be defined using a preference criterion induced by some priorities or penalties on atoms. Finally, we may prefer recoveries that preserve as much as preferred pieces of information as possible (like in the coherence-based approaches to paraconsistent reasoning).

Formally, each preference criterion associates to every base B and every forgetting context \mathcal{C} for B a preorder (i.e., a reflexive and transitive relation) \sqsubseteq on $\mathcal{F}_{\mathcal{C}}(B)$, called a preference relation.

Definition 3.9 (*Preference relation on forgetting vectors*). Given a base $B = \langle \phi_1, \dots, \phi_n \rangle$ and a forgetting context \mathcal{C} for it, a *preference relation* is a reflexive and transitive relation \sqsubseteq on $\mathcal{F}_{\mathcal{C}}(B)$.

As usual, we write $\vec{V} \sqsubset \vec{V}'$ for $\vec{V} \sqsubseteq \vec{V}'$ and $\vec{V}' \not\sqsubseteq \vec{V}$, and $\vec{V} \sim \vec{V}'$ for $\vec{V} \sqsubseteq \vec{V}'$ and $\vec{V}' \sqsubseteq \vec{V}$.

The most preferred elements from $\mathcal{R}_{\mathcal{C}}(B)$ are defined as the minimal ones for \sqsubseteq .

One could have thought of selecting the most expected recoveries by strengthening the forgetting context \mathcal{C} under consideration. However, such an approach is generally not feasible, because determining these most expected recoveries cannot be done *a priori*, that is, these recoveries are of course dependent on the base B , and their computation from B is far from easy.

Typically, the preference relation \sqsubseteq will satisfy the following monotonicity, or even strong monotonicity, condition (although we do not need to require it):

- \sqsubseteq satisfies *monotonicity* if and only if $\forall \vec{V}, \vec{V}' \in \mathcal{F}_{\mathcal{C}}(B)$, if $\vec{V} \subseteq_p \vec{V}'$ (i.e., for every i we have $V_i \subseteq V'_i$), then $\vec{V} \sqsubseteq \vec{V}'$.
- \sqsubseteq satisfies *strong monotonicity* if and only if $\forall \vec{V}, \vec{V}' \in \mathcal{F}_{\mathcal{C}}(B)$, if $\vec{V} \subset_p \vec{V}'$ (i.e., for every i we have $V_i \subset V'_i$), then $\vec{V} \sqsubset \vec{V}'$.

Note that \subseteq_p satisfies strong monotonicity (and *a fortiori*, monotonicity). It is easy to prove that for any $\vec{V}, \vec{V}' \in \mathcal{F}_{\mathcal{C}}(B)$, if $\vec{V} \subseteq_p \vec{V}'$ then $B \mid \vec{V} \models B \mid \vec{V}'$ holds, since we forget at least as much in \vec{V}' as in \vec{V} . Accordingly, among the recoveries from $\mathcal{R}_{\mathcal{C}}(B)$, the minimal ones w.r.t. $\sqsubseteq = \subseteq_p$ lead to projections that preserve as much information as possible given \mathcal{C} . This is the rationale for the monotonicity property.

Example 3.3. Let B and \mathcal{C} be as in Example 3.1. We have $\vec{V}^6 \subseteq_p \vec{V}^4$. Accordingly, $B \mid \vec{V}^6 \models B \mid \vec{V}^4$.

Depending on the problem at hand, many other properties on \sqsubseteq can be imposed so as to capture various intuitions about the result of the merging process. We give two of them, which echo the properties imposed on forgetting contexts and will play an important role further.

Definition 3.10 (*Homogeneity*). Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base and \mathcal{C} a forgetting context for B . A preference relation \sqsubseteq is said to satisfy the *homogeneity* property if and only if for any $\vec{V} \in \mathcal{F}_{\mathcal{C}}(B)$, if $(\bigcup \vec{V}, \dots, \bigcup \vec{V})$ is in $\mathcal{F}_{\mathcal{C}}(B)$ then $\vec{V} \sim (\bigcup \vec{V}, \dots, \bigcup \vec{V})$.

The intuitive meaning of homogeneity is that as soon as an atom is forgotten in some of the ϕ_i 's then it is exactly as bad as if it were forgotten everywhere.

Definition 3.11 (*Binaricity*). Given a base $B = \langle \phi_1, \dots, \phi_n \rangle$, and a forgetting context \mathcal{C} for B , a preference relation \sqsubseteq is said to satisfy the *binaricity* property if and only if, for any $\vec{V}, \vec{V}' \in \mathcal{F}_{\mathcal{C}}(B)$ such that for every $i \in \{1, \dots, n\}$ we have $V_i \neq \emptyset$ if and only if $V'_i \neq \emptyset$, then $\vec{V} \sim \vec{V}'$.

The intuitive meaning of binaricity is that no matter how many atoms we forget in ϕ_i , as soon as we forget at least one of them it is exactly as bad as if we were to forget all of them (or more precisely, all that we are allowed to forget). Observe that binaricity is equivalent to: $\vec{V} \sim \vec{V}^*$, where $V^* = \langle V_1^*, \dots, V_n^* \rangle$ is defined by for every $i \in \{1, \dots, n\}$, $V_i^* = \emptyset$ if $V_i = \emptyset$ and $V_i^* = \text{Var}(\phi_i)$ otherwise.

In many cases, it is desirable to assume that the preference relation is a *complete* preorder. In this situation, it can be equivalently represented by a *ranking function* μ from $\mathcal{F}_{\mathcal{C}}(B)$ to \mathbb{N} .

The preference relation \sqsubseteq_{μ} induced by μ is the complete preorder defined by $\forall \vec{V} \in \mathcal{F}_{\mathcal{C}}(B)$, $\vec{V} \sqsubseteq_{\mu} \vec{V}'$ if and only if $\mu(\vec{V}) \leq \mu(\vec{V}')$. We say that a ranking function μ is *decomposable* if and only if there exists a total function $\mu_0 : 2^{PS} \rightarrow \mathbb{N}$ and a total function $H : \mathbb{N}^n \rightarrow \mathbb{N}$ such that $\mu(\vec{V}) = H(\mu_0(V_1), \dots, \mu_0(V_n))$.

Intuitively, $\mu_0(X)$ is the penalty for forgetting X in some ϕ_i , and $H(p_1, \dots, p_n)$ is the aggregation of penalties p_1, \dots, p_n .

3.4. Forget-based inferences

We are now in position to introduce the family of inference relations $\models_{\sqsubseteq}^{\mathcal{C}}$ that can be defined in our framework by letting \mathcal{C} and \sqsubseteq to vary. We call them *forget-based* inference relations.

Definition 3.12 (*Skeptical forget-based inference*). Given a base B and a forgetting context \mathcal{C} for it, let \sqsubseteq be a preference relation on $\mathcal{F}_{\mathcal{C}}(B)$ (possibly induced by a ranking function μ).

- A recovery $\vec{V} \in \mathcal{R}_{\mathcal{C}}(B)$ is said to be *preferred* (w.r.t. \sqsubseteq) if and only if it is minimal in $\mathcal{R}_{\mathcal{C}}(B)$ with respect to \sqsubseteq , i.e., there is no $\vec{V}' \in \mathcal{R}_{\mathcal{C}}(B)$ such that $\vec{V}' \sqsubset \vec{V}$. The set of all preferred recoveries in $\mathcal{R}_{\mathcal{C}}(B)$ (w.r.t. \sqsubseteq) is denoted by $\text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \sqsubseteq)$.
- Let ψ be any formula from PROP_{PS} . ψ is said to be (skeptically) inferred from B w.r.t. \sqsubseteq , denoted by $B \models_{\sqsubseteq}^{\mathcal{C}} \psi$, if and only if for any preferred recovery $\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \sqsubseteq)$, we have $B \mid \vec{V} \models \psi$.

As far as skeptical inference is concerned, among the preferred recoveries for B , only the maximal ones with respect to \sqsubseteq_p are relevant for inference, in the sense that the other ones can be ignored without modifying the inference relation $\models_{\sqsubseteq}^{\mathcal{C}}$. Indeed, such maximal elements correspond to the logically weakest projections.

Clearly, there is no reason for considering only skeptical inferences and not defining also *adventurous* forget-based inferences, where ψ is inferred if and only if $B \mid \vec{V} \models \psi$ holds for *some* preferred recovery $\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \sqsubseteq)$. We can also consider *prudent* forget-based inferences, obtained by first determining the pointwise union of all preferred recoveries for B w.r.t. \sqsubseteq , that is,

$$\vec{V}_{\text{prudent}} = \left\langle \bigcup_{\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \sqsubseteq)} V_1, \dots, \bigcup_{\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \sqsubseteq)} V_n \right\rangle$$

and then inferring ψ if and only if (a) \vec{V}_{prudent} satisfies the constraints imposed by \mathcal{C} and (b) $B \mid \vec{V}_{\text{prudent}} \models \psi$. Intuitively, one first determines all preferred recoveries, and then, in each ϕ_i of B , if \mathcal{C} allows it, we forget all atoms that are mentioned in *at least one* recovery in $\mathcal{R}_{\mathcal{C}}(B)$ at rank i .

However, for the sake of brevity, we will focus mainly on skeptical forget-based inference; in the rest of this section, we present some interesting properties it satisfies.

First, when \sqsubseteq is a complete relation, one can derive a characterization of skeptical forget-based inference in terms of preferred models, à la Shoham [48]. We need the following definition:

Definition 3.13 (\leq_{\sqsubseteq}). Let B be a base, \mathcal{C} a forgetting context for B and \sqsubseteq a complete preference relation on $\mathcal{F}_{\mathcal{C}}(B)$. We first define $\Omega_B = \{\omega \in \Omega \mid \exists \vec{V} \in \mathcal{R}_{\mathcal{C}}(B), \omega \models B \mid \vec{V}\}$.

- For any $\omega \in \Omega_B$, we note

$$\text{Pref}(\mathcal{R}_{\mathcal{C}}(B)(\omega), \sqsubseteq) = \min(\{\vec{V} \in \mathcal{R}_{\mathcal{C}}(B) \mid \omega \models B \mid \vec{V}\}, \sqsubseteq).$$

- For any $\omega, \omega' \in \Omega_B$, we note $\omega \leq_{\sqsubseteq} \omega'$ if and only if \vec{V}_ω and $\vec{V}_{\omega'}$ are in, respectively, $\text{Pref}(\mathcal{R}_C(B)(\omega), \sqsubseteq)$ and $\text{Pref}(\mathcal{R}_C(B)(\omega'), \sqsubseteq)$ (note that because \sqsubseteq is complete, whether $\omega \leq_{\sqsubseteq} \omega'$ holds is independent of the particular choice of \vec{V}_ω and $\vec{V}_{\omega'}$).

Proposition 3.1. *Let B be a base, C a forgetting context for B , \sqsubseteq a complete preference relation on $\mathcal{F}_C(B)$, and ψ a formula. We have $B \approx_{\sqsubseteq}^C \psi$ if and only if $\omega \models \psi$ holds for every $\omega \in \min(\Omega_B, \leq_{\sqsubseteq})$.*

Whether this property can be generalized to the case \sqsubseteq is not complete is a non-trivial question – since it is not central to the paper, we leave it out.

Now, it is easy to check that, whatever the forgetting context, two equivalent bases have the same skeptical forget-based consequence; furthermore, every logical consequence of a skeptical forget-based consequence of a base B also is a skeptical forget-based consequence of B (stated otherwise, skeptical forget-based inference satisfies the properties of “syntax independence” and “right weakening”). Other standard properties like here because a base is not a simple set of formulas, but a vector of formulas (hence an implicit “comma” connective is at work in our setting [30]). Such properties could be extended (possibly in several ways) to the case of vectors of formulas (see [30]) but would fail to be satisfied even in the simple case when B contains a single formula (i.e., B is a vector of dimension 1) and the standard forgetting context for it is considered.

It is also easy to show that \approx_{\sqsubseteq}^C satisfies:

- If $\mathcal{R}_C(B) \neq \emptyset$ then $B \not\approx_{\sqsubseteq}^C \perp$;
- If \sqsubseteq is such that $\vec{\emptyset} = \langle \emptyset, \dots, \emptyset \rangle$ is the unique minimal element of $\text{Pref}(\mathcal{R}_C(B), \sqsubseteq)$ when it belongs to it, then if B is consistent, we have for any formula ψ , $B \models \psi$ if and only if $B \approx_{\sqsubseteq}^C \psi$.

The first property shows that a weak form of paraconsistency is achieved by skeptical forget-based inference, while the second one shows that under some conditions, a valuable information preservation property is also satisfied. These conditions are guaranteed provided that the base, the context and/or the preference relation satisfy some rather weak conditions:

- If every ϕ_i in B is consistent and C is upward normal then $B \not\approx_{\sqsubseteq}^C \perp$;
- If B is consistent, \sqsubseteq is strictly monotonic and C is downward normal, then for any formula ψ , $B \models \psi$ if and only if $B \approx_{\sqsubseteq}^C \psi$.

For the first point, if C is upward normal then $\langle PS, \dots, PS \rangle$ is a recovery for B given C . For the second point, if C is downward normal then $\langle \emptyset, \dots, \emptyset \rangle$ is a recovery for a consistent B given C ; since \sqsubseteq is strictly monotonic, $\langle \emptyset, \dots, \emptyset \rangle$ is the unique preferred recovery for B given C .

Observe that considering the standard forgetting context C_S and \sqsubseteq_p as a preference relation on forgetting vectors is enough to ensure both conditions.

We may want to compare the sets of forget-based inference relations obtained respectively by letting the set of forgetting contexts vary and by letting the set of preference vary (that we may call the “discriminative power” offered respectively by forgetting contexts and by preference relations). As a matter of fact, when skeptical inference is considered, the discriminative power offered by forgetting contexts is the same as the discriminative power offered and by preference relations coincide: every skeptical consequence of a base B given a forgetting context C for B and a preference relation \sqsubseteq on \mathcal{F}_C can be recovered as a skeptical consequence of B given the standard forgetting context C_S and a preference relation \sqsubseteq' on \mathcal{F}_{C_S} which includes \sqsubseteq and satisfies the following two conditions: (i) for any two forgetting vectors \vec{V}, \vec{V}' such that \vec{V} satisfies C and \vec{V}' does not, then $\vec{V} \sqsubseteq' \vec{V}'$; and (ii) for any two forgetting vectors \vec{V}, \vec{V}' such that both \vec{V} and \vec{V}' satisfy C , then $\vec{V} \sqsubseteq' \vec{V}'$ holds if and only if $\vec{V} \sqsubseteq \vec{V}'$. (Clearly, such a preference relation \sqsubseteq' exists.) Then by construction the preferred recoveries for B given C w.r.t. \sqsubseteq coincide with the preferred recoveries for B given C_S w.r.t. \sqsubseteq' .

This result typically shows that the two approaches pointed out in order to focus on expected recoveries (namely, forgetting contexts and preference relations) are equivalent from an expressiveness point of view when skeptical inference is considered. However, keeping both concepts in the setting makes sense. On the one hand, as mentioned previously, some significant computational effort must be spent to go from a preference criterion (e.g. “prefer homogeneous recoveries”) and a base to the corresponding set of preferred recoveries; stated otherwise, since preference relations are typically not part of the input, generating a forgetting context which qualifies only the preferred recoveries can be very expensive, especially from the point of view of computational space. On the other hand, forgetting contexts are more explicit (they are part of the input) and are useful for specifying hard constraints on the recoveries.

4. On the generality of our framework

We now show how several well-known paraconsistent inference relations belong to the family of forget-based relations. The key idea is that forget-based weakening is fine-grained enough to emulate other weakening mechanisms at work in

previous approaches to reasoning under inconsistency, especially formula inhibition (considered in approaches based on preferred consistent subbases), formula dilation (considered in distance-based merging) and abstraction-based techniques.

4.1. Reasoning from preferred consistent subbases

Let us first recall the definition of inferences drawn from preferred consistent subbases also called *maxcons* inference relations (see e.g. [46,24,25,11,3,44,4]).

For avoiding heavy definitions, we consider only bases consisting of consistent formulas. This is without loss of generality when considering maxcons inference relations, since inconsistent formulas do not participate to any consistent subbases; accordingly, they can be removed from the base at start without any change on the inference relation.

Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a finite and non-empty vector of formulas such that each ϕ_i ($1 \leq i \leq n$) is consistent. For any $X \subseteq \{1, \dots, n\}$ we denote $B_X = \{\phi_i \mid i \in X\}$. Let $\text{Cons}(B) = \{X \subseteq \{1, \dots, n\} \mid B_X \text{ is consistent}\}$ – $\text{Cons}(B)$ is isomorphic to the set of consistent subbases of B .

Let \succcurlyeq be a preorder on subsets of $\{1, \dots, n\}$ (also called preference relation). Typically, \succcurlyeq is monotonic, and even strictly monotonic, with respect to set containment; however, we do not have to require this. Among usual preference relations, we find set inclusion: $X \succcurlyeq Y$ if and only if $X \supseteq Y$, and cardinality: $X \succcurlyeq Y$ if and only if $|X| \geq |Y|$. Let $\text{Pref}(\text{Cons}(B), \succcurlyeq)$ be the set of all maximal (w.r.t. \succcurlyeq) elements of $\text{Cons}(B)$. Then, by definition:

An inference relation $\approx_{\succcurlyeq}^{\forall}$ is a *maxcons* inference relation if and only if for every base B consisting of consistent formulas and every formula ψ , we have $B \approx_{\succcurlyeq}^{\forall} \psi$ if and only if $\forall X \in \text{Pref}(\text{Cons}(B), \succcurlyeq)$, $B_X \models \psi$.

There are two ways of characterizing maxcons inference relations as forget-based inference relations: either by imposing a condition on the preference relation (Proposition 4.1) or by imposing a condition on the forgetting context (Proposition 4.2).

Proposition 4.1. *An inference relation \approx is a maxcons inference relation if and only if for every base B consisting of consistent formulas, there exists a preference relation \sqsubseteq on $\mathcal{F}_{\mathcal{C}_S}(B)$ satisfying binaricity and such that for any formula ψ , $B \approx \psi$ if and only if $B \approx_{\sqsubseteq}^{\mathcal{C}_S} \psi$, where \mathcal{C}_S is the standard forgetting context.*

Proposition 4.2. *An inference relation \approx is a maxcons inference relation if and only if for every base B consisting of consistent formulas, there exists a preference relation \sqsubseteq on $\mathcal{F}_{\mathcal{C}_B}(B)$ such that for any formula ψ , $B \approx \psi$ if and only if $B \approx_{\sqsubseteq}^{\mathcal{C}_B} \psi$, where \mathcal{C}_B is the logically weakest binary context, defined by $\bigwedge_{i=1}^n \bigwedge_{x,y \in \mathcal{P}S} \text{forget}(x, i) \Rightarrow \text{forget}(y, i)$.*

The first direction of Proposition 4.1 (respectively Proposition 4.2) expresses that with a preference relation satisfying binaricity and the standard forgetting context (respectively, with the logically weakest binary forgetting context), inference from preferred recoveries comes down to inference from preferred maximal consistent subbases. The other direction of Propositions 4.1 and 4.2 states that the set of all maxcons inference relations is contained in the set of forget-based inference relations, and therefore that the latter family is at least as general as the former. In particular, if preference over forgetting vectors is defined by $\forall \vec{V}, \vec{V}' \in \mathcal{F}_{\mathcal{C}_S}(B)$, $\vec{V} \sqsubseteq \vec{V}'$ if and only if $\{i \mid V_i \neq \emptyset\} \subseteq \{i \mid V'_i \neq \emptyset\}$ then $B \approx_{\sqsubseteq}^{\mathcal{C}_S} \psi$ if and only if $S \models \psi$ holds for any maximal (w.r.t. set inclusion) consistent subbase S of B . Other criteria such that maximum cardinality, “discrimin” or “leximin”, or minimum penalty, can be recovered as well. The assumption that each ϕ_i ($1 \leq i \leq n$) in B is consistent is necessary (and sufficient) to ensure that B is recoverable given its standard forgetting context. It can be made without loss of generality, just because B and $B^* = \langle \phi_1^*, \dots, \phi_n^* \rangle$, where $\phi_i^* = \phi_i$ if ϕ_i^* is consistent and $\phi_i^* = \top$ otherwise, have the same consistent subbases, and each formula in B^* is consistent. As an immediate consequence of Propositions 4.1 and 4.2, our framework encompasses other important frameworks, like supernormal default theories with priorities [11] and syntax-based belief revision [41], as specific cases.

It is easy to show that a generalization of Proposition 4.1 (resp. Proposition 4.2) holds, replacing the forgetting context \mathcal{C}_S (resp. \mathcal{C}_B) by any logically stronger context \mathcal{C}_P . The resulting forget-based inferences correspond to maxcons inferences with protected formulas. Here is an example of such relations:

Example 4.1. Let us step back to Example 3.1. Let \mathcal{C}_P be defined by $\bigwedge_{x \in \mathcal{P}S} \neg \text{forget}(x, 1)$. Define \sqsubseteq_I and \sqsubseteq_C by

- $\vec{V} \sqsubseteq_I \vec{V}'$ if and only if $\{i \mid V_i \neq \emptyset\} \subseteq \{i \mid V'_i \neq \emptyset\}$;
- $\vec{V} \sqsubseteq_C \vec{V}'$ if and only if $|\{i \mid V_i \neq \emptyset\}| \leq |\{i \mid V'_i \neq \emptyset\}|$.

Clearly, \sqsubseteq_I and \sqsubseteq_C both satisfy the binaricity property.

Consider the preferred recoveries for \sqsubseteq_I listed below, together with their projection on B (in fact, there exist other preferred recoveries, but each of them is pointwise included in one of these listed below):

$$\left[\begin{array}{l} \langle \emptyset, PS, \emptyset, \emptyset, \emptyset \rangle \\ \langle \emptyset, \emptyset, PS, \emptyset, \emptyset \rangle \\ \langle \emptyset, \emptyset, \emptyset, PS, PS \rangle \end{array} \middle| \begin{array}{l} \phi_1 \wedge s \wedge t \wedge s_b \wedge c_2 \\ \phi_1 \wedge s \wedge s_r \wedge \neg t \wedge c_1 \\ \phi_1 \wedge \neg s \wedge t \wedge c_1 \end{array} \right].$$

Therefore we have $B \approx_{\sqsubseteq_I}^{\mathcal{C}_P} (s \wedge t) \Rightarrow s_b$ and $B \approx_{\sqsubseteq_I}^{\mathcal{C}_P} \neg t \Rightarrow s_r$.

With \sqsubseteq_C , the two preferred recoveries are $\langle \emptyset, PS, \emptyset, \emptyset, \emptyset \rangle$ and $\langle \emptyset, \emptyset, PS, \emptyset, \emptyset \rangle$. Therefore, $B \approx_{\sqsubseteq_C}^{\mathcal{C}_P} s$ holds.

4.2. Belief merging

A merging operator (see e.g. [37,47,31,40]) maps any vector $B = \langle \phi_1, \dots, \phi_n \rangle$ of consistent propositional formulas into a new consistent propositional formula $Merge(B)$, also viewed as its set of models. When n is fixed, we define an n -merging operator $Merge_n(B)$ as the restriction of $Merge$ to n -uples of formulas. In order to simplify notations, we will write $Merge(\langle \phi_1, \dots, \phi_n \rangle)$ instead of $Merge_n(\langle \phi_1, \dots, \phi_n \rangle)$ in the following.

4.2.1. Distance-based belief merging

A special class of n -merging operators is the class of decomposable merging operators induced by (pseudo-)distances and aggregation functions (see [29]). Let $d : \Omega \times \Omega \rightarrow \mathbb{N}$ be a total function, called a pseudo-distance, satisfying $\forall \omega, \omega' \in \Omega, d(\omega, \omega') = 0$ if and only if $\omega = \omega'$, and $d(\omega, \omega') = d(\omega', \omega)$. Let \star be a total function from \mathbb{N}^n to \mathbb{N} , monotonic in each of its arguments; it is called an aggregation function.

For every $B = \langle \phi_1, \dots, \phi_n \rangle$ such that each ϕ_i ($1 \leq i \leq n$) is consistent, the decomposable n -merging operator $Merge_d^\star$ induced by d and \star is defined by $Merge_d^\star(B) = \{\omega \in \Omega \mid d(\omega, B) \text{ is minimal}\}$, where $d(\omega, B) = \star(d(\omega, \phi_1), \dots, d(\omega, \phi_n))$ and $d(\omega, \phi_i) = \min_{\omega' \models \phi_i} d(\omega, \omega')$. Then, by definition, an inference relation $\approx_{d, \star}^M$ is the n -merging inference relation $\approx_{d, \star}^M$ induced by d and \star if and only if for every base B consisting of n consistent formulas and every formula ψ , we have:

$$B \approx_{d, \star}^M \psi \quad \text{if and only if} \quad Merge_d^\star(B) \subseteq Mod(\psi).$$

In order to characterize the decomposable n -merging inference relations as forget relations, we need to focus on those based on a differential pseudo-distance [32]. A pseudo-distance d is said to be differential if and only if there exists a total function $f : 2^{PS} \rightarrow \mathbb{N}$ satisfying (i) $f(\emptyset) = 0$ and (ii) $\forall A, B \subseteq PS$, if $A \subseteq B$, then $f(A) \leq f(B)$, and such that $\forall \omega, \omega' \in \Omega, d(\omega, \omega') = f(Diff(\omega, \omega'))$. In simpler words, if d is a differential pseudo-distance then $d(\omega, \omega')$ is determined by the set of propositional atoms on which ω and ω' differ. The literature on merging mainly makes use of two pseudo-distances: the Hamming distance d_H , which counts the number of symbols in which the two models differ, and the so-called drastic distance d_D , which is 1 as soon as the two models differ (and 0 otherwise). These two distances are differential: for d_H we have $f(X) = |X|$, and for d_D we have $f(\emptyset) = 0$ and $f(X) = 1$ if $X \neq \emptyset$.

Proposition 4.3. \approx is a decomposable n -merging inference relation induced by a differential pseudo-distance and an aggregation function if and only if for every base B containing n formulas, there exists a complete preference relation \sqsubseteq_μ on $\mathcal{F}_{\mathcal{C}_S}(B)$, induced by a ranking function μ satisfying decomposability, such that for any formula ψ , $B \approx \psi$ if and only if $B \approx_{\sqsubseteq_\mu}^{\mathcal{C}_S} \psi$.

The first direction of Proposition 4.3 expresses that under the assumptions of completeness of \sqsubseteq_μ , decomposability, and under the standard forgetting context, forget-based inferences come down to inferences defined from a merging operator. The other direction expresses that a particularly interesting subclass of inferences from merging operators is contained in the set of forget-based inference relations. In particular, usual arbitration and majority merging operators (see e.g. [37,31]) are recovered by letting (1) $\mu_i(A) = |A|$ for each $i \in \{1, \dots, n\}$, which implies that the induced distance d is the Hamming distance between interpretations, and (2) \star is a standard aggregation function, like max, leximax or +. Integrity constraints can also be expressed by adding to the context a specific constraint whose effect is to protect formulas (similarly as in the last paragraph of Section 4.1).

Example 4.2. Let us consider Example 3.1 again. We now take $\mathcal{C} = \mathcal{C}_S$ and

$$\mu(V) = \max(10 \cdot |V_1|, |V_2|, |V_3|, |V_4|, |V_5|).$$

The role of factor 10 in $10 \cdot |V_1|$ is to ensure that the integrity constraint ϕ_1 has priority over all other formulas: since $|PS| < 10$, it will always be worse to forget a single atom from ϕ_1 than forgetting any number of variables from ϕ_2, \dots, ϕ_5 .

μ satisfies decomposability: $\mu_0(X) = |X|$ and $H(p_1, p_2, p_3, p_4, p_5) = \max(10 \cdot p_1, p_2, p_3, p_4, p_5)$.

Because ϕ_1 is consistent, the preferred recoveries are those for which $V_1 = \emptyset$ and the maximum number of variables to be forgotten in each of the ϕ_i 's is minimal. Forgetting at most one variable in each ϕ_i (for $i \geq 2$) allows for recovering consistency: it suffices either (a) to forget c_0 or c_1 in ϕ_2 and s or s_r in ϕ_3 , or (b) to forget t in ϕ_5 and either s or

s_r in ϕ_2 or s or s_b in ϕ_3 . This gives many preferred recoveries (all satisfying $\mu(V) = 1$), among which $\langle \emptyset, \{s_r\}, \emptyset, \emptyset, \{t\} \rangle$, $\langle \emptyset, \{c_0\}, \{s\}, \{s\}, \{s\} \rangle$, $\langle \emptyset, \{c_0\}, \{s\}, \{s\}, \{s\} \rangle$, and so on.

We can check that $B \approx_{\sqsubseteq_{\mu}}^{C_S} s \vee t$ and $B \approx_{\sqsubseteq_{\mu}}^{C_S} (s \wedge t) \Rightarrow s_r$. This last inference is an effect of equity (reflected by the aggregation function): if agent 2 is forced to give up her subgoal $c_0 \vee c_1$ (by forgetting c_0 or c_1) then she is ensured to see her other subgoal $s \Rightarrow s_r$ satisfied.

Finally, while distance-based merging can be extended to non-classical propositional settings (by considering multiple-valued semantics as in [2]), Proposition 4.3 cannot be generalized directly to those settings; indeed, the notion of variable forgetting we use is anchored on classical logic; it prevents from deriving non-trivial information from an inconsistent formula, while some propositional settings based on multiple-valued semantics do not suffer from this drawback.

4.2.2. Belief merging à la Delgrande and Schaub

Another interesting merging operator which proves to be a particular case of forget-based skeptical inference is one of the operators proposed by Delgrande and Schaub [17] and based on variable renaming. For the sake of simplicity, we do not consider integrity constraints.

Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base where each ϕ_i ($i \in \{1, \dots, n\}$) is consistent. Let $B^r = \bigwedge_{i=1}^n \text{rename}(\phi_i, i)$ where $\text{rename}(\phi_i, i)$ is the formula obtained from ϕ_i by replacing each occurrence of each atom x appearing in ϕ_i by x_i . Let $EQ = \bigcup_{x \in \text{Var}(B)} \{x_i \Leftrightarrow x_0 \mid i \in \{1, \dots, n\}\}$. A *fit* for B is a subset X of EQ consistent with B^r . X is a *maximal fit* for B if and only if it is a fit for B and none of its strict supersets is a fit for B . Lastly, the Delgrande–Schaub merging $\Delta^{DS}(B)$ is defined by, for any formula ψ ,

$$\Delta^{DS}(B) \models \psi \quad \text{if and only if for any maximal fit } X \text{ for } B, \quad B^r \wedge \bigwedge_{x \in X} x \models \text{rename}(\psi, 0).$$

Example 4.3. Let $B = \langle a \wedge b \wedge c, \neg a, \neg a \vee \neg b \rangle$. There are four maximal fits for B : $EQ \setminus \{a_1 \Leftrightarrow a_0\}$; $EQ \setminus \{a_2 \Leftrightarrow a_0, a_3 \Leftrightarrow a_0\}$; $EQ \setminus \{a_2 \Leftrightarrow a_0, b_3 \Leftrightarrow b_0\}$; $EQ \setminus \{a_2 \Leftrightarrow a_0, b_1 \Leftrightarrow b_0\}$. We can check that $\Delta^{DS}(B) \equiv c \wedge (a \vee b)$.

This merging operator turns out to have a very simple characterization in terms of forgetting-based inference:

Proposition 4.4. Let B be a base consisting of consistent formulas and ψ be a formula. $\Delta^{DS}(B) \models \psi$ if and only if $B \approx_{\sqsubseteq}^{C_S} \psi$, where C_S is the standard forgetting context and $\sqsubseteq = \subseteq_p$.

A similar result would hold with integrity constraints: let I be a formula representing some (hard) integrity constraints, then B would contain a formula $\phi_{n+1} = \text{rename}(I, n+1)$ in B , and the forgetting context would forbid to forget any atom in ϕ_{n+1} .

4.2.3. Conflict-based belief merging

In [23], conflict-based merging operators have been introduced. In this setting, the “distance” between an interpretation and a formula is not evaluated as a numerical measure of the conflicts between them (for instance the minimum number of variables to be switched in the interpretation to make it a model of the formula) as usual in distance-based merging, but by the conflicts themselves, i.e., the minimal sets (for set inclusion) of propositional atoms which differ between the interpretation and the closest models of the formula. Given a profile of such formulas (a base B), each interpretation ω can be associated with a set $\text{diff}(\omega, B)$ of vectors of sets of propositional atoms (i.e., a set of forgetting vectors) such that ω is a model of the corresponding recoveries. This set of vectors expresses in some sense how close ω is to B and preference relations \sqsubseteq_B over such sets are defined in order to define the models of the merging $\Delta_{\mu}^{\text{diff}}(B)$ (as the minimal models of a given integrity constraint μ w.r.t. \sqsubseteq_B). Formally:

Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base consisting of consistent formulas ϕ_i ($i \in \{1, \dots, n\}$). Let ω, ω' be interpretations from Ω , and let μ be a formula.

- A *conflict vector* \vec{V} for ω given B is a recovery for B given C_S such that $\omega \models B \mid \vec{V}$ and for no recovery \vec{V}' for B given C_S such that $\omega \models B \mid \vec{V}'$ we have $\vec{V}' \subset_p \vec{V}$;
- $\text{diff}(\omega, B)$ is the set of all conflict vectors for ω given B ;
- Let \sqsubseteq be a binary relation on forgetting vectors for B given C_S . $\omega \sqsubseteq_B \omega'$ if and only if $\exists \vec{V}_{\omega} \in \text{diff}(\omega, B)$ s.t. $\forall \vec{V}_{\omega'} \in \text{diff}(\omega', B)$, we have $\vec{V}_{\omega} \sqsubseteq \vec{V}_{\omega'}$;
- $\Delta_{\mu}^{\text{diff}}(B)$ is given by $\text{Mod}(\Delta_{\mu}^{\text{diff}}(B)) = \min(\text{Mod}(\mu), \sqsubseteq_B)$.

Conflict-based merging can be recovered as a specific forget-based inference relation whenever \sqsubseteq is a complete preorder including \subseteq_p . For the sake of simplicity, we assume that there are no integrity constraints, i.e., μ is a valid formula. Then we have:

Table 2

Preference relation \sqsubseteq	Preferred recoveries that are maximal w.r.t. \subseteq_p	$\approx_{\sqsubseteq}^{\mathcal{C}}$
\sqsubseteq_1	$\vec{V}^1, \vec{V}^2, \vec{V}^3$	$s \vee t$
\sqsubseteq_2	$\vec{V}^1, \vec{V}^2, \vec{V}^6$	$s \oplus t$

Proposition 4.5. *Let B be a base consisting of consistent formulas and let \sqsubseteq be a complete preorder on $\mathcal{F}_{\mathcal{C}_S}(B)$ which includes \subseteq_p . For any formula ψ , we have $\Delta_{\top}^{diff}(B) \models \psi$ if and only if $B \approx_{\sqsubseteq}^{\mathcal{C}_S} \psi$.*

Again, a similar result would hold in presence of an integrity constraint (which should be treated as a protected formula).

4.3. Homogeneous forget-based inferences

Up to now, we showed in this section that reasoning from preferred consistent subbases and some forms of merging are particular cases of forget-based inference relations. Now, there are many other forget-based inference relations which do not belong to any of the previous families.

A first example is the family of *homogeneous forget-based inference relations*. Such relations can be obtained by requiring the forgetting context to be homogenous, or alternatively, requiring the preference relation \sqsubseteq to satisfy the homogeneity property (which is less demanding).

When \mathcal{C} is homogeneous, it is sufficient to focus on homogeneous recoveries, of the form $\vec{V} = \langle V, \dots, V \rangle$. In such a context, we abuse notations and simply denote \vec{V} by the set V of variables that are uniformly forgotten. Thus, $(\langle \phi_1, \dots, \phi_n \rangle \mid V) = \bigwedge_{i=1}^n \exists V. \phi_i$. Likewise, when homogeneous inference relations are considered, it is enough to define a preference relation on subsets of PS (instead of vectors of subsets of PS). Many such relations can be obtained by letting the preference relation on 2^{PS} vary, in quite the same manner as maxcons inference relations are obtained by letting the preference relation on subsets vary. We may for instance minimize the *set* of forgotten variables (i.e., $V \sqsubseteq V'$ if and only if $V \subseteq V'$); we call the associated inference relation the *prototypical homogeneous forget-based inference relation* given \mathcal{C} . Alternatively, we may minimize the *number* of forgotten variables, or, more generally, make use of a predefined penalty function or priority preordering on variables.

Example 4.4. We step back to Example 3.1 again. Consider the preference relations \sqsubseteq_1 and \sqsubseteq_2 below on 2^{PS} and the induced homogeneous inference relations $\approx_{\sqsubseteq_1}^{\mathcal{C}}$ and $\approx_{\sqsubseteq_2}^{\mathcal{C}}$. We define the mapping $k : 2^{Var(B)} \rightarrow \mathbb{N}$ by $k = k_c + k_s + k_t$ where $k_c(V) = 3$ if $\{c_0, c_1, c_2\} \subseteq V$, $k_c(V) = 0$ otherwise; $k_s(V) = 3$ if $\{s, s_b, s_r\} \subseteq V$, $k_s(V) = 1$ if $V \cap \{s, s_b, s_r\} = \{s_b, s_r\}$, and $k_s(V) = 0$ otherwise; $k_t(V) = 2$ if $t \in V$, $k_t(V) = 0$ otherwise.

- \sqsubseteq_1 : $V \sqsubseteq_1 V'$ if and only if $V \subseteq V'$;
- \sqsubseteq_2 : $V \sqsubseteq_2 V'$ if and only if $k(V) \leq k(V')$.

\sqsubseteq_1 minimizes the set of atoms forgotten everywhere while \sqsubseteq_2 minimizes their cumulated cost.

The results are synthesized in Table 2. The rightmost column gives the most general consequence of B w.r.t. $\approx_{\sqsubseteq}^{\mathcal{C}}$, which can be represented using the two atoms, s and t , only.

4.4. Abstraction-based inferences

In many situations, the propositional variables have different levels of abstraction, and inconsistency can be avoided by forgetting more specific variables and keeping more abstract ones.

Example 4.5. A theft has been committed in a house, three witnesses saw somebody leaving it through a window. The first witness saw a young (y) man (m) with a grey (g) cap (c), the second one an old (o) man with a black (bl) bonnet (bo) and the third one a lady (l) with a brown (br) beret (be). The three witnesses are highly contradictory since man/lady, old/young, cap/bonnet/beret, grey/black/brown are conflicting pieces of information. Nevertheless some significant pieces of information can be derived by abstracting away the concepts occurring in the three witnesses: all witnesses agree that they saw a person (p) with a dark (d) hat (h).

The point is that such inferences are forget-based inferences. Indeed, the replacement scheme at work can be achieved by first conjoining the contents of each source of information with a formula, representing for instance a taxonomy and linking specific concepts to more general ones, then forgetting concepts in each resulting formula, from the most specific ones to the most general ones, until consistency is recovered.

Here the initial base is $B = \langle \phi_1, \phi_2, \phi_3 \rangle$ with $\phi_1 = y \wedge m \wedge g \wedge c$, $\phi_2 = o \wedge m \wedge bl \wedge bo$, $\phi_3 = l \wedge br \wedge be$. The theory linking the concepts together can be represented by $\phi = (\neg m \vee \neg l) \wedge (\neg o \vee \neg y) \wedge (\neg c \vee \neg bo) \wedge (\neg c \vee \neg be) \wedge (\neg bo \vee \neg be) \wedge (\neg g \vee \neg bl) \wedge (\neg g \vee \neg br) \wedge (\neg bl \vee \neg br) \wedge ((m \vee l) \Rightarrow p) \wedge ((c \vee bo \vee be) \Rightarrow h) \wedge ((g \vee bl \vee br) \Rightarrow d)$.

B is first turned into the base $\langle \phi_1 \wedge \phi, \phi_2 \wedge \phi, \phi_3 \wedge \phi \rangle$. A stratification (i.e., a totally ordered partition) $S = \langle C_1, \dots, C_k \rangle$ of $\text{Var}(\phi)$ can then be defined so that the less significant concepts (e.g. the most specific ones) are ranked first, for instance it can be $S = \langle \{y, o, c, bo, be, g, bl, br\}, \{m, l\}, \{p, d, h\} \rangle$. The abstraction of level i ($i \in \{1, \dots, k\}$) of each $\phi_j \wedge \phi$ ($j \in \{1, \dots, n\}$) w.r.t. S is the most general consequence of $\phi_j \wedge \phi$ which is independent of $\bigcup_{l=1}^{i-1} C_l$, namely $\exists \bigcup_{l=1}^{i-1} C_l. (\phi_j \wedge \phi)$. Now, a forgetting vector \vec{V} is abstraction-based given S when each of the V_j ($j \in \{1, \dots, n\}$) in it is equal to $\bigcup_{l=1}^{i-1} C_l$ for some i ($i \in \{1, \dots, k\}$), called the abstraction level of V_j , and noted $a(V_j)$. Back to the example, we have $C_1 = \{y, o, c, bo, be, g, bl, br\}$, $C_2 = \{m, l\}$ and $C_3 = \{p, d, h\}$; and the following forgetting vectors are abstraction-based:

- $\vec{V}_1 = \langle C_1 \cup C_2, C_1 \cup C_2, C_1 \cup C_2 \rangle$, for which we have $(B \mid \vec{V}_1) \equiv \exists(C_1 \cup C_2). (\varphi \wedge \varphi_1) \wedge \exists(C_1 \cup C_2). (\varphi \wedge \varphi_2) \wedge \exists(C_1 \cup C_2). (\varphi \wedge \varphi_3) \equiv p \wedge d \wedge h$;
- $\vec{V}_2 = \langle C_1, C_1, C_1 \cup C_2 \rangle$, for which we have $(B \mid \vec{V}_2) \equiv \exists C_1. (\varphi \wedge \varphi_1) \wedge \exists C_1. (\varphi \wedge \varphi_2) \wedge \exists(C_1 \cup C_2). (\varphi \wedge \varphi_3) \equiv m \wedge p \wedge d \wedge h$;
- $\vec{V}_3 = \langle C_1 \cup C_2, C_1 \cup C_2 \cup C_3, \emptyset \rangle$, for which we have $(B \mid \vec{V}_3) \equiv \exists(C_1 \cup C_2). (\varphi \wedge \varphi_1) \wedge \exists(C_1 \cup C_2 \cup C_3). (\varphi \wedge \varphi_2) \wedge \exists \emptyset. (\varphi \wedge \varphi_3) \equiv l \wedge br \wedge be \wedge p \wedge d \wedge h$.

Note that \vec{V}_1 is homogeneous, whereas \vec{V}_2 and \vec{V}_3 are not.

The stratification S can be exploited in different ways in our setting. Considering the standard forgetting context for the resulting base, a preference relation \sqsubseteq_S over forgetting vectors can be defined so that it is such that for any abstraction-based forgetting vectors \vec{V} and \vec{V}' , $\vec{V} \sqsubseteq_S \vec{V}'$ when $\star(a(V_1), \dots, a(V_n)) \leq \star(a(V'_1), \dots, a(V'_n))$ for some aggregation function \star (like max, leximax or +). In addition to this, the forgetting context may also be required to be homogeneous. The consequences p, d, h (and sometimes others) can then be derived using skeptical forget-based inference.

Apart of the families of forget-based inferences to which we devoted some attention in Sections 4.1, 4.2 and 4.3, many other interesting families of forget-based inference relations can be defined. For the sake of illustration, we give here an example of such a family, which is suitable in the situation where reliability between sources is a matter of *topic*. (Note that this family does not degenerate into one of the previous types of inference.)

Here, a topic simply is a propositional atom x and a source i conveys some information about it as soon as ϕ_i is not independent of x , i.e., $\phi_i \not\equiv \exists\{x\}. \phi_i$,¹ as explained in the following example:

Example 4.6. Let $B = \langle \phi_1, \phi_2 \rangle$, with $\phi_1 = a \wedge b$, $\phi_2 = \neg a \wedge \neg b$. Let us consider the standard forgetting context C_S for B . Assume that source 1 is more reliable to what concerns a than to what concerns b and conversely, source 2 is more reliable to what concerns b than to what concerns a . To capture this in our framework, the following ranking function can be considered (among other possible choices): $\mu((V_1, V_2)) = \mu_0(V_1) + \mu_0(V_2)$ where $\mu_0(V_i) = \sum_{v \in V_i} k_i(v)$ and $k_1(a) = 2$, $k_1(b) = 1$, $k_2(a) = 1$, $k_2(b) = 2$. Then the (unique) preferred recovery for B is $\langle \{b\}, \{a\} \rangle$ and therefore we have $B \approx_{\sqsubseteq_{\mu}^{C_S}} a \wedge \neg b$.

5. Computational complexity

We now investigate the computational complexity of forget-based inference relations. We suppose the reader familiar with computational complexity (see e.g. [43]), and especially the complexity classes NP, Δ_2^P and Π_2^P of the polynomial hierarchy.

We have obtained the following results:

Proposition 5.1.

- (1) Given a base B and a forgetting context C for it, determining whether B is recoverable is NP-complete.
- (2) Provided that the preference relation \sqsubseteq on $\mathcal{F}_{C_S}(B)$ can be decided in polynomial time, the inference problem associated to \approx_{\sqsubseteq}^C is in Π_2^P .
- (3) If moreover \sqsubseteq is induced by a ranking function μ computable in polynomial time, the inference problem associated to \approx_{\sqsubseteq}^C is in Δ_2^P .

These results show that the gain in generality and flexibility offered by our framework does not induce a complexity shift compared with many more specific approaches, such as maxcons inferences, that are already Π_2^P -complete when maximality is with respect to set-inclusion. This holds as well for distance-based merging where the inference problems associated to some operators are Δ_2^P -complete in the general case (see [38,27]), while the problem of forget-based inference with a complete preference relation induced by a polytime ranking function is not harder. Hardness results for our forget-based inference relations (in the general case) can be easily obtained from known results for maxcons inference relations, belief revision or merging (see e.g. [18,13,42,27]), thanks to the forget-based characterizations of such approaches, as reported in Section 4. Hardness also holds for some prototypical homogeneous forget-based inference relations. Since homogeneous forget-based inferences are new, we establish and prove the corresponding result.

¹ For a different approach to inconsistency handling where topics are taken into account, see [14].

Proposition 5.2. *Skeptical inference for the prototypical homogeneous forget-based inference given a forgetting context \mathcal{C} is in Π_2^P . It is Π_2^P -complete for some forgetting contexts \mathcal{C} .*

In the general case (especially, when \sqsubseteq is not induced by a ranking function), our complexity results show that two independent sources of complexity must be dealt with. One of them lies in the possibly exponential number of preferred recoveries which must be taken into account. The other one originates from the complexity of classical entailment in propositional logic.

One way to circumvent this intractability consists in *compiling* the base (see e.g. [12,15]). Such a compilation may consist in computing off-line the set of preferred recoveries, or only the maximal elements of it w.r.t. \subseteq_p . In this situation, the source of complexity due to the number of preferred recoveries disappears (in practice, this is significant only when the number is small enough, which is always the case when prudent inference is considered); accordingly, the complexity of (skeptical or prudent) inference goes down to coNP. The other source of complexity can also be removed by imposing some restrictions on the base B . In this situation, on-line forget-based inference is tractable.

Proposition 5.3.

- (1) *Provided that $\text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \sqsubseteq)$ – or only the maximal elements of it w.r.t. \subseteq_p – is part of the input, the inference problem associated to $\approx_{\sqsubseteq}^{\mathcal{C}}$ (or $\approx_{\sqsubseteq}^{\mathcal{C}_{\text{prudent}}}$) is in coNP.*
- (2) *If moreover each formula from B belongs to a propositional fragment (a subset of PROP_{PS}) that is tractable for clausal entailment, stable for conjunction, and stable for new variable renaming,² the inference problem associated to $\approx_{\sqsubseteq}^{\mathcal{C}}$ (or $\approx_{\sqsubseteq}^{\mathcal{C}_{\text{prudent}}}$) is in P when CNF queries are considered.*

Interestingly, several well-known classes that are tractable for clausal entailment are also stable for conjunction and for new variable renaming; for example, this is the case for the class of Krom formulas (CNF formulas consisting of binary clauses) and for the class of CNF Horn formulas.

6. Discussion

Since the family of forget-based inferences covers many different subfamilies of ways of remedying inconsistency, it is worth to comparing them on simple examples and discuss their respective merits informally.

Let us start by considering the following two bases:

Example 6.1. Let $B = \langle \phi_1, \phi_2 \rangle$, and let $B' = \langle \phi'_1, \phi'_2 \rangle$, with:

	$\phi_1 = (a \vee b) \wedge c$	$\phi'_1 = a \wedge b$
	$\phi_2 = \neg a \wedge \neg b$	$\phi'_2 = \neg a \wedge \neg b$
\sqsubseteq_1	$\phi_1 \vee \phi_2$	$a \Leftrightarrow b$
\sqsubseteq_2	$c \wedge (\neg a \vee \neg b)$	\top
\sqsubseteq_3	$(a \vee b \vee c) \wedge (\neg a \vee b)$	$a \not\approx b$
\sqsubseteq_4	$c \wedge (\neg a \vee \neg b)$	\top
\sqsubseteq_5	$c \wedge (\neg a \vee \neg b)$	\top

Each cell of the second column (resp. third column) of this table contains a formula equivalent to the conjunctively-interpreted set gathering all consequences of B (resp. B') w.r.t. skeptical forget-based inference.

What should be concluded from the inconsistent bases B and B' , especially when no information about the (absolute or relative to topics) reliability of the two sources are available? There is, of course, no absolute answer to this question. The result depends on various assumptions on the merging process, and the variety of conclusions obtained through many paraconsistent inference relations, as reported in the table, exemplifies it. The important point here is that our framework enables to capture each of them (here the standard forgetting context is considered for both bases):

- \sqsubseteq_1 corresponds to inference from maximal (w.r.t. \subseteq) consistent subsets.
- \sqsubseteq_2 (resp. \sqsubseteq_3) corresponds to distance-based merging where d is the Hamming distance and $\star = +$ (resp. $\star = \max$).
- $\sqsubseteq_4 = \subseteq_p$ corresponds to merging à la Delgrande–Schaub.
- \sqsubseteq_5 is the prototypical homogeneous inference relation given $\mathcal{C}_{\mathcal{S}}$.

² This means that for every formula ϕ from PROP_{PS} and for every substitution $\text{rename}(\dots)$ such that $\text{rename}(\phi, V)$ is the formula obtained by renaming in a uniform way every variable from V occurring in ϕ into a new variable (not occurring in ϕ) if ϕ belongs to the class under consideration, this is also the case for $\text{rename}(\phi, V)$ whatever V .

Now, let us have a critical examination of the two well-known families of weakening-based approaches to paraconsistent reasoning, namely coherence-based approaches and distance-based merging.

Coherence-based approaches. This quite popular family of approaches relies on a drastic weakening mechanism: the *inhibition* of some formulas from the base. It has been studied in depth, both from the logical side and from the computational one, and its success can be explained by several factors, like the fact that it is quite simple in essence and that it encompasses other important frameworks, including supernormal default theories with priorities [11] and syntax-based belief revision [41], as specific cases.

Nevertheless, the coherence-based approach suffers in its actual turn from several important drawbacks which originate from both its high-level of syntax sensitivity and the simple weakening mechanism it relies on. Especially, it is not adapted to situations where several sources of information must be merged (and not only because two different sources can convey the same information).

In order to illustrate such limitations, let us consider the following scenario. Let B be a base gathering information stemming from two different sources:

- The first source states $\phi_1 = a \wedge b \wedge c$;
- The second source states $\phi_2 = \neg a \wedge \neg b$.

Defining what are the expected consequences of B is not so simple (despite the simplicity of B). It may depend on the reliability of each source, and the relative independence of each piece of information conveyed by the sources. Specifically, it can be the case that $\gamma_1 = \neg a \wedge b \wedge c$ is expected, or that $\gamma_2 = \neg(a \Leftrightarrow b) \wedge c$ is expected. Indeed:

- γ_1 is expected whenever the first source is very reliable to what concerns b , but not so much to what concerns a , and the converse holds for the second source. Furthermore, c is expected because there is no argument against it.
- γ_2 is expected because its models are in-between those of each source.

The point is that none of these two consequences can be derived through B skeptically interpreted under the standard coherence-based approach. Actually, there are only two maximal (w.r.t. \subseteq or cardinality) subsets of B : $\{\phi_1\}$ and $\{\phi_2\}$. If ϕ_1 is strictly preferred to ϕ_2 , then ϕ_1 is the consequence; if ϕ_2 is strictly preferred to ϕ_1 , then ϕ_2 is the consequence; if both sources are equally preferred, then $\phi_1 \vee \phi_2$ would be the consequence.

The problem stems from the fact that each formula from B is taken as a whole, so if it is preferred, it is preferred as a whole and if it must be weakened so as to restore consistency, it is inhibited as a whole. Thus, if ϕ_1 is preferred to ϕ_2 , every piece of information conveyed by ϕ_1 is preferred to the possibly conflicting corresponding piece of information in ϕ_2 while this is not always expected. Furthermore, if ϕ_1 is inhibited, every piece of information conveyed by ϕ_1 is inhibited as well, even those that are not involved in any contradiction.

Of course, the problem can be partially solved by replacing each ϕ_i by a (logically equivalent) subset of its consequences, referred to as a decomposition of ϕ_i . However, this is not a panacea. Indeed, this solution leads to give much importance to the syntax since several decompositions may easily give rise to different consequences. While the choice of a decomposition for ϕ_i can be viewed as an implicit (and compact w.r.t. representation) way to express some preferences over the logical consequences of ϕ_i , the problem of finding out the “right decomposition” of ϕ_i (given as a single formula) remains unsolved. For instance, assume that $\phi_i = a \wedge (\neg a \vee b)$ is known as more reliable to what concerns b than to what concerns a ; if ϕ_i must be weakened because it contradicts the piece of information $\neg a \vee \neg b$ (stemming from a more priority source), b is an expected consequence but it is not a direct subformula of the conjunctive formula ϕ_i , hence some computational effort must be spent to derive it. Here, it would be required to derive explicitly the most general consequence of ϕ_i that is about b so as to give it some priority. The problem is that the computational resource to be spent so as to achieve this goal is high, even if we neglect the time which is consumed (this could be reasonable, provided that the decomposition is performed only once and off-line): indeed, in the general case and under the usual assumptions of complexity theory, there is no formula ϕ_i^X equivalent to $\exists (PS \setminus X). \phi_i$, namely the most general consequence of ϕ_i that is about a given set X of variables and s.t. $|\phi_i^X|$ is polynomial in $|\phi_i|$ (see e.g. [33]). Another dimension that is not handled in a satisfying way by the coherence-based approach concerns equity in the weakening process. Indeed, in some situations, it can be expected that consistency is recovered by weakening each source in the same way. Removing the same number of formulas from each source does not solve this problem at all since what is important is the logical contents of each source, while the number of formulas in a decomposition mainly is irrelevant.

Compared with the weakening by inhibition mechanism at work in the coherence-based approaches, weakening by variable forgetting is typically less drastic; indeed, instead of inhibiting a whole formula – or equivalently, replacing it by \top – it is possible to keep all its consequences that are not involved in any contradiction (see e.g. the piece of information c in Example 6.1). Subsequently, more information can be preserved.

Distance-based merging. As recalled in Section 4.2.1, most approaches to distance-based merging are based on a notion of (pseudo-)distance between worlds (see e.g. [37,47,31,40,29]). Distance-based merging consists in weakening the pieces of information by *dilating* them: the piece of information ϕ , instead of expressing that the real world is for sure among

the models of ϕ , now expresses that it is close to be a model of ϕ , with respect to some given distance d (the further a world ω from the models of ϕ , the less plausible it is that ω is the real world). A merging operator then maps a profile $B = \langle \phi_1, \dots, \phi_n \rangle$ of n consistent propositional formulas into a merged base whose models are the closest to B .

Distance-based merging operators from the literature are typically *decomposable*, that is, the distance between an interpretation ω and a base $B = \langle \phi_1, \dots, \phi_n \rangle$ is computed as $d(\omega, B) = \star(d(\omega, \phi_1), \dots, d(\omega, \phi_n))$, where \star is an aggregation function. Furthermore, a decomposable merging operator is *anonymous* if \star is symmetric (which means that all formulas ϕ_i in B are treated equally).

Even if this framework has a good level of generality, due to the variety of distance functions that can be chosen, almost all papers we know focus on the following two distances: the Hamming distance d_H and the so-called drastic distance d_D . However, even considering wider classes of distances will not really help solving simple examples such as the following one.

Let $B = \langle \phi_1, \phi_2, \phi_3, \phi_4, \phi_5, \phi_6 \rangle$ where $\phi_1 = (a \wedge b \wedge c) \vee (\neg a \wedge \neg b)$, $\phi_2 = \neg c \wedge (a \Leftrightarrow b)$, $\phi_3 = \neg a \vee \neg c$, $\phi_4 = c \wedge (a \vee b)$, $\phi_5 = c$ and $\phi_6 = \neg c$. Assume also that the six sources are equally reliable, which requires \star to be symmetric.

A quick examination on B reveals that c seems to be the source of conflict. Forgetting the information about c in each of the ϕ_i 's would lead to the consistent base $\langle a \Leftrightarrow b, a \Leftrightarrow b, \top, a \vee b, \top, \top \rangle$, whose conjunction is simply equivalent to $a \wedge b$. Note that forgetting any other single variable (a or b) in B will not help solving the inconsistency (even forgetting both a and b will not). Therefore, first ignoring c in the process of merging the sources of information is worth doing: it gives some useful information about everything except c (if information about c is really needed and beliefs (resp. goals) are considered, then one may perform knowledge-gathering actions (resp. start a negotiation process) so as to make a decision about c). This process is even more significant if c is a variable of secondary importance. For instance, in a decision-theoretic context, where the ϕ_i 's represent goals, each being held by an agent, and the interpretations of Ω represent the set of alternatives which can be jointly chosen by the group of agents, an inconsistency simply means that the agents cannot jointly satisfy their goals; the fact that forgetting c is simple way of getting rid of inconsistency means that a reasonable decision can be made about a and b , leaving the decision about c undecided (again, this is even more relevant when c is not of primary importance).

However, we will have a hard time obtaining this result $a \wedge b$ with distance-based merging. For instance, take d to be the Hamming distance d_H . For any interpretation $\omega \in \Omega$, let $\delta_H(\omega, B) = \langle d_H(\omega, \phi_i) \mid i = 1, \dots, 6 \rangle$. We have:

$$\begin{aligned} \delta_H(abc, B) &= \langle 0, 1, 1, 0, 0, 1 \rangle; & \delta_H(ab\bar{c}, B) &= \langle 1, 0, 0, 1, 1, 0 \rangle; \\ \delta_H(a\bar{b}c, B) &= \langle 1, 2, 1, 0, 0, 1 \rangle; & \delta_H(a\bar{b}\bar{c}, B) &= \langle 1, 1, 0, 1, 1, 0 \rangle; \\ \delta_H(\bar{a}bc, B) &= \langle 1, 2, 0, 0, 0, 1 \rangle; & \delta_H(\bar{a}b\bar{c}, B) &= \langle 0, 1, 0, 1, 1, 0 \rangle; \\ \delta_H(\bar{a}\bar{b}c, B) &= \langle 1, 1, 0, 1, 0, 1 \rangle; & \delta_H(\bar{a}\bar{b}\bar{c}, B) &= \langle 0, 0, 0, 2, 1, 0 \rangle. \end{aligned}$$

We see that three interpretations have three 1's and three 0's: abc , $ab\bar{c}$ and $\bar{a}b\bar{c}$. Therefore, no anonymous decomposable merging operator based on d_H can output $a \wedge b$ as the merging of B . In particular, if $\star = +$ then the merging of B is equivalent to $(a \wedge b) \vee (\neg a \wedge \neg c)$; if $\star = \text{leximax}$ then the merging of B is equivalent to $(a \wedge b) \vee (\neg a \wedge b \wedge \neg c)$; and if $\star = \text{max}$ then it is $(a \wedge b) \vee (a \wedge \neg b \wedge \neg c) \vee (\neg a \wedge b \wedge \neg c) \vee (\neg a \wedge \neg b \wedge c)$.

It is important noticing that the failure of distance-based merging to get simple results such as $a \wedge b$ in the above example is not due to the specific choice of the Hamming distance. Similar counterexamples could be obtained with any distance (although it does not seem easy to construct a counterexample that would work for all distances). The reason why we fail to get such simple outputs is the decomposability of the merging operator.

7. Other related work

Our approach consists in exploiting variable forgetting in order to reason in a non-trivial way from inconsistent bases. Both the problem under consideration (reasoning under inconsistency) and the method used to solve it (variable forgetting) have received much attention for years, and are described in depth in several previous papers, especially [5] and [33]. This explains why we refrain from presenting them here again exhaustively and focus instead on recent and/or closely related work about forgetting and/or reasoning under inconsistency.

Conflict resolving by forgetting. The principle we make use of, namely, forgetting a preferred set of atoms so as to make a base consistent, has been applied recently in a few other places.

[50,52,51] and [20,21] use forgetting for resolving conflicts in *logic programs*. These two approaches differ on the nature of the logic programs considered (classical logic programs with negation as failure in [50,52,51]; disjunctive logic programs with strong negation and negation as failure in [20,21]), but the principle at work for solving conflicts is similar: starting from a finite collection of logic programs (one for each of a set of agents), one looks for sets of atoms or literals to forget in each of the logic programs so that the resulting set of logic programs is consistent.

In [19], using mappings from logic programs to ontologies, notions and techniques for forgetting in logic programs are adapted to forgetting concepts in ontologies, which in turn can be a way similar to ours for resolving conflicts when merging

ontologies: for instance, if the conflict between two ontologies is caused by some concept C , C will be forgotten from one of these ontologies or from both.

[22] define the forgetting of actions in a domain description (expressed in a propositional action language), and mention as a potential application the search for joint plans for several agents, each of whom has her own goals about which actions should be in the joint plan. When agents' goals conflict, a minimal set of actions to forget is looked for, and a joint plan (not mentioning these actions) is constructed. Once again, the methodology at work is in the same vein as ours, the difference being in the nature of the base and the objects to be forgotten.

Reasoning under inconsistency. The closest work to our own one is [6] which elaborates on our previous work [35]. In both works, forgetting is used as a key mechanism for reasoning under inconsistency. While the present work includes a number of characterization results and complexity results which were not available in previous papers, it goes also further conceptually speaking by considering a more general notion of forgetting context (as discussed previously).

In [16], an inference relation in the setting of 3-valued paraconsistent logics is defined and studied. Using inconsistency forgetting as a key mechanism for recovering consistency, it guarantees that the deductive closure of any formula is classically consistent and classically closed. This allows to interpret a classically inconsistent formula as a set of classical worlds (hence to reason classically from it). The notion of forgetting considered in this paper is not the one from classical propositional logic, as used here. The approach enables to reason in a non-trivial way from a single inconsistent formula, which is not feasible using forget-based inference. Conversely, since it basically considers a single formula as input, this approach is not suited to merging (for which distinguishing the various sources of information is important). Furthermore, it relies on quite a drastic morphology restriction of the language (only \neg , \vee and \wedge are allowed).

[9] gives a family of paraconsistent inference relations based on so-called signed systems. The formal framework they consider is a propositional framework not containing irreducible contradictions (like \perp or $\neg\top$) and they focus on formulas in Negation Normal Form (NNF).³ Every formula Σ is associated to a default theory $\langle \Sigma^\pm, D_\Sigma \rangle$ where:

- Σ^\pm is a formula in the language $PROP_{PS^\pm}$ where $PS^\pm = \{x^+ \mid x \in PS\} \cup \{x^- \mid x \in PS\}$; Σ^\pm is obtained by replacing in Σ every occurrence of a positive literal x by the positive literal x^+ and every occurrence of a negative literal $\neg x$ by the positive literal x^- .
- $D_\Sigma = \{\delta_x \mid x \in \text{Var}(\Sigma)\}$ is a set of prerequisite-free default rules

$$\delta_x = \left\{ \frac{x^+ \Leftrightarrow \neg x^-}{(x \Leftrightarrow x^+) \wedge (\neg x \Leftrightarrow x^-)} \right\}$$

Negation is given a special treatment in such signed systems; first, every literal is rendered independent from its negation through renaming, this ensures the consistency of Σ^\pm , given the restriction put on the language; then, the corresponding dependence relations are re-introduced in a parsimonious way, so that consistency is preserved. Several inference relations can be defined from $\langle \Sigma^\pm, D_\Sigma \rangle$ (see [9] for details). Among them is *skeptical unsigned inference* (resp. *skeptical signed inference*): a formula ψ of $PROP_{PS}$ is a skeptical unsigned (resp. signed) consequence of Σ , denoted $\Sigma \vdash_u \psi$ (resp. $\Sigma \vdash_s \psi$) if and only if ψ (resp. ψ^\pm) belongs to every extension of $\langle \Sigma^\pm, D_\Sigma \rangle$.

None of these inference relations is a forget-based inference relation. First of all, in signed systems, non-trivial inferences can be drawn starting from any single inconsistent formula Σ . For instance, we have $a \wedge \neg a \wedge b \not\vdash_u c$ and $a \wedge \neg a \wedge b \not\vdash_s c$, while forget-based inference trivializes when starting from $B = \langle a \wedge \neg a \wedge b \rangle$. One may object that this is because variable forgetting cannot be used to restore consistency and that the problem would disappear if the input formula Σ were a conjunction of consistent formulas. However, this is not the case. On the one hand, as far as signed inference is considered, the set of consequences of an inconsistent Σ is not guaranteed to be consistent; for instance, both a and $\neg a$ are skeptical signed consequences of $a \wedge \neg a \wedge b$. Contrastingly, the set of consequences w.r.t. any skeptical forget-based inference relation of any base consisting of consistent formulas is consistent, unless it is the language itself. Hence there is no way to associate every formula Σ like $a \wedge \neg a \wedge b$ to a vector B of consistent formulas such that the skeptical signed consequences of Σ coincide with the consequences of B w.r.t. a skeptical forget-based inference relation. On the other hand, as far as unsigned inference is considered, the choice of B has a great impact on the consequences that can be drawn w.r.t. forget-based inferences. If $a \wedge \neg a \wedge b$ is associated to $B = \langle a, \neg a \wedge b \rangle$, then the skeptical unsigned consequences of Σ (i.e., the classical consequences of b) can be recovered as the skeptical consequences of B w.r.t. the prototypical homogeneous forget-based inference relation given \mathcal{C}_S . If $a \wedge \neg a \wedge b$ is associated to $B = \langle a, \neg a, a \vee b \rangle$, b can no longer be derived using the prototypical homogeneous forget-based inference relation given \mathcal{C}_S .

Similar conclusions can be drawn for settings extending signed systems, especially those based on signed quantified Boolean formulas [1].

³ This is without loss of expressiveness since every propositional formula can be turned into an equivalent NNF, but not without loss of succinctness when connectives like \Leftrightarrow belongs to the morphology of the language.

8. Conclusion

We have proposed a new framework for reasoning from inconsistent propositional bases, which proceeds by selecting subsets of atoms to be forgotten in the pieces of information, respecting some fixed constraints. We have shown that this framework encompasses many existing frameworks as specific cases, including reasoning from preferred consistent subsets, and several forms of merging, and that the gain obtained in flexibility is not balanced by a complexity shift w.r.t. inference.

As evoked in the introduction, our approach is based on classical entailment so that the logical contents of the base are weakened instead of the inference relation itself. This gives to our family of forget-based relations some pros and some cons. Thus, if the base under consideration consists of a single inconsistent formula, our approach is not suited (as well as the coherence-based approaches and the merging ones, which are also based on classical entailment). Paraconsistent logics based on subclassical entailment relations can then be used with profit. Conversely, since weakening bears on the inference relation in paraconsistent logics, it is uniform and the fact that information comes from several sources typically⁴ is not exploited. For instance, it is not possible to derive b from the (conjunctively-interpreted) set of formulas $\{a \wedge (\neg a \vee b), \neg a\}$ using the paraconsistent logic LP_m [45] (this is because $\neg a \vee b$ is a consequence of $\neg a$), while we can easily get this consequence in our setting.

Further work can be conducted along several directions.

First, inconsistency handling by variable forgetting can be easily generalized to other settings than propositional logic. All we need is a language involving several variables equipped with a forgetting operation. We have already mentioned in Section 7 how conflict solving via forgetting can be done in various forms of logic programs as well as in propositional action descriptions. There is no reason why we should stop here. For instance, variable forgetting has been defined in multi-agent settings [49] as well as in the propositional modal logic of knowledge $S5$ [53], which pave the way towards conflict solving via forgetting in these settings. More generally, we could generalize the principle to many settings where variable forgetting (or equivalently, projection or marginalization) makes sense; this is the case for other propositional settings (e.g. when modalities or multiple-valued semantics are considered) but also for many non-propositional settings: for instance, the pieces of information given by the sources could be constraints, joint probability distributions, etc.

Finally, weakening formulas by variable forgetting can be useful in many practical situations. For instance, if a series of tests has to be performed so as to “solve” inconsistency [34,36], preferred recoveries give some hints about the variables that should be tested first, namely, the most important ones that have to be forgotten. Similarly, in a decision-theoretic context where the ϕ_i 's are no longer *beliefs* but *preferences* of several agents, variables which do not need to be forgotten are these the agents agree on, and preferred recoveries help finding the variables about which negotiation should start.

Acknowledgements

The authors thank Philippe Besnard and Torsten Schaub for many helpful discussions about the topic of this paper. They are also grateful to the anonymous reviewers for a number of interesting comments and suggestions.

Appendix A

Proof of Proposition 3.1. Ω_B is empty if and only if B is not recoverable. When Ω_B is empty, every ψ is such that $\forall \omega \in \min(\Omega_B, \leq_{\square})$, $\omega \models \psi$. Similarly, when B is not recoverable, we have $B \models_{\square}^{\subseteq} \psi$ for every ψ . We assume now that $\Omega_B \neq \emptyset$.

- Let us first consider $\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \square)$. Let ω be any model of $B \mid \vec{V}$, and $\vec{V}_{\omega} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B)(\omega), \square)$. Let us show that ω is minimal w.r.t. \leq_{\square} in Ω_B . First, we obviously have that $\omega \in \Omega_B$. Now, towards a contradiction, assume that there exists an $\omega' \in \Omega_B$ such that $\omega' <_{\square} \omega$. Let $\vec{V}_{\omega'} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B)(\omega'), \square)$. Since $\omega' <_{\square} \omega$, we have that $\vec{V}_{\omega'} \sqsubset \vec{V}_{\omega}$; and because $\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \square)$, we have $\vec{V} \sim \vec{V}_{\omega}$, henceforth $\vec{V}_{\omega'} \sqsubset \vec{V}$, which contradicts the fact that $\vec{V} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \square)$.
- Conversely, let $\omega \in \min(\Omega_B, \leq_{\square})$, and $\vec{V}_{\omega} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B)(\omega), \square)$. Let us show that $\vec{V}_{\omega} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B), \square)$. Towards a contradiction, assume not. Then there must exist an $\omega' \in \Omega_B$ and a $\vec{V}_{\omega'} \in \text{Pref}(\mathcal{R}_{\mathcal{C}}(B)(\omega'), \square)$ such that $\vec{V}_{\omega'} \sqsubset \vec{V}_{\omega}$. Therefore, we have $\omega' <_{\square} \omega$, which contradicts $\omega \in \min(\Omega_B, \leq_{\square})$. \square

Proof of Proposition 4.1. We say that a forgetting vector $\vec{V} = \langle V_1, \dots, V_n \rangle$ is *binary* if for every i , either $V_i = PS$ or $V_i = \emptyset$.

Let us first consider the following notations:

- Let $B = \langle \phi_1, \dots, \phi_n \rangle$. For any subset S of $\{1, \dots, n\}$, we define $f(S)$ as the binary vector $\vec{V}_S = \langle V_1, \dots, V_n \rangle$ where for each $i \in \{1, \dots, n\}$, $V_i = \emptyset$ if $\phi_i \in S$ and $V_i = \text{Var}(\phi_i)$ otherwise.
- For any forgetting vector $\vec{V} = \langle V_1, \dots, V_n \rangle$, we define $g(\vec{V}) = S_{\vec{V}}$ as the subset of $\{1, \dots, n\}$ defined by $S_{\vec{V}} = \{i \mid V_i = \emptyset\}$, and the *binary forgetting vector associated with \vec{V}* is $\vec{V}^* = f(g(\vec{V}))$, i.e., $V_i^* = \emptyset$ if $V_i = \emptyset$ and $V_i^* = \text{Var}(\phi_i)$ otherwise.

⁴ There are nevertheless some exceptions, like [8].

Clearly, f is a one-to-one mapping between binary forgetting vectors and subsets of $\{1, \dots, n\}$, and if \vec{V} is binary then $f(g(\vec{V})) = \vec{V}$.

Now, for any $S \subseteq B$,

$$(B \mid \vec{V}_S) = \{\phi_i \mid i \in S\} = B_S. \quad (1)$$

Therefore,

$$S \in \text{Cons}(B) \text{ if and only if } \vec{V}_S \in \mathcal{R}_{\mathcal{C}_S}(B) \quad (2)$$

and the restriction of f to $\text{Cons}(B)$ is a one-to-one mapping between $\text{Cons}(B)$ and $\mathcal{R}_{\mathcal{C}_S}(B)$.

Next, we define the following one-to-one mapping between preference relations on $\text{Cons}(B)$ and preference relations on $\mathcal{R}_{\mathcal{C}_S}(B)$ satisfying binaricity. Let \sqsubseteq be a preference relation over $\mathcal{R}_{\mathcal{C}_S}(B)$ satisfying binaricity. Let $\preceq = h(\sqsubseteq)$ be the relation on $\text{Cons}(B)$ defined by: for any $X, Y \in \text{Cons}(B)$, $X \preceq Y$ if and only if $\vec{V}_X \sqsubseteq \vec{V}_Y$. Clearly, \preceq is a preorder because \sqsubseteq is a preorder, henceforth \preceq is a preference relation over consistent subsets of $\{1, \dots, n\}$.

Now, we show

$$\text{for any } \vec{V}, \vec{W} \in \mathcal{R}_{\mathcal{C}_S}(B) \text{ we have } \vec{V} \sqsubseteq \vec{W} \text{ if and only if } S_{\vec{V}} \preceq S_{\vec{W}}. \quad (3)$$

First, for any $\vec{V} \in \mathcal{R}_{\mathcal{C}_S}(B)$, $\vec{V} \in \mathcal{R}_{\mathcal{C}_S}(B)$ implies $\vec{V}^* \in \mathcal{R}_{\mathcal{C}_S}(B)$. Indeed, $\vec{V} \sqsubseteq_p \vec{V}^*$, henceforth $(B \mid \vec{V}) \models (B \mid \vec{V}^*)$, which shows that the consistency of $(B \mid \vec{V})$ implies the consistency of $(B \mid \vec{V}^*)$. Now, \preceq satisfies binaricity, therefore $\vec{V} \sim \vec{V}^*$ and $\vec{W} \sim \vec{W}^*$, from which we get $\vec{V} \sqsubseteq \vec{W}$ if and only if $\vec{V}^* \sqsubseteq \vec{W}^*$, which by definition of h is equivalent to $S_{\vec{V}^*} \preceq S_{\vec{W}^*}$, which is now equivalent to $S_{\vec{V}} \preceq S_{\vec{W}}$, because $S_{\vec{V}} = S_{\vec{V}^*}$ and $S_{\vec{W}} = S_{\vec{W}^*}$.

This implies that h is a bijection between the set of preference relations on $\mathcal{R}_{\mathcal{C}_S}(B)$ satisfying binaricity and the set of preference relations on $\text{Cons}(B)$. For any preference relation \preceq on $\text{Cons}(B)$, let $k(\preceq) = \sqsubseteq$ be the relation on $\mathcal{R}_{\mathcal{C}_S}(B)$ defined by $\vec{X} \sqsubseteq \vec{Y}$ if and only if $S_{\vec{X}} \preceq S_{\vec{Y}}$. Then $k(\preceq)$ is a preference relation on $\mathcal{R}_{\mathcal{C}_S}(B)$ satisfying binaricity, and due to (3) we have $k = h^{-1}$.

Now we show that

$$S \in \text{Pref}(\text{Cons}(B), \preceq) \text{ if and only if } \vec{V}_S \in \text{Pref}(\mathcal{R}_{\mathcal{C}_S}(B), \sqsubseteq). \quad (4)$$

For (\Rightarrow) : assume $\vec{V}_S \notin \text{Pref}(\mathcal{R}_{\mathcal{C}_S}(B), \sqsubseteq)$. This means that either $\vec{V}_S \notin \mathcal{R}_{\mathcal{C}_S}(B)$, in which case $S \notin \text{Cons}(B)$, or that there exists $W \in \mathcal{R}_{\mathcal{C}_S}(B)$ such that $W \sqsubset \vec{V}_S$; in the latter case, $S_W \prec S$, and $S_W \in \text{Cons}(B)$ because $W \in \mathcal{R}_{\mathcal{C}_S}(B)$, therefore $S \notin \text{Pref}(\text{Cons}(B), \preceq)$.

For (\Leftarrow) : assume S is not a \preceq -preferred consistent subset of $\{1, \dots, n\}$. This means that either S is not consistent, in which case $\vec{V}_S \notin \mathcal{R}_{\mathcal{C}_S}(B)$, or there exists $S' \in \text{Cons}(B)$ such that $S' \prec S$, in which case $\vec{V}_{S'} \prec \vec{V}_S$ and \vec{V}_S is not a preferred recovery for B .

We are now ready for proving Proposition 4.1.

(\Rightarrow) Let \approx be a maxcons inference relation, which by definition means that there exists a preorder \preceq on subsets of $\{1, \dots, n\}$ such that $\approx = \approx_{\preceq}^{\vee}$, i.e., such that $B \approx \psi$ if and only if for any $X \in \text{Pref}(\text{Cons}(B), \preceq)$ we have $B_X \models \psi$. Let $\sqsubseteq = h^{-1}(\preceq)$. \sqsubseteq satisfies binaricity. From (4) we have $X \in \text{Pref}(\text{Cons}(B), \preceq)$ if and only if $\vec{V}_X \in \text{Pref}(\mathcal{R}_{\mathcal{C}_S}(B), \sqsubseteq)$. From (1) we have $(B \mid \vec{V}_X) = B_X$, therefore $(B \mid \vec{V}_X) \models \psi$ if and only if $B_X \models \psi$. Therefore, $B \approx \psi$ if and only if $B \approx_{\sqsubseteq}^{\vee} \psi$.

(\Leftarrow) Let \sqsubseteq be a preference relation on $\mathcal{R}_{\mathcal{C}_S}(B)$ satisfying binaricity. Let $\preceq = h(\sqsubseteq)$. Then \preceq is a preference relation on $2^{\{1, \dots, n\}}$, and again due to (1) and (4) we have $B \approx_{\sqsubseteq}^{\vee} \psi$ if and only if $B \approx_{\preceq}^{\vee} \psi$, i.e., $\approx_{\sqsubseteq}^{\vee}$ is a maxcons inference relation. \square

Proof of Proposition 4.2. The proof is very similar to that of Proposition 4.1. For every $S \subseteq \{1, \dots, n\}$, define the forgetting vector \vec{V}_S by $\vec{V}_S = (V_1, \dots, V_n)$, where $V_i = \emptyset$ if $\phi_i \in S$ and $V_i = \text{Var}(\phi_i)$ otherwise. Note that \vec{V}_S satisfies the logically weakest binary context. Conversely, for every forgetting vector $\vec{V} = (V_1, \dots, V_n)$ satisfying the logically weakest binary context, let $S_{\vec{V}} = \{i \mid V_i = \emptyset\}$. Note that $(B \mid \vec{V}_S) = B_S$ and $(B \mid \vec{V}) = B_{S_{\vec{V}}}$.

Let \approx be a maxcons inference relation, associated with the preference relation \preceq , that is, $\approx = \approx_{\preceq}^{\vee}$. Define the preference relation on forgetting vectors by $\vec{V} \sqsubseteq \vec{V}'$ if and only if $S_{\vec{V}} \preceq S_{\vec{V}'}$. Then we have $\vec{V} \sqsubseteq \vec{V}'$ if and only if $S_{\vec{V}} \preceq S_{\vec{V}'}$. With similar arguments as those used in the proof of Proposition 4.1, we get that \approx_{\preceq}^{\vee} and $\approx_{\sqsubseteq}^{\vee}$ coincide.

Conversely, let $\approx_{\sqsubseteq}^{\vee}$ be a forget-based inference relation associated with a preference relation \sqsubseteq and the logically weakest binary context. Define the preference relation \preceq over subsets of $\{1, \dots, n\}$ by $S \preceq S'$ if and only if $\vec{V}_S \sqsubseteq \vec{V}_{S'}$. Again, with similar arguments as those used in the proof of Proposition 4.1, we get that \approx_{\preceq}^{\vee} and $\approx_{\sqsubseteq}^{\vee}$ coincide. \square

Proof of Proposition 4.3. Let us first notice that any decomposable ranking function μ can be equivalently represented by a pair $\langle \mu_0, H \rangle$. Now, there exists a straightforward one-to-one mapping between decomposable ranking functions $\langle \mu_0, H \rangle$

and pairs $\langle f, \star \rangle$ defining a merging operator based on a differential distance. Let us call t this mapping, defined by $t(\mu) = t(\langle \mu_0, H \rangle) = \langle f, \star \rangle$ with $f = \mu_0$ and $\star = H$.

Now, let μ be a decomposable ranking function on subsets of PS and d be its corresponding differential distance defined by $d(\omega, \omega') = \mu_0(\text{Diff}(\omega, \omega'))$. Let $\mu_B(\omega) = \min(\{\mu(\vec{V}) \mid \vec{V} \in \mathcal{R}_{\mathcal{C}_S}(B) \text{ and } \omega \models B \mid \vec{V}\})$. We first establish that $d(\omega, B) = \mu_B(\omega)$ holds for any $\omega \in \Omega$.

Let $\omega \in \Omega$ and $B = \langle \phi_1, \dots, \phi_n \rangle$. Since \star is monotonic, we have:

$$\begin{aligned} d(\omega, B) &= \min(\{\star(d(\omega, \phi_1), \dots, d(\omega, \phi_n))\}) \\ &= \min(\{\star(d(\omega, \omega_1), \dots, d(\omega, \omega_n)) \mid \omega_1 \models \phi_1, \dots, \omega_n \models \phi_n\}) \\ &= \min(\{H(\mu_0(\text{Diff}(\omega, \omega_1)), \dots, \mu_0(\text{Diff}(\omega, \omega_n))) \mid \omega_1 \models \phi_1, \dots, \omega_n \models \phi_n\}). \end{aligned}$$

Let $\vec{V} = \langle \text{Diff}(\omega, \omega_1), \dots, \text{Diff}(\omega, \omega_n) \rangle$. Using Corollary 5 of [33], for every $i \in \{1, \dots, n\}$ we have $\omega \models \exists \text{Diff}(\omega, \omega_i). \phi_i$ because $\omega_i \models \phi_i$. Therefore, $\omega \models (B \mid \vec{V})$, hence, $\vec{V} \in \mathcal{R}_{\mathcal{C}_S}(B)$. Now, $H(\mu_0(\text{Diff}(\omega, \omega_1)), \dots, \mu_0(\text{Diff}(\omega, \omega_n))) = \mu(\vec{V})$ by definition of μ . Therefore, $d(\omega, B) = \min(\{\mu(\vec{V}) \mid \vec{V} \in \mathcal{R}_{\mathcal{C}_S}(B) \text{ and } \omega \models B \mid \vec{V}\}) = \mu_B(\omega)$.

Now, since the standard forgetting context \mathcal{C} is considered, $\langle PS, \dots, PS \rangle$ is a forgetting vector for B . Since each ϕ_i ($i \in \{1, \dots, n\}$) is consistent, every interpretation $\omega \in \Omega$ is a model of $B \mid \langle PS, \dots, PS \rangle$, showing that $\Omega_B = \Omega$. Furthermore, given $\omega, \omega' \in \Omega$, $\mu_B(\omega) \leq \mu_B(\omega')$ if and only if each \vec{V} in $\min(\{\mu(\vec{V}) \mid \vec{V} \in \mathcal{R}_{\mathcal{C}_S}(B) \text{ and } \omega \models B \mid \vec{V}\})$ and each \vec{V}' in $\min(\{\mu(\vec{V}') \mid \vec{V}' \in \mathcal{R}_{\mathcal{C}_S}(B) \text{ and } \omega' \models B \mid \vec{V}'\})$ are such that $\vec{V} \sqsubseteq_{\mu} \vec{V}'$ if and only if $\omega \sqsubseteq_{\mu} \omega'$.

Lastly, $B \approx_{d, \star}^M \psi$ if and only if for every $\omega \in \Omega$ such that $d(\omega, B) = \min_{\omega' \in \Omega} d(\omega', B)$ we have $\omega \models \psi$, which, using a previous equality $d(\omega, B) = \mu_B(\omega)$, is equivalent to: for every $\omega \in \Omega$ such that $\mu_B(\omega) = \min_{\omega' \in \Omega} \mu_B(\omega')$ we have $\omega \models \psi$, which is also equivalent to: for every $\omega \in \min(\Omega_B, \sqsubseteq_{\mu})$ we have $\omega \models \psi$.

Since \sqsubseteq_{μ} is complete (by construction), Proposition 3.1 shows that this is equivalent to $B \approx_{\sqsubseteq_{\mu}}^{\mathcal{C}_S} \psi$. \square

Proof of Proposition 4.4. We start by introducing the following notations: for any $X \subseteq EQ$, let $F(X) = \langle V_1, \dots, V_n \rangle$ where each V_i is defined by $V_i = \{x \mid x_0 \Leftrightarrow x_i \notin X\}$. Clearly, F is bijective and for any vector $\vec{V} = \langle V_1, \dots, V_n \rangle$ of subsets of PS we have $F^{-1}(\vec{V}) = \{x_i \Leftrightarrow x_0 \mid x \in PS, 1 \leq i \leq n, x \notin V_i\}$.

We associate with PS the sets PS_0, \dots, PS_n s.t. for each $i \in \{0, \dots, n\}$, $PS_i = \{x_i \mid x \in PS\}$. Every interpretation M over $\bigcup_{i=0}^n PS_i$ can be represented by a tuple $(\omega_0, \dots, \omega_n)$ s.t. for each $i \in \{0, \dots, n\}$, ω_i is the restriction of M to PS_i . Finally, given an interpretation ω_i over PS_i , we can associate to it in a bijective way the interpretation ω_i^* over PS s.t. $\forall x \in PS$, $\omega_i^*(x) = \omega_i(x_i)$. Clearly enough, for each $i \in \{1, \dots, n\}$ and each formula $\phi \in \text{PROP}_{PS}$, $\omega_i^* \models \phi$ if and only if $\omega_i \models \text{rename}(\phi, i)$.

The proof now goes by successively proving the following lemmata.

Lemma 8.1. $X \subseteq EQ$ is a fit for B if and only if $F(X)$ is a recovery for B given \mathcal{C}_S .

Lemma 8.2. $X \subseteq EQ$ is a maximal fit for B if and only if $F(X)$ is a preferred recovery w.r.t. \sqsubseteq_p for B given \mathcal{C}_S .

Lemma 8.3.

$$\text{Mod}(B \mid F(X)) = \left\{ \omega_0^* \in 2^{PS} \mid M \models B^r \wedge \bigwedge_{x \in X} x \text{ for some } M = (\omega_0, \omega_1, \dots, \omega_n) \right\}.$$

Lemma 8.4. $B^r \wedge \bigwedge_{x \in X} x \models \text{rename}(\psi, 0)$ if and only if $(B \mid F(X)) \models \psi$.

Proof of Lemma 8.1. Let $X \subseteq EQ$ be a fit for B , which means that $B^r \wedge \bigwedge_{x \in X} x$ is consistent, therefore there exists a model $M = (\omega_0, \omega_1, \dots, \omega_n)$ of $B^r \wedge \bigwedge_{x \in X} x$ over $\bigcup_{i=0}^n PS_i$. By definition of B^r we have

$$\text{for each } i \in \{1, \dots, n\}, \quad \omega_i \models \text{rename}(\phi_i, i) \tag{1}$$

and $M \models \bigwedge_{x \in X} x$ is equivalent to

$$\text{for all } x_i \Leftrightarrow x_0 \text{ in } X, \quad \omega_0(x_0) = \omega_i(x_i).$$

This shows that for each $i \in \{1, \dots, n\}$, ω_0^* coincides with ω_i^* (a model of ϕ_i) on every variable from PS except possibly on those x s.t. $x_i \Leftrightarrow x_0 \notin X$, namely those $x \in V_i$. Therefore, ω_0^* is a model of $\exists V_i. \phi_i$ for each $i \in \{1, \dots, n\}$. Since every forgetting vector is admissible for the standard forgetting context, this shows that $\vec{V} = \langle V_1, \dots, V_n \rangle$ is a recovery for B given \mathcal{C}_S .

Conversely, let $V = \langle V_1, \dots, V_n \rangle$ be a recovery for B given \mathcal{C}_S , then $\bigwedge_{i=1}^n \exists V_i. \phi_i$ is consistent. Hence it has a model ω and there exist n models $\omega_1^*, \dots, \omega_n^*$ over PS such that (a) $\omega_i^* \models \phi_i$ for each $i \in \{1, \dots, n\}$ and (b) ω and ω_i^* coincide on $PS \setminus V_i$ for every $i \in \{1, \dots, n\}$. Let ω_0 be the interpretation over PS_0 s.t. $\forall x_0 \in PS_0$ we have $\omega_0(x_0) = \omega(x)$. Let $M = (\omega_0, \omega_1, \dots, \omega_n)$. Condition (a) implies $M \models B^r$ and condition (b) implies $M \models \bigwedge_{x \in X} x$, therefore $M \models B^r \wedge \bigwedge_{x \in X} x$, which shows that X is a fit for B . \square

Proof of Lemma 8.2. Let X be a fit for B . If X is not maximal then there exists a fit X' for B such that $X' \supset X$. Then, by Lemma 8.1, $F(X')$ is a recovery for B given \mathcal{C}_S and obviously, $F(X') \sqsubset_p F(X)$, therefore $F(X)$ cannot be a maximal fit for B .

Conversely, let \vec{V} be a recovery for B given \mathcal{C}_S . If \vec{V} is not preferred w.r.t. $\sqsubseteq = \sqsubseteq_p$ then there exists a recovery \vec{V}' for B given \mathcal{C}_S such that $\vec{V}' \sqsubset_p \vec{V}$, which entails that the associated fit $F^{-1}(\vec{V}')$ for B (thanks to Lemma 8.1) is such that $F^{-1}(\vec{V}') \supset F^{-1}(\vec{V})$, therefore $F^{-1}(\vec{V})$ is not a maximal fit for B . \square

Proof of Lemma 8.3. Let $M = (\omega_0, \omega_1, \dots, \omega_n)$ be a model of $B^r \wedge \bigwedge_{x \in X} x$ over $\bigcup_{i=0}^n PS_i$. Using the same proof as in Lemma 8.1, we derive that ω_0^* is a model of $\exists V_i. \phi_i$ for each $i \in \{1, \dots, n\}$. Hence, ω_0^* is a model of $B \mid \vec{V}$ with $\vec{V} = F(X)$.

Conversely, let ω_0^* be any model of $B \mid \vec{V}$ where \vec{V} is a recovery for B given \mathcal{C}_S . Then for all $i \in \{1, \dots, n\}$, $\omega_0^* \models \exists V_i. \phi_i$, which by Corollary 5 of [33], entails that there exists an interpretation ω_i^* over PS such that $\omega_i^* \models \phi_i$ and ω_0^* and ω_i^* coincide on all variables from PS except possibly on those of V_i . Then let $M = (\omega_0, \omega_1, \dots, \omega_n)$. The fact that $\omega_i^* \models \phi_i$ for all $i \in \{1, \dots, n\}$ entails that $\omega_i \models \text{rename}(\phi_i, i)$ for all $i \in \{1, \dots, n\}$, hence $M \models B^r$. Finally, the fact that for all $i \in \{1, \dots, n\}$, ω_0^* and ω_i^* coincide on all variables from PS except possibly on those of V_i shows that $M \models \bigwedge_{x \in X} x$, and this concludes the proof. \square

Proof of Lemma 8.4. Assume $B^r \wedge \bigwedge_{x \in X} x \models \text{rename}(\psi, 0)$. Let $\vec{V} = F(X)$. Then for all $M = (\omega_0, \omega_1, \dots, \omega_n) \models B^r \wedge \bigwedge_{x \in X} x$ we have $M \models \text{rename}(\psi, 0)$, which is equivalent to $\omega_0 \models \text{rename}(\psi, 0)$ and also equivalent to $\omega_0^* \models \psi$.

By Lemma 8.3, every model ω_0^* of $B \mid \vec{V}$ over PS is such that there exists a model $M = (\omega_0, \omega_1, \dots, \omega_n)$ of $B^r \wedge \bigwedge_{x \in X} x$ over $\bigcup_{i=0}^n PS_i$. Hence we must have $\omega_0^* \models \psi$, and since this must hold for every model ω_0^* of $B \mid \vec{V}$, we conclude that $B \mid \vec{V} \models \psi$.

Conversely, assume that $B \mid \vec{V} \models \psi$. Towards a contradiction, assume that $B^r \wedge \bigwedge_{x \in X} x \not\models \text{rename}(\psi, 0)$. Then there exists an interpretation $M = (\omega_0, \omega_1, \dots, \omega_n)$ which satisfies $B^r \wedge \bigwedge_{x \in X} x$ but does not satisfy $\text{rename}(\psi, 0)$. Hence $\omega_0^* \not\models \psi$. However, Lemma 8.3 shows that ω_0^* must be a model of $B \mid \vec{V}$. This contradicts the fact that $B \mid \vec{V} \models \psi$. \square

Proof of Proposition 4.4. By definition, $B \approx_{DS} \psi$ is equivalent to $B^r \wedge \bigwedge_{x \in X} x \models \text{rename}(\psi, 0)$ for any maximal fit X for B , which, using both Lemmata 8.2 and 8.4, is equivalent to $B \mid \vec{V} \models \psi$ for every recovery \vec{V} for B given \mathcal{C} , that is preferred w.r.t. \sqsubseteq_p . \square

Proof of Proposition 4.5. We take advantage of Proposition 3.1. We obviously have $\Omega_B = \Omega$ since by assumption B consists of consistent formulas and the forgetting context is the standard one (indeed $\langle PS, \dots, PS \rangle$ is a forgetting vector for B given \mathcal{C}_S). Hence it is enough to show that for any interpretations $\omega, \omega' \in \Omega$, we have $\omega \sqsubseteq_B \omega'$ if and only if $\omega \sqsubseteq \omega'$. First, observe that since \sqsubseteq includes \sqsubseteq_p , for any interpretation $\omega \in \Omega$, we have $\text{Pref}(\mathcal{R}_C(B)(\omega), \sqsubseteq) \subseteq \text{diff}(\omega, B)$, and for any interpretations $\omega, \omega' \in \Omega$, for every \vec{V} (resp. \vec{V}') $\in \text{diff}(\omega, B)$ (resp. $\in \text{diff}(\omega', B)$), there exists \vec{W} (resp. \vec{W}') $\in \text{Pref}(\mathcal{R}_C(B)(\omega), \sqsubseteq)$ (resp. $\in \text{Pref}(\mathcal{R}_C(B)(\omega'), \sqsubseteq)$) such that $\vec{W} \sqsubseteq \vec{V}$ (resp. $\vec{W}' \sqsubseteq \vec{V}'$). Now:

- If $\omega \sqsubseteq_B \omega'$ then by definition $\exists \vec{V}_\omega \in \text{diff}(\omega, B)$ s.t. $\forall \vec{V}_{\omega'} \in \text{diff}(\omega', B)$, we have $\vec{V}_\omega \sqsubseteq \vec{V}_{\omega'}$. Hence there exists $\vec{W}_\omega \in \text{Pref}(\mathcal{R}_C(B)(\omega), \sqsubseteq)$ s.t. $\vec{W}_\omega \sqsubseteq \vec{V}_\omega$. Furthermore, $\text{diff}(\omega', B)$ is not empty since every ω' is a model of $(B \mid \langle PS, \dots, PS \rangle)$. Hence for each $\vec{V}_{\omega'} \in \text{diff}(\omega', B)$, we can find a $\vec{W}_{\omega'} \in \text{Pref}(\mathcal{R}_C(B)(\omega'), \sqsubseteq)$ s.t. $\vec{W}_{\omega'} \sqsubseteq \vec{V}_{\omega'}$. Since $\vec{W}_{\omega'}$ belongs to $\text{diff}(\omega', B)$, we must have $\vec{V}_\omega \sqsubseteq \vec{W}_{\omega'}$. Since \sqsubseteq is a preorder, from $\vec{W}_\omega \sqsubseteq \vec{V}_\omega$ and $\vec{V}_\omega \sqsubseteq \vec{W}_{\omega'}$, we get that $\vec{W}_\omega \sqsubseteq \vec{W}_{\omega'}$. This shows that $\omega \sqsubseteq \omega'$.
- Conversely, if $\omega \sqsubseteq \omega'$, then by definition $\exists \vec{V}_\omega \in \text{Pref}(\mathcal{R}_C(B)(\omega), \sqsubseteq)$ $\exists \vec{V}_{\omega'} \in \text{Pref}(\mathcal{R}_C(B)(\omega'), \sqsubseteq)$ s.t. $\vec{V}_\omega \sqsubseteq \vec{V}_{\omega'}$. We have $\vec{V}_\omega \in \text{diff}(\omega, B)$. Since \sqsubseteq is a complete preorder, $\vec{V}_{\omega'} \in \text{Pref}(\mathcal{R}_C(B)(\omega'), \sqsubseteq)$ is such that $\vec{V}_{\omega'} \sqsubseteq \vec{V}'_{\omega'}$ for every $\vec{V}'_{\omega'} \in \text{diff}(\omega', B)$. Hence, by transitivity of \sqsubseteq , $\vec{V}_\omega \in \text{diff}(\omega, B)$ is such that $\vec{V}_\omega \sqsubseteq \vec{V}'_{\omega'}$ for every $\vec{V}'_{\omega'} \in \text{diff}(\omega', B)$. This shows that $\omega \sqsubseteq_B \omega'$. \square

Proof of Proposition 5.1.

- Point 1.
 - Membership. Let $B = \langle \phi_1, \dots, \phi_n \rangle$ be a base and let \mathcal{C} be a forgetting context for B . The following non-deterministic algorithm can be used to determine whether B is recoverable:
 1. Guess V_1, \dots, V_n n subsets of $\text{Var}(B)$;
 2. Guess $\omega, \omega_1, \dots, \omega_n$ $n+1$ interpretations over $\text{Var}(B)$;
 3. For each $i \in \{1, \dots, n\}$ do
 - Check that $F_{\vec{V}} \models \mathcal{C}$ where $\vec{V} = \langle V_1, \dots, V_n \rangle$;
 - Check that ω_i is a model of ϕ_i ;
 - Check that ω coincides with ω_i on $\text{Var}(\phi_i) \setminus V_i$.

The last point in the algorithm above enables to check that $\omega \models \exists V_i. \phi_i$ holds for each $i \in \{1, \dots, n\}$. This implies that $B \mid \vec{V}$ is consistent, hence B is recoverable.

- **Hardness.** Easy, by reduction from SAT. To every CNF formula Σ consisting of n clauses $\gamma_1, \dots, \gamma_n$, we associate in polynomial time the base $B = \langle \gamma_1, \dots, \gamma_n \rangle$ and the forgetting context $C = \bigwedge_{x \in PS} \bigwedge_{i=1}^n \neg \text{forget}(x, i)$. Clearly enough, only one recovery is possible given C : $\langle \emptyset, \dots, \emptyset \rangle$. This vector actually is a recovery for B if and only if Σ is satisfiable.
- **Point 2.** Let us consider the complementary problem and show that it belongs to Σ_2^P . The following non-deterministic algorithm can be used to determine whether $B \not\approx_{\subseteq}^C \psi$:
 1. Guess V_1, \dots, V_n n subsets of $\text{Var}(B)$
 s.t. $\vec{V} = \langle V_1, \dots, V_n \rangle$ is a recovery
 for B given C and $B \mid \vec{V} \not\approx \psi$;
 2. Check that \vec{V} is a preferred recovery
 for B w.r.t. \subseteq .

The proof of Point 1 above shows that the first step of this algorithm can be done in non-deterministic polynomial time (just add a final step “check that $\omega \not\approx \psi$ holds” to the non-deterministic algorithm given in the proof of Point 1). This is also the case for the second step. Indeed, once \vec{V} has been guessed, in order to check that it is not a preferred recovery for B , it is sufficient to guess a recovery \vec{V}' for B (again, just one call to an NP oracle) and to check in deterministic polynomial time that $\vec{V}' \sqsubset \vec{V}$.

- **Point 3.**
 - (i) First of all, for any possible value $k \in \mathbb{N}$ which can be taken by $\mu(\vec{V})$, it is possible to guess V_1, \dots, V_n n subsets of $\text{Var}(B)$ s.t. $\vec{V} = \langle V_1, \dots, V_n \rangle$ is a recovery for B given C (see the proof of Point 1 above), and s.t. $\mu(\vec{V}) = k$ (just add a final step “check that $\mu(\vec{V}) = k$ holds” to the non-deterministic algorithm in the proof of Point 1). Since $\mu(\vec{V})$ can be computed in deterministic polynomial time, we obtain a non-deterministic algorithm to check in time polynomial in the cumulated size of B, C and k whether there exists a recovery \vec{V} for B s.t. $\mu(\vec{V}) = k$.
 - (ii) Now, if μ can be computed in time polynomial in the size of the input then the size of the binary representation of $\mu(\vec{V})$ for any $\vec{V} \in \mathcal{R}_C(B)$ is bounded by a polynomial in the input size. Accordingly, the value of $\mu(\vec{V})$ is bounded by a (simple) exponential in the input size. Through binary search, a polynomial number of calls to the non-deterministic algorithm described above is sufficient to find out the minimal value $\min \in \mathbb{N}$ s.t. there exists a recovery \vec{V} for B given C and $\mu(\vec{V}) = \min$.
 - (iii) Once \min has been computed, a last call to an NP oracle is sufficient to guess a preferred recovery \vec{V} for B s.t. $B \mid \vec{V} \not\approx \psi$ (just add a final step “check that $\omega \not\approx \psi$ holds” to the non-deterministic algorithm given in the proof of Point 1).
 - (iv) The fact that Δ_2^P is closed for the complement concludes the proof. \square

Proof of Proposition 5.2. Membership comes directly from item 2 of Proposition 5.1. Hardness is obtained by a reduction from skeptical inference from supernormal default theories without background theory. The latter problem consists in determining, given a set $\Delta = \{\phi_1, \dots, \phi_n\}$ of propositional formulas and a propositional formula α , whether $S \models \alpha$ holds for every maximal consistent subset S of Δ . It is known to be Π_2^P -complete [26]. We map any instance $\langle \Delta, \alpha \rangle$ of this problem to the following instance of skeptical inference from prototypical homogeneous forget-based inference:

- $B = \langle h_1 \Rightarrow \phi_1, \dots, h_n \Rightarrow \phi_n, h_1, \dots, h_n \rangle$, where each h_i ($i \in \{1, \dots, n\}$) is a fresh propositional symbol (not belonging to $\text{Var}(\Delta) \cup \text{Var}(\alpha)$);
- $\psi = \alpha$;
- $C = \bigwedge_{i=1}^{2n} (\bigwedge_{x \in \text{Var}(\Delta) \cup \text{Var}(\alpha)} \neg \text{forget}(x, i)) \wedge \bigwedge_{i,j \in \{1, \dots, 2n\}} \bigwedge_{k=1}^n (\text{forget}(h_k, i) \Rightarrow \text{forget}(h_k, j))$.

Since C is homogeneous, any forgetting vector for B given C is of the form $\vec{V}_J = \langle V_J, \dots, V_J \rangle$, where $V_J \subseteq \{h_1, \dots, h_n\}$. Let $J = \{i \in \{1, \dots, n\} \mid h_i \in V_J\}$.

We have

$$\exists V_J. (h_i \Rightarrow \phi_i) \equiv \begin{cases} \top & \text{if } i \in J, \\ h_i \Rightarrow \phi_i & \text{if } i \notin J \end{cases}$$

whereas

$$\exists V_J. h_i \equiv \begin{cases} \top & \text{if } i \in J, \\ h_i & \text{if } i \notin J. \end{cases}$$

Therefore,

$$(B \mid \vec{V}_J) \equiv \bigwedge_{i=1}^n (\exists V_J. (h_i \Rightarrow \phi_i) \wedge \exists V_J. h_i) \equiv \bigwedge_{j \in J} \top \wedge \bigwedge_{j \notin J} ((h_j \Rightarrow \phi_j) \wedge h_j) \equiv \bigwedge_{j \notin J} (h_j \wedge \phi_j).$$

Since $\{h_1, \dots, h_n\} \cap \text{Var}(\Delta) = \emptyset$, \vec{V}_J is a recovery for B given C if and only if $\{\phi_j \in \Delta \mid j \notin J\}$ is consistent, and \vec{V}_J is a preferred recovery w.r.t. \subseteq_p for B given C if and only if $\{\phi_j \in \Delta \mid j \notin J\}$ is a maximal (w.r.t. \subseteq) consistent subset of Δ . Therefore, we have that for every maximal consistent subset S of Δ , $S \models \alpha$ holds if and only if $B \approx_{\subseteq_p}^C \psi$. \square

Proof of Proposition 5.3. We first show that

$$\bigwedge_{i=1}^n \exists V_i . \phi_i \models \psi \quad \text{if and only if} \quad \bigwedge_{i=1}^n \text{rename}(\phi_i, V_i) \models \psi.$$

Consider any quantified formula of the form $\exists V_i . \phi_i$. Since quantified variables are dummy ones, $\exists V_i . \phi_i$ is equivalent to $\exists V_i^i . \text{rename}(\phi_i, V_i)$, where $V_i^i = \{x_i \mid x \in V_i\}$ and $\text{rename}(\phi_i, V_i)$ is the formula obtained by renaming in ϕ_i every $x \in V_i$ by x_i . As a consequence, we get that $\bigwedge_{i=1}^n \exists V_i . \phi_i$ is equivalent to $\bigwedge_{i=1}^n \exists V_i^i . \text{rename}(\phi_i, V_i)$. Now, for every $i, j \in \{1, \dots, n\}$ with $i \neq j$, no variable from V_i^i occurs in any $\exists V_j^j . \text{rename}(\phi_j, V_j)$. Hence, we have that $\bigwedge_{i=1}^n \exists V_i^i . \text{rename}(\phi_i, V_i)$ is equivalent to $\exists \bigcup_{i=1}^n V_i^i . (\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i))$. Finally, since ψ is independent from the new variables $\bigcup_{i=1}^n V_i^i$, we have that $\exists \bigcup_{i=1}^n V_i^i . (\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i)) \models \psi$ if and only if $\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i) \models \psi$.

- (1) If $\text{Pref}(\mathcal{R}_C(B), \sqsubseteq)$ – or only the maximal elements of it w.r.t. \subseteq_p – are provided then deciding whether $B = \langle \phi_1, \dots, \phi_n \rangle$ is such that $B \approx_{\sqsubseteq}^C \psi$ holds simply amounts to deciding whether $B \mid \vec{V} \models \psi$ for every given \vec{V} . Now, each $B \mid \vec{V} \models \psi$ can be polynomially reduced to an instance of the prototypical coNP-complete problem UNSAT. Indeed, $\langle \phi_1, \dots, \phi_n \rangle \mid \vec{V} \models \psi$ holds if and only if $\bigwedge_{i=1}^n \exists V_i . \phi_i \models \psi$ if and only if $\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i) \models \psi$ if and only if $\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i) \wedge \neg \psi$ is unsatisfiable. To conclude, it is enough to observe that a linear number of UNSAT instances can be polynomially reduced to a single one (through variable renaming), and since coNP is closed under polynomial reduction, the conclusion follows. A similar demonstration can be achieved for prudent inference (the first step consists in computing $\vec{V}_{prudent}$ which can be done easily in polynomial time when $\text{Pref}(\mathcal{R}_C(B), \sqsubseteq)$ – or only the maximal elements of it w.r.t. \subseteq_p – are given as part of the input; the second step is to polynomially reduce $B \mid \vec{V}_{prudent} \models \psi$ to an instance of UNSAT, as above).
- (2) If in addition every ϕ_i of B belongs to a propositional fragment which is stable by conjunction and by new variable renaming, then for each given \vec{V} (resp. for $\vec{V} = \vec{V}_{prudent}$), $\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i)$ belongs as well to this fragment. If this fragment is tractable for clausal entailment, deciding whether $\bigwedge_{i=1}^n \text{rename}(\phi_i, V_i) \models \psi$ holds can be achieved in polynomial time when ψ is a CNF formula. \square

References

- [1] O. Arieli, Paraconsistent reasoning and preferential entailments by signed quantified boolean formulae, *ACM Transactions on Computational Logic* 8 (3) (2007).
- [2] O. Arieli, Distance-based paraconsistent logics, *International Journal on Approximate Reasoning* 48 (3) (2008) 766–783.
- [3] C. Baral, S. Kraus, J. Minker, Combining multiple knowledge bases, *IEEE Transactions on Knowledge and Data Engineering* 3 (2) (1991) 208–220.
- [4] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, H. Prade, Inconsistency management and prioritized syntax-based entailment, in: *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI'93)*, 1993, pp. 640–645.
- [5] L.E. Bertossi, A. Hunter, T. Schaub (Eds.), *Inconsistency Tolerance*, Lecture Notes in Computer Science, vol. 3300, Springer, 2005.
- [6] Ph. Besnard, Remedying inconsistent sets of premises, *International Journal on Approximate Reasoning* 45 (2) (2007) 308–320.
- [7] Ph. Besnard, A. Hunter, Introduction to actual and potential contradictions, in: *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, vol. 2, Kluwer Academic Publishers, 1998, pp. 1–11.
- [8] Ph. Besnard, T. Schaub, Circumscribing inconsistency, in: *Proceedings of the 15th International Joint Conference on Artificial Intelligence (IJCAI'97)*, 1997, pp. 150–155.
- [9] Ph. Besnard, T. Schaub, Signed systems for paraconsistent reasoning, *Journal of Automated Reasoning* 20 (1) (1998) 191–213.
- [10] G. Boole, *An Investigation of the Laws of Thought*, Macmillan Publishers, Basingstoke, 1854.
- [11] G. Brewka, Preferred subtheories: An extended logical framework for default reasoning, in: *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI'89)*, 1989, pp. 1043–1048.
- [12] M. Cadoli, F.M. Donini, A survey on knowledge compilation, *AI Communications* 10 (1997) 137–150 (printed in 1998).
- [13] C. Cayrol, M.-C. Lagasque-Schiex, Th. Schiex, Nonmonotonic reasoning: From complexity to algorithms, *Annals of Mathematics and Artificial Intelligence* 22 (3–4) (1998) 207–236.
- [14] L. Cholvy, Automated reasoning with merged contradictory information whose reliability depends on topics, in: *Proceedings of the European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU'95)*, 1995, pp. 125–132.
- [15] S. Coste-Marquis, P. Marquis, Compiling stratified belief bases, in: *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI'00)*, 2000, pp. 23–27.
- [16] S. Coste-Marquis, P. Marquis, Recovering consistency by forgetting inconsistency, in: *Proceedings of the 11th European Conference on Logics in Artificial Intelligence (JELIA'08)*, in: *Lecture Notes in Artificial Intelligence*, vol. 5293, Springer, 2008, pp. 113–125.
- [17] J. Delgrande, T. Schaub, Two approaches for merging belief bases, in: *Proceedings of the 9th European Conference on Logics in Artificial Intelligence (JELIA'04)*, in: *Lecture Notes in Computer Science*, vol. 3229, 2004, pp. 426–438.
- [18] Th. Eiter, G. Gottlob, On the complexity of propositional knowledge base revision, updates and counterfactuals, *Artificial Intelligence* 57 (1992) 227–270.
- [19] Th. Eiter, G. Ianni, R. Schindlauer, H. Tompits, K. Wang, Forgetting in managing rules and ontologies, in: *Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence (WI'06)*, 2006, pp. 411–419.
- [20] Th. Eiter, K. Wang, Forgetting and conflict resolving in disjunctive logic programming, in: *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI'06)*, 2006, pp. 238–243.
- [21] Th. Eiter, K. Wang, Semantic forgetting in answer set programming, *Artificial Intelligence* 172 (14) (2008) 1644–1672.
- [22] E. Erdem, P. Ferraris, Forgetting actions in domain descriptions, in: *Proceedings of 22nd AAAI Conference on Artificial Intelligence (AAAI'07)*, 2007, pp. 409–414.
- [23] P. Everaere, S. Konieczny, P. Marquis, Conflict-based merging operators, in: *Proceedings of the 11th International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, 2008, pp. 348–357.

- [24] R. Fagin, J.D. Ullman, M.Y. Vardi, On the semantics of updates in databases, in: Proceedings of the 2nd ACM Symposium on Principles of Database Systems (PODS'83), 1983, pp. 352–355.
- [25] M.L. Ginsberg, Counterfactuals, *Artificial Intelligence* 30 (1986) 35–79.
- [26] G. Gottlob, Complexity results for nonmonotonic logics, *Journal of Logic and Computation* 2 (1992) 397–425.
- [27] S. Konieczny, J. Lang, P. Marquis, Distance-based merging: A general framework and some complexity results, in: Proceedings of the 8th International Conference on Principles and Knowledge Representation and Reasoning (KR'02), 2002, pp. 97–108.
- [28] S. Konieczny, J. Lang, P. Marquis, Quantifying information and contradiction in propositional logic through test actions, in: Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI'03), 2003, pp. 106–111.
- [29] S. Konieczny, J. Lang, P. Marquis, DA^2 merging operators, *Artificial Intelligence* 157 (2004) 45–79.
- [30] S. Konieczny, J. Lang, P. Marquis, Reasoning under inconsistency: The forgotten connective, in: Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI'05), 2005, pp. 484–489.
- [31] S. Konieczny, R. Pino Pérez, On the logic of merging, in: Proceedings of the 6th International Conference on Knowledge Representation and Reasoning (KR'98), 1998, pp. 488–498.
- [32] C. Lafage, J. Lang, Propositional distances and compact preference representation, *European Journal of Operation Research* 160 (2005) 741–761.
- [33] J. Lang, P. Liberatore, P. Marquis, Propositional independence – formula-variable independence and forgetting, *Journal of Artificial Intelligence Research* 18 (2003) 391–443.
- [34] J. Lang, P. Marquis, In search of the right extension, in: Proceedings of the 7th International Conference on Knowledge Representation and Reasoning (KR'00), 2000, pp. 625–636.
- [35] J. Lang, P. Marquis, Resolving inconsistencies by variable forgetting, in: Proceedings of the 8th International Conference on Principles and Knowledge Representation and Reasoning (KR'02), 2002, pp. 239–250.
- [36] P. Liberatore, F.M. Donini, Verification programs for abduction, in: Proceedings of the 14th European Conference on Artificial Intelligence (ECAI'00), 2000, pp. 166–170.
- [37] P. Liberatore, M. Schaerf, Arbitration (or how to merge knowledge bases), *IEEE Transactions on Knowledge and Data Engineering* 10 (1) (1998) 76–90.
- [38] P. Liberatore, M. Schaerf, Brels: A system for the integration of knowledge bases, in: Proceedings of the 7th International Conference on Knowledge Representation and Reasoning (KR'00), 2000, pp. 145–152.
- [39] F. Lin, R. Reiter, Forget it!, in: Proceedings of the AAAI Fall Symposium on Relevance, 1994, pp. 154–159.
- [40] J. Lin, A.O. Mendelzon, Knowledge base merging by majority, in: *Dynamic Worlds: From the Frame Problem to Knowledge Management*, Kluwer, 1999.
- [41] B. Nebel, Syntax-based approaches to belief revision, in: *Belief Revision*, in: Cambridge Tracts in Theoretical Computer Science, vol. 29, Cambridge University Press, 1992, pp. 52–88.
- [42] B. Nebel, How hard is it to revise a belief base?, in: *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, vol. 3, Belief Revision, Kluwer Academic, 1998, pp. 77–145.
- [43] Ch.H. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.
- [44] G. Pinkas, R.P. Loui, Reasoning from inconsistency: A taxonomy of principles for resolving conflict, in: Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning (KR'92), 1992, pp. 709–719.
- [45] G. Priest, Minimally inconsistent LP, *Studia Logica* 50 (1991) 321–331.
- [46] N. Rescher, R. Manor, On inference from inconsistent premises, *Theory and Decision* 1 (1970) 179–219.
- [47] P.Z. Revesz, On the semantics of arbitration, *International Journal of Algebra and Computation* 7 (2) (1997) 133–160.
- [48] Y. Shoham, A semantical approach to nonmonotonic logics, in: Proceedings of the 2nd IEEE Symposium on Logic in Computer Science (LICS'87), 1987, pp. 275–279.
- [49] K. Su, G. Lv, Y. Zhang, Reasoning about knowledge by variable forgetting, in: Proceedings of the 9th International Conference on Principles and Knowledge Representation and Reasoning (KR'04), 2004, pp. 576–586.
- [50] K. Wang, A. Sattar, K. Su, A theory of forgetting in logic programming, in: Proceedings of the 20th National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference (AAAI'05), 2005, pp. 682–688.
- [51] Y. Zhang, N.Y. Foo, Solving logic program conflict through strong and weak forgettings, *Artificial Intelligence* 170 (8–9) (2006) 739–778.
- [52] Y. Zhang, N.Y. Foo, K. Wang, Solving logic program conflict through strong and weak forgettings, in: Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI'05), 2005, pp. 627–634.
- [53] Y. Zhang, Y. Zhou, Properties of knowledge forgetting, in: Proceedings of the 12th International Workshop on Non-Monotonic Reasoning (NMR'08), 2008, pp. 68–75.