Quantifying information and contradiction in propositional logic through test actions

Sébastien Konieczny
IRIT - UPS
F-31062 Toulouse - France
konieczny@irit.fr

Jérôme Lang
IRIT - UPS
F-31062 Toulouse - France
lang@irit.fr

Pierre Marquis
CRIL - Université d'Artois
F-62307 Lens - France
marquis@cril.univ-artois.fr

Abstract

Degrees of information and of contradiction are investigated within a uniform propositional framework, based on test actions. We consider that the degree of information of a propositional formula is based on the cost of actions needed to identify the truth values of each atomic proposition, while the degree of contradiction of a formula is based on the cost of actions needed to make the formula classically consistent. Our definitions are to a large extent independent of the underlying propositional logic; this flexibility is of prime importance since there is no unique, fully accepted logic for reasoning under inconsistency.

1 Introduction

Information and contradiction are two fundamental aspects of knowledge processing. Quantifying them is an important issue when reasoning about beliefs (or preferences) stemming from one or different sources. Here are some contexts where quantifying information and contradiction is relevant:

• *diagnosis and testing*. In model-based diagnosis, some initial assumptions that each component works correctly are made; these assumptions may conflict with actual observations. Measuring the conflict of the resulting base may be a good hint about how hard it will be to identify the faulty components.

• *preference elicitation*. In the interactive process of elicitating the preference profile of an individual (user) about a set of possible alternatives, it is not unfrequent that contradictory preferences arise. In this situation, it is useful to quantify and localize the contradictions as well as the information about the user's preferences, so as to be in position to choose the next questions to ask.

• *belief merging*. In this framework, degrees of information and contradiction can be the basis on which one can decide whether to take or not into account the data conveyed by an agent. If the degree of contradiction of the data given by an agent is high, it may be relevant to reject the information, since there is a significant evidence that the source is not reliable; however, this must be balanced by the quantity of information furnished by the agent, especially when she also gives some important and uncontroversial pieces of information.

• *group decision making*. Contradictions arise frequently when trying to reach a compromise among several agents who have possibly conflictual preferences about a common decision (like voting, resource sharing, public goods buying). In this context, not only it is relevant to compute a global degree of conflict (among the set of agents) but also degrees of conflicts associated with small groups of agents (*coalitions*) so as to localize as precisely as possible where the conflicts are.

Now, what do "degree of information" and "degree of contradiction" mean? There is no consensus about it. The main features shared by existing definitions (and there are not numerous, cf. Section 7) is that (1) such degrees are numerical values, and (2) they vary depending on the representation language. Thus, one may consider a as fully informative in the case where a is the single atomic proposition of the language but surely not fully informative when the vocabulary also contains b (provided that a and b are independent propositions).

In this paper, our point of view is that it is inadequate to quantify the information/contradiction conveyed by some data without considering at the same time a set of available actions and a goal to be reached. Accordingly, our degrees of information and contradiction are defined in an "active" way. Acting so as to reduce inconsistency or to gain information often relies on performing *knowledge-gathering* actions (also called *tests*). We consider that the degree of information of an information base is based on the number (or the cost) of actions needed to identify the truth value of each atomic proposition (the lower the cost the more informative the base); and that the degree of contradiction of an information base is based on the number (or the cost) of actions needed to render the base classically consistent. Thus, both degrees are dependent on the language but also on the given set of tests and the way plans costs are computed.

The rest of this paper is organized as follows. After some formal preliminaries in Section 2, we present our framework in Section 3. In order to show the generality of our framework, we instantiate it to three different propositional logics: classical logic (Section 4), the paraconsistent logic LP_m (Section 5) and a syntax-based approach to inconsistency handling (Section 6). Related work is given in Section 7, and some conclusions in Section 8.

2 Formal preliminaries and notations

We consider a propositional language $\mathcal{L}_{PS}$ based on a finite set of propositional symbols PS and a set of connectives that may vary depending on the logic used. Well-formed formulas are denoted by φ, ψ, etc. The available information is represented by a formula Σ of $\mathcal{L}_{PS}$ (called *information base*).

Since we wish to keep a reasonable level of generality, we do not want to commit right now to a particular propositional logic. Indeed, flexibility is a major feature of our framework, and it is essential at least for two reasons. On the one hand, classical logic is inadequate to handle contradiction: *ex falso quodlibet sequitur* (a single, local contradiction is sufficient to pollute the whole information base). On the other hand, there is no unique, fully accepted, logic for reasoning under inconsistency. This is reflected by the various approaches that can be found in the literature under different names like belief revision, merging, paraconsistent logics, etc.

$\mathcal{L}_{PS}$ is required to possess the following ingredients:

1. A *consequence relation* $\models_L$ on $\mathcal{L}_{PS} \times \mathcal{L}_{PS}$.
2. An *acceptance function* $A_L \subseteq \mathcal{L}_{PS} \times \mathcal{L}_{PS}$: $A_L(\Sigma, \varphi)$ means that given the information base Σ, φ is accepted as a true information (we say that Σ *accepts* φ). By default, acceptance is defined by: $A_L(\Sigma, \varphi)$ iff ($\Sigma \models_L \varphi$ and $\Sigma \not\models_L \neg\varphi$). We say that Σ is *informative about* φ iff exactly one of $A_L(\Sigma, \varphi)$ or $A_L(\Sigma, \neg\varphi)$ holds, and that Σ is *fully informative* iff for any $\varphi \in \mathcal{L}_{PS}$, Σ is informative about φ.
3. A *contradiction indicator* $C_L \subseteq \mathcal{L}_{PS} \times \mathcal{L}_{PS}$: if $C_L(\Sigma, \varphi)$ holds, then we say that Σ *is contradictory about* φ. By default, we define $C_L(\Sigma, \varphi)$ iff ($\Sigma \models_L \varphi$ and $\Sigma \models_L \neg\varphi$). Σ is said to be *contradiction-free* iff for every $\varphi \in \mathcal{L}_{PS}$, we do not have $C_L(\Sigma, \varphi)$.
4. A *weak revision* operator $\star : \mathcal{L}_{PS} \times \mathcal{L}_{PS} \to \mathcal{L}_{PS}$: $\Sigma \star \varphi$ represents the new information base obtained once taking account of the observation φ into the information base Σ. For the sake of generality, we are not very demanding about $\star$. The only requirement is that $\Sigma \star \varphi \models_L \varphi$, which expresses that our tests are assumed reliable (each test outcome must be true in the actual world)[1]. In the following we will simply refer to those operators as *revision* operators (omitting the *weak*). It would be interesting to explore in more details what happens when one puts more requirements on $\star$ (this is left for further research).

3 Degree of ignorance and degree of contradiction

In this section, we give general definitions of degrees of contradiction and ignorance, i.e., of lack of information, which are to a large extent independent of the logic chosen to represent information. These definitions will be specified further in Sections 4 to 6 where the logic will be fixed.

Definition 3.1 (test contexts) *A* test context $\mathcal{C}_{\mathcal{L}_{PS}}$ *(w.r.t. $\mathcal{L}_{PS}$) is a pair $\langle T, c\rangle$ where T is a finite set of tests and c is a cost function from T to $\mathbb{N}^*$. The* outcome *to any test $t_\varphi \in T$ is one of φ, $\neg\varphi$, where $\varphi \in \mathcal{L}_{PS}$. We say that t_φ is the test about φ, and we assume w.l.o.g. that at most one test t_φ of T is about φ for each $\varphi \in \mathcal{L}_{PS}$. A context is* standard *iff $\forall t_\varphi \in T$, we have $c(t_\varphi) = 1$ (every test has a unit cost). A context is* universal *iff for every $\varphi \in \mathcal{L}_{PS}$, there is a test $t_\varphi \in T$. A context is* atomic *iff the testable formulas are exactly the atomic propositions ($t_x \in T$ iff $x \in PS$).*

Definition 3.2 (test plans, trajectories) *Given a test context $\mathcal{C}_{\mathcal{L}_{PS}}$, a* test plan *$\pi$ is a finite binary tree; each of its non-terminal nodes is labelled with a test action t_φ; the left and right arcs leaving every non-terminal node labelled with t_φ are respectively labelled with the outcomes φ and $\neg\varphi$. An* (outcome) trajectory *$\langle o_1, \ldots, o_n\rangle$ with respect to π is the sequence of test outcomes on a branch of π. The* cost *of a trajectory $\langle o_1, \ldots, o_n\rangle$ with respect to π is defined as $\Sigma_{i=1}^n c(t_{\varphi_i})$, where each t_{φ_i} is the test labelling the node of π reached by following the path $\langle o_1, \ldots, o_{i-1}\rangle$ from the root of π.*

Definition 3.3 (disambiguation, purification) *Let π be a test plan and Σ the initial information base.*

- *The application of π on Σ is the tree $apply(\pi, \Sigma)$, isomorphic to π, whose nodes are labelled with information bases defined inductively as follows:*
  - *the root ϵ of $apply(\pi, \Sigma)$ is labelled with $\Sigma(\epsilon) = \Sigma$;*
  - *let n be a node of $apply(\pi, \Sigma)$, labelled with the information base $\Sigma(n)$, whose corresponding node in π is non-terminal and labelled with t_φ; then n has two children in $apply(\pi, \Sigma)$, labelled respectively with $\Sigma(n) \star \varphi$ and $\Sigma(n) \star \neg\varphi$.*
- *π* disambiguates *φ* given *Σ iff for every terminal node n of $apply(\pi, \Sigma)$, $\Sigma(n)$ is informative about φ (i.e., either $A_L(\Sigma, \varphi)$ or $A_L(\Sigma, \neg\varphi)$). π* (fully) disambiguates *Σ iff it disambiguates all formulas of $\mathcal{L}_{PS}$, i.e., iff for any terminal node n of $apply(\pi, \Sigma)$, $\Sigma(n)$ is fully informative.*
- *π* purifies *φ* given *Σ iff for every terminal node n of $apply(\pi, \Sigma)$, $\Sigma(n)$ is not contradictory about φ (i.e., not $C_L(\Sigma(n), \varphi)$). π* (fully) purifies *Σ iff it eliminates all contradictions in Σ, i.e., iff for any terminal node n of $apply(\pi, \Sigma)$, $\Sigma(n)$ is contradiction-free.*

Clearly enough, it can be the case that there is no plan to purify or disambiguate a formula; especially, there is no plan for purifying the constant $\perp$ (false) in any of the three logics considered in the following.

In our framework, degrees of ignorance and contradiction are defined in a uniform way: the degree of ignorance (resp. of contradiction) of Σ measures the minimal effort necessary to disambiguate (resp. to purify) Σ.

Definition 3.4 (degree of contradiction, of ignorance) *Let us define the* cost *$c(\pi)$ of a test plan π as the maximum of the costs of its trajectories. Then*

- *The* degree of contradiction *of Σ is defined by $d_C(\Sigma) = min(\{c(\pi) \mid \pi \text{ purifies } \Sigma\})$. When no plan purifies Σ, we let $d_C(\Sigma) = +\infty$.*
- *The* degree of ignorance *of Σ is defined by $d_I(\Sigma) = min(\{c(\pi) \mid \pi \text{ disambiguates } \Sigma\})$. When no plan disambiguates Σ, we let $d_I(\Sigma) = +\infty$.*

Clearly, both degrees depend on the test context; by default, we consider the standard atomic context.

In the previous definition, we actually define *pessimistic degrees* of contradiction and ignorance (because the cost of

[1]While this condition corresponds to the "success" postulate in the AGM framework [Alchourrón *et al.*, 1985], we do not ask our operators to be AGM ones, especially because the AGM framework is not suited to characterize the revision of an inconsistent formula.

a plan is defined as the *maximum* cost among its trajectories); this principle, consisting in assuming the worst outcome, is known in decision theory as *Wald criterion*. Other criteria could be used instead, such as the optimistic criterion obtained by replacing max by min. More interesting, the criterion obtained by first using max and then min for tie-breaking, or the $leximax$ criterion, allow for a better discrimination than the pure pessimistic criterion. The choice of a criterion is fairly independent from the other issues discussed in this paper, which gives our framework a good level of flexibility and generality. Due to space limitations, however, we consider only the pessimistic criterion in the rest of the paper.

The definitions below concern the cost of being informative about a formula φ (resp. of getting rid of contradictions about φ) given an initial information base Σ.

Definition 3.5 (disambiguation cost, purification cost)

- *The* disambiguation cost *of φ given Σ, denoted by $C_D(\Sigma, \varphi)$, is the minimum cost of a test plan applied to Σ and disambiguating φ.*
- *The* purification cost *of φ given Σ, denoted by $C_P(\Sigma, \varphi)$, is the minimum cost of a test plan applied to Σ and purifying φ.*

The disambiguation cost of φ given Σ can be seen as the price one is ready to pay for being informed about φ. The purification cost of φ given Σ is a bit less intuitive: in a database context, for instance, it represents the cost that the database holder is ready to pay to ensure that the database does not contain any contradictory information about a given formula.

Proposition 3.1 *The degrees defined above are related as follows:*

- $d_I(\Sigma) \geq d_C(\Sigma)$;
- $C_D(\Sigma, \varphi) = C_D(\Sigma, \neg\varphi)$;
- $C_P(\Sigma, \varphi) = C_P(\Sigma, \neg\varphi)$;
- $C_D(\Sigma, \varphi) \geq C_P(\Sigma, \varphi)$;
- $d_I(\Sigma) \geq C_D(\Sigma, \varphi)$;
- $d_C(\Sigma) \geq C_P(\Sigma, \varphi)$.

4 Case study 1: classical logic

In this section, we fix $\mathcal{L}_{PS}$ to classical propositional logic CL, which means that we fix the following:

- The language of CL is built up from the connectives $\wedge$, $\vee$, $\neg$, $\rightarrow$ and the constants $\top$, $\bot$;
- Interpretations are mappings from PS to $\{\mathsf{T}, \mathsf{F}\}$ and the consequence relation is defined as usual by $\Sigma \models_{CL} \varphi$ iff $Mod_{CL}(\Sigma) \subseteq Mod_{CL}(\psi)$, where $Mod_{CL}(\Sigma)$ denotes the set of classical models of Σ.

Clearly enough, $A_{CL}(\Sigma, \varphi)$ holds iff $\Sigma \models_{CL} \varphi$ and Σ is consistent; and $C_{CL}(\Sigma, \varphi)$ holds iff Σ is inconsistent. Finally, let $\Sigma \star \varphi = \Sigma \wedge \varphi$ (i.e., the revision operator is simple expansion).

Example 4.1 *Figure 1 reports a test plan of minimal cost that fully disambiguates $\Sigma = (a \vee b) \wedge (a \rightarrow c)$ (w.r.t. the standard atomic test context).*

Accordingly, the degree of ignorance of Σ is 2 given the standard atomic context; as Figure 1 also shows, such a degree can be strictly lower than the number of literals built up from PS that are not accepted by Σ (here, the information given by the test outcomes within the test plan are crucial).

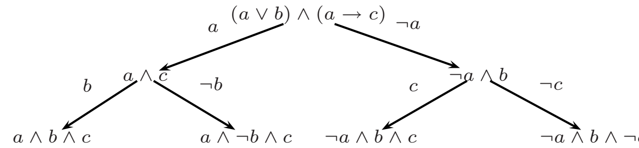

Figure 1: Degree of ignorance of $(a \vee b) \wedge (a \rightarrow c)$

Proposition 4.1 *Given any test context:*

- $C_D(\Sigma, \varphi) = 0$ *iff Σ is consistent and ($\Sigma \models \varphi$ or $\Sigma \models \neg\varphi$);*
- *if Σ is inconsistent, then $C_D(\Sigma, \varphi) = +\infty$;*
- $C_D(\Sigma, \varphi \vee \psi) \leq C_D(\Sigma, \varphi) + C_D(\Sigma, \psi)$;
- $C_D(\Sigma, \varphi \wedge \psi) \leq C_D(\Sigma, \varphi) + C_D(\Sigma, \psi)$;
- $d_I(\Sigma) \leq \sum_{x \in PS} C_D(\Sigma, x)$.

Degrees of conflict and purification are not relevant when $\mathcal{L}_{PS}$ is classical logic, just because inconsistency is an extremely rough notion in classical logic. Indeed, we have:

$$d_C(\Sigma) = \begin{cases} 0 & \text{if } \Sigma \text{ is consistent} \\ +\infty & \text{if } \Sigma \text{ is inconsistent} \end{cases}$$

and $C_P(\Sigma, \varphi) = d_C(\Sigma)$ for all φ.

Example 4.2 *Let a, b be two (independent) atomic propositions we focus on. Starting from an information base represented by $\top$ (no contradiction, no information), stepwise expansions progressively lead to a base Σ having a single model over $\{a, b\}$ like ($a = \mathsf{T}, b = \mathsf{T}$). In this situation, the degrees of ignorance and of contradiction (given the standard atomic context) are minimal (no contradiction, full information). But each additional expansion that does not lead to an equivalent information base leads to an inconsistent one; for instance, the resulting base $\Sigma \star \neg a$ has a maximal degree of contradiction. This is counterintuitive since it seems that $a \wedge b \wedge \neg a$ is both more informative and less contradictory than $a \wedge b \wedge \neg a \wedge \neg b$. This is not reflected by our degrees when classical logic is considered.*

- $d_I(\{\top\}) = 2$, $d_C(\{\top\}) = 0$
- $d_I(\{a\}) = 1$, $d_C(\{a\}) = 0$
- $d_I(\{a \wedge b\}) = 0$, $d_C(\{a \wedge b\}) = 0$
- $d_I(\{a \wedge b \wedge \neg a\}) = +\infty$, $d_C(\{a \wedge b \wedge \neg a\}) = +\infty$
- $d_I(\{a \wedge b \wedge \neg a \wedge \neg b\}) = +\infty$, $d_C(\{a \wedge b \wedge \neg a \wedge \neg b\}) = +\infty$

This example also shows that mere expansion is not a very satisfying revision operator. Indeed, since it does not enable to purify any inconsistent base (whatever the test context), expansion does not enable as well to disambiguate any inconsistent base. Furthermore, it may lead to degrees of contradiction (or purification costs) that are not intuitively correct. Thus, on the example, we have $d_I(\{a \wedge b \wedge \neg a\}) = +\infty$, while given the standard atomic context, two tests are sufficient to determine the actual world (over $\{a, b\}$). The reason of this discrepancy between what is expected and what is achieved is that expanding an inconsistent information base always leads to an inconsistent base, while it would be necessary to restore consistency[2] for achieving purification and disambiguation in classical logic. Note that using AGM revision instead of expansion would not help a lot since AGM operators do not behave well when applied to inconsistent bases.

[2] A way to do it consists in *forgetting information* [Lang and Marquis, 2002] (possibly everything) in $\{a \wedge b \wedge \neg a\}$.

5 Case study 2: the paraconsistent logic LP_m

Paraconsistent logics have been introduced to avoid *ex falso quodlibet sequitur* of classical logic, hence handling inconsistent bases in a much more satisfying way. While many paraconsistent logics have been defined so far and could be used in our framework, we focus here on the LP_m logic as defined in [Priest, 1991]. This choice is mainly motivated by the fact that this logic is simple enough and has an inference relation that coincides with classical entailment whenever the information base is classically consistent (this feature is not shared by many paraconsistent logics).

- The language of LP_m is built up from the connectives $\wedge$, $\vee$, $\neg$, $\rightarrow$ and the constants $\top$, $\bot$.
- An interpretation ω for LP_m maps each propositional atom to one of the three "truth values" $\mathsf{F}, \mathsf{B}, \mathsf{T}$, the third truth value B meaning intuitively "both true and false". 3^{PS} is the set of all interpretations for LP_m. "Truth values" are ordered as follows: $\mathsf{F} <_t \mathsf{B} <_t \mathsf{T}$.
  - $\omega(\top) = \mathsf{T}, \omega(\bot) = \mathsf{F}$
  - $\omega(\neg\alpha) = \mathsf{B}$ iff $\omega(\alpha) = \mathsf{B}$
    $\omega(\neg\alpha) = \mathsf{T}$ iff $\omega(\alpha) = \mathsf{F}$
  - $\omega(\alpha \wedge \beta) = \min_{\leq_t}(\omega(\alpha), \omega(\beta))$
  - $\omega(\alpha \vee \beta) = \max_{\leq_t}(\omega(\alpha), \omega(\beta))$
  - $\omega(\alpha \rightarrow \beta) = \begin{cases} \mathsf{T} & \text{if } \omega(\alpha) = \mathsf{F} \\ \omega(\beta) & \text{otherwise} \end{cases}$
- The set of models of a formula φ is $Mod_{LP}(\varphi) = \{\omega \in 3^{PS} \mid \omega(\varphi) \in \{\mathsf{T}, \mathsf{B}\}\}$. Define $\omega! = \{x \in PS \mid \omega(x) = \mathsf{B}\}$. Then $min(Mod_{LP}(\varphi)) = \{\omega \in Mod_{LP}(\varphi) \mid \nexists \omega' \in Mod_{LP}(\varphi)$ s.t. $\omega'! \subset \omega!\}$. The consequence relation is defined by $\Sigma \models_{LP_m} \varphi$ iff $min(Mod_{LP}(\Sigma)) \subseteq Mod_{LP}(\varphi)$.
- The definitions of $A_{LP_m}(\Sigma, \varphi)$ and $C_{LP_m}(\Sigma, \varphi)$ are those by default; $C_{LP_m}(\Sigma, \varphi)$ holds only if Σ has no classical model.

Now, what about the revision operator? Actually, the issue of revision in paraconsistent logic has never been considered so far. Expansion is not satisfactory as a revision operator for LP_m because it enables neither the purification task nor the disambiguation one when Σ has no classical model ω (i.e., such that $\omega(x) \neq \mathsf{B}$ for each $x \in PS$), whatever the test context. Among the many possible choices, we have considered the following revision operator, defined model-theoretically (for the sake of brevity, we characterize only its restriction to the case the revision formula φ is a literal l).

Let $force(\omega, l)$ be the interpretation of 3^{PS} defined by (for every literal $l = x$ or $\neg x$):

$$\begin{cases} force(\omega, x)(x) = \mathsf{T} \\ \forall y \in PS, y \neq x, force(\omega, x)(y) = \omega(y) \end{cases}$$

$$\begin{cases} force(\omega, \neg x)(x) = \mathsf{F} \\ \forall y \in PS, y \neq x, force(\omega, x)(y) = \omega(y) \end{cases}$$

Then the revision operator is defined by:

$$Mod_{LP}(\Sigma \star l) = \begin{cases} \{\omega \models \Sigma \mid \omega(l) = \mathsf{T}\} \text{ if this set is non-empty, otherwise} \\ \{force(\omega, l) \mid \omega \models \Sigma \text{ and } \omega(l) = \mathsf{B}\}. \end{cases}$$

Example 5.1 *Stepping back to Example 4.2, we can check that LP_m leads to more intuitive values for the degrees of ignorance and contradiction of the bases we considered. Given the standard atomic test context, we now have:*

- $d_I(\{\top\}) = 2$, $d_C(\{\top\}) = 0$
- $d_I(\{a\}) = 1$, $d_C(\{a\}) = 0$
- $d_I(\{a \wedge b\}) = 0$, $d_C(\{a \wedge b\}) = 0$
- $d_I(\{a \wedge b \wedge \neg a\}) = 1$, $d_C(\{a \wedge b \wedge \neg a\}) = 1$
- $d_I(\{a \wedge b \wedge \neg a \wedge \neg b\}) = 2$, $d_C(\{a \wedge b \wedge \neg a \wedge \neg b\}) = 2$

Thus, starting from $\top$, expansion (i.e. the logical strengthening of an information base) leads to more information up to a point where full information is reached, then progressively to more contradiction and less information, up to an information base with full contradiction and no information.

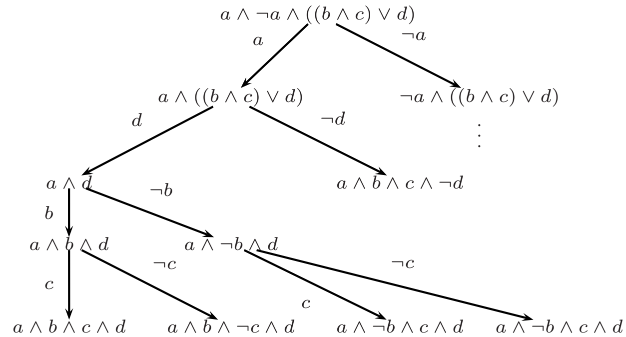

Figure 2: Degrees of ignorance and contradiction in LP_m

Example 5.2 *Let us now consider the base $\Sigma = \{a \wedge \neg a \wedge ((b \wedge c) \vee d)\}$. Figure 2 reports a plan of minimal cost (given the standard atomic context) which disambiguates Σ (the subtree rooted below the $\neg a$ outcome is similar to the one of a and is not represented). Here, the degree of ignorance of Σ is 4 and the degree of contradiction of Σ is 1 (after testing a there is no contradiction any more).*

Contrariwise to expansion, this revision operator is adequate to our objective:

Proposition 5.1 *Every information base that has a model $\in 3^{PS}$ has a finite degree of ignorance and a finite degree of contradiction given any atomic or universal test context.*

6 Case study 3: "syntax-based" information bases

Many approaches to reasoning under inconsistency make use of the selection of maximal consistent subsets of a belief base. This principle comes down to [Rescher and Manor, 1970] and is known under different names such as the possible-world approach [Ginsberg and Smith, 1988], assumption-based theories [Poole, 1988], supernormal default theories [Brewka, 1989], syntax-based approaches [Nebel, 1991] etc. We retain the last name and call the logic SBL[3].

The language of SBL is defined as the set of all pairs $\Sigma = \langle \alpha, \Delta \rangle$, where $\Delta = \{\delta_1, \ldots, \delta_n\}$; α, δ_1, ..., δ_n are propositional formulas from the language of CL. α represent the set of hard facts and $\delta_1, \ldots, \delta_n$ the defaults. Formulas of the form $\langle \alpha, \emptyset \rangle$ are called *simple formulas* and, slightly abusing notations, are abbreviated into α – thus the language of SBL can be seen as an extension of that of classical logic. A

[3] Saying that the approach is "syntax-based" actually means that $\{\varphi_1, \ldots, \varphi_n\}$ should not be identified with $\varphi_1 \wedge \ldots \wedge \varphi_n$, or, in other words, that the comma is viewed as a (non truth-functional) connective.

maximal scenario of Σ is a subset Δ' of Δ such that $\alpha \wedge \Delta'$ is consistent and there is no $\Delta'' \supset \Delta'$ such that $\alpha \wedge \Delta''$ is consistent. We denote by $MS(\Sigma)$ the set of all maximal scenarios of Σ.

The semantics of SBL is defined by a preference relation over classical interpretations. Let $\Sigma = \langle \alpha, \Delta \rangle$. If $\omega \in 2^{PS}$, we define $Sat(\omega, \Sigma) = \{i \mid \omega \models_{CL} \delta_i\}$. For $\omega, \omega' \in Mod_{CL}(\alpha)$, we say that $\omega \geq_\Sigma \omega'$ iff $Sat(\omega, \Sigma) \supseteq Sat(\omega', \Sigma)$. Now, $Mod_{SBL}(\Sigma) = \{\omega \mid \omega \models_{CL} \alpha$ and there is no $\omega' \models_{CL} \alpha$ such that $\omega' >_\Sigma \omega\}$. Consequence in SBL is defined by $\Sigma \models_{SBL} \Sigma'$ iff $Mod_{SBL}(\Sigma) \subseteq Mod_{SBL}(\Sigma')$ – especially, if φ is a simple formula, then $\Sigma \models_{SBL} \varphi$ iff φ is a skeptical consequence of Σ in the sense of default logic.

Let φ be a simple formula. Acceptance is defined by the default formulation. Equivalently, $A_{SBL}(\langle \alpha, \Delta \rangle, \varphi)$ holds iff $\langle \alpha, \Delta \rangle \models_{SBL} \varphi$ and α is consistent. Contradiction refers to maximal scenarios: $C_{SBL}(\Sigma, \varphi)$ iff $MS(\Sigma) = \emptyset$ or there exists $\Delta', \Delta'' \in MS(\Sigma)$ such that $\alpha \wedge \Delta' \models_{CL} \varphi$ and $\alpha \wedge \Delta'' \models_{CL} \neg\varphi$. Lastly, $\langle \alpha, \Delta \rangle \star \varphi$ is defined as $\langle \alpha \wedge \varphi, \Delta \rangle$, i.e., the hard facts are expanded with the revision formula.

Example 6.1 $\Sigma = \langle \emptyset, \{a, \neg a \wedge b, a \rightarrow c, \neg c, c \rightarrow a\} \rangle$. *Here is a purification plan of minimal cost for* Σ *with respect to the standard atomic context. We start by testing a;* $\Sigma \star \neg a = \langle \neg a, \Delta \rangle$ *is contradiction-free because it has a single maximal scenario* $\{\neg a \wedge b, a \rightarrow c, \neg c, c \rightarrow a\}$*;* $\Sigma \star a = \langle a, \Delta \rangle$ *is not contradiction-free because it has two maximal scenarios* $\{a, a \rightarrow c, c \rightarrow a\}$ *and* $\{a, \neg c, c \rightarrow a\}$*; then testing c leads to* $\Sigma \star (a \wedge c)$ *and* $\Sigma \star (a \wedge \neg c)$*, both being contradiction-free. Hence, given the standard atomic context,* $d_C(\Sigma) = 2$ *and* $d_I(\Sigma) = 3$ *(since disambiguating b will require one additional test).*

Proposition 6.1

- Σ *is contradiction-free iff* $MS(\Sigma)$ *is a singleton.*
- $\Sigma = \langle \alpha, \Delta \rangle$ *is fully informative iff* $MS(\Sigma)$ *is a singleton* $\{\Delta'\}$ *and* $Mod_{CL}(\alpha \wedge \Delta')$ *is a singleton.*

Despite it mainly amounts to expansion of the hard facts, if the hard constraint α of Σ is true for sure in the actual world (i.e., it is consistent and consistency cannot be questioned by expansion with test outcomes), then Σ has a finite degree of contradiction and a finite degree of ignorance given any atomic or universal context.

7 Related work

To the best of our knowledge, only few proposals for a notion of degree of information can be found in the literature, and things are even worse to what concerns the notion of degree of contradiction. All existing approaches are stuck to specific propositional logics with the corresponding consequence relations, which address only some aspects of the paraconsistency issue, if any (as evoked previously, there is no undebatable paraconsistent inference relation).

Shannon's information theory [Shannon, 1948] provides the most famous approach on which notions of quantity of information can be defined, but it relies on the assumption that the available information is given under the form of a probability distribution; furthermore, it cannot directly address inconsistent data. Interestingly, our definition of degree of information is general enough to recover classical entropy, applied to classical logic[4].

Lozinskii [1994a] gives a set of properties that a measure of quantity of information should satisfy. Our degree of ignorance is fully compatible with Lozinskii's requirements in several cases. The degree of information defined by Lozinskii corresponds to the notion in Shannon's theory, assuming a uniform distribution over the set of propositional interpretations[5]. It is thus required that the input information base Σ has a classical model. [Lozinskii, 1994b] extends [Lozinskii, 1994a] by considering as well some inconsistent logical systems, through a more general notion of model. It is specifically focused on so-called quasi-models of the information set Σ, which are the models of the maximal (w.r.t. $\subseteq$) consistent subsets of Σ. This is sufficient to avoid the notion of degree of information to trivialize when applied to an inconsistent Σ (unless it is a singleton). However, no notion of degree of contradiction is specifically introduced in Lozinskii's approach. He claims that "an inconsistent information base always contains less semantic information than any of its maximal consistent subsets". In our point of view, this may lead to counterintuitive results. For instance, given $PS = \{a, b, c\}$, according to Lozinskii, $\{a\}$ contains more information than $\{a \wedge \neg a \wedge b \wedge c\}$. That is not always the case in our framework: it depends on the context under consideration. Thus, while the degree of contradiction of the former is lower than the degree of contradiction of the latter, the degree of information of the latter is greater than the one of the former (w.r.t. the standard atomic test context).

Wong and Besnard [2001] criticize the syntax-sensitivity of Lozinskii's approach. In particular, the presence of tautologies in Σ may unexpectedly change the quantity of information. As they note this can be easily repaired by considering the models of Σ over the set of variables on which Σ depends [Lang and Marquis, 1998], instead of the set of variables occurring in Σ. Wong and Besnard also adhere to Lozinskii's definition of degree of information; what changes is the underlying notion of quasi-model, since the paraconsistent logic they considered is quasi-classical logic [Besnard and Hunter, 1995; Hunter, 2000].

[Knight, 2003] reports some other postulates for a measure of quantity of information. Our measure d_I does not satisfy all of them, even in simple cases (for space reasons, we cannot detail it here). This contrasts with the two measures introduced by Knight, which generalize in an elegant way Shannon's entropy-based measure to the case the information base is an inconsistent set of formulas. However, both measures trivialize when the information set is an inconsistent singleton.

The only two approaches we are aware of, which consider (non-trivial) degrees of inconsistency defined for clas-

[4] Let pr_Σ be the probability distribution stating that all models of Σ over PS are equiprobable (and the other ones impossible), i.e., for any $\omega \models \Sigma$, $pr_\Sigma(\omega) = \frac{1}{|Mod_{CL}(\Sigma)|}$. The entropy of pr_Σ is defined as $H(pr_\Sigma) = \sum_{\omega \models \Sigma} -pr_\Sigma(\omega).log_2 pr_\Sigma(\omega) = log_2|Mod_{CL}(\Sigma)|$. Therefore, the integer upper part of $H(pr_\Sigma)$ is nothing but the minimal number of tests that have to be performed to identify the actual world given the standard universal test context.

[5] Such a degree was already known by Kemeny [1953] and Hintikka [1970].

sical formulas (i.e., without additional information, like a preference preorder), are [Knight, 2002] and [Hunter, 2002]. Knight finds a set Σ of classical formulas maximally α-consistent whenever α is the greatest number for which a probability measure P exists, satisfying $P(\varphi) \geq \alpha$ for every $\varphi \in \Sigma$. Accordingly, his measure makes sense when the logic SBL is considered, but trivializes whenever Σ is a singleton. Hunter appeals to quasi-classical logic as a framework for dealing with inconsistent information. Considering the minimal quasi-classical models of the information base Σ (roughly, the "most classical" ones), Hunter defines the degree of coherence of Σ as the ratio between the amount of contradiction (conflict) and the amount of information (opinion) of the formula. Such a degree of coherence does not always give what is expected. Consider for instance the two information sets $\Sigma = \{e \wedge \neg e, a \vee (b \wedge c \wedge d)\}$ and $\Sigma' = \{e \wedge \neg e, (a \wedge b) \vee (c \wedge d)\}$. [Hunter, 2002] shows that *Coherence*(Σ) > *Coherence*(Σ'). We feel it counterintuitive since the two bases shares the same contradiction $e \wedge \neg e$, which is not related to the other formulas of both bases[6].

Finally, Hunter [2003] defines a degree of *significance* of the contradictions, which makes sense when some information about the importance of potential conflicts is available.

8 Conclusion

The main contribution of the paper is a uniform action-based framework for quantifying both degrees of information and of contradiction. The framework is parameterized by a propositional logic (together with the corresponding notions of consequence, acceptance, contradiction and a revision operator), a test context and an aggregation criterion for computing plan costs. These parameters enable a great flexibility.

There are many interesting notions that can be easily defined in our framework but that we cannot mention here for space reasons. Let us note that through the notion of purification plan, our approach for quantifying contradiction also allows to *localize* conflicts. Note also that notions of *joint degrees* and *conditional degrees* of information / contradiction can be easily defined. Another simple extension would consist in taking advantage of additional knowledge about the sources of information and the origin of conflicts (e.g., in a diagnosis setting, it can be the case that the failure of a component physically causes the failure of other components).

Many other extensions of our approach can be envisioned. For instance, coping with preferences over the goal variables (determining whether a holds is more important than determining whether b holds). Another possible extension concerns the case where ontic actions are available and the objective is to let the actual world as unchanged as possible (i.e., we can execute invasive actions but we prefer not to do it).

Acknowledgements

The third author has been partly supported by the IUT de Lens, the Université d'Artois, the Région Nord / Pas-de-Calais under the TACT-TIC project, and by the European Community FEDER Program.

[6] If we instantiate our framework with LP_m logic (under standard atomic context), we get $d_C(\Sigma) = d_C(\Sigma') = 1$, $d_I(\Sigma) = 5$, and $d_I(\Sigma') = 4$, showing that there is the same amount of contradiction in the two bases and that there is less information in Σ than in Σ'.

References

[Alchourrón *et al.*, 1985] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[Besnard and Hunter, 1995] P. Besnard and A. Hunter. Quasi-classical logic: Non-trivializable classical reasoning from inconsistent information. In *ECSQARU'95*, pages 44–51, 1995.

[Brewka, 1989] G. Brewka. Preferred subtheories: an extended logical framework for default reasoning. In *IJCAI'89*, pages 1043–1048, 1989.

[Ginsberg and Smith, 1988] M. Ginsberg and D. E. Smith. Reasoning about action I: A possible world approach. *Artificial Intelligence*, 35:165–195, 1988.

[Hintikka, 1970] J. Hintikka. On semantic information. *Information and Inference*, pages 3–27, 1970.

[Hunter, 2000] A. Hunter. Reasoning with conflicting information using quasi-classical logic. *Journal of Logic and Computation*, 10:677–703, 2000.

[Hunter, 2002] A. Hunter. Measuring inconsistency in knowledge via quasi-classical models. In *AAAI'02*, pages 68–73, 2002.

[Hunter, 2003] A. Hunter. Evaluating significance of inconsistencies. In *IJCAI'03 (this issue)*, 2003.

[Kemeny, 1953] J.G. Kemeny. A logical measure function. *Journal of Symbolic Logic*, 18:289–308, 1953.

[Knight, 2002] K.M. Knight. Measuring inconsistency. *Journal of Philosophical Logic*, 31:77–98, 2002.

[Knight, 2003] K.M. Knight. Two information measures for inconsistent sets. *Journal of Logic, Language and Information*, 12:227–248, 2003.

[Lang and Marquis, 1998] J. Lang and P. Marquis. Complexity results for independence and definability in propositional logic. In *KR'98*, pages 356–367, 1998.

[Lang and Marquis, 2002] J. Lang and P. Marquis. Resolving inconsistencies by variable forgetting. In *KR'02*, pages 239–250, 2002.

[Lozinskii, 1994a] E. Lozinskii. Information and evidence in logic systems. *Journal of Experimental and Theoretical Artificial Intelligence*, 6:163–193, 1994.

[Lozinskii, 1994b] E. Lozinskii. Resolving contradictions: a plausible semantics for inconsistent systems. *Journal of Automated Reasoning*, 12:1–31, 1994.

[Nebel, 1991] B. Nebel. Belief revision and default resoning: syntax-based approaches. In *KR'91*, pages 417–428, 1991.

[Poole, 1988] D. Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

[Priest, 1991] G. Priest. Minimally inconsistent LP. *Studia Logica*, 50:321–331, 1991.

[Rescher and Manor, 1970] N. Rescher and R. Manor. On inference from inconsistent premises. *Theory and Decision*, 1:179–219, 1970.

[Shannon, 1948] C.E. Shannon. A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423,623–656, 1948.

[Wong and Besnard, 2001] P. Wong and P. Besnard. Paraconsitent reasoning as an analytic tool. *Journal of the Interest Group in Propositional Logic*, 9:233–246, 2001.