Shapley Inconsistency Values

Anthony Hunter
Department of Computer Science
University College London, UK
a.hunter@cs.ucl.ac.uk

Sébastien Konieczny
CRIL - CNRS
Université d'Artois, Lens, France
konieczny@cril.univ-artois.fr

Abstract

There are relatively few proposals for inconsistency measures for propositional belief bases. However inconsistency measures are potentially as important as information measures for artificial intelligence, and more generally for computer science. In particular, they can be useful to define various operators for belief revision, belief merging, and negotiation. The measures that have been proposed so far can be split into two classes. The first class of measures takes into account the number of formulae required to produce an inconsistency: the more formulae required to produce an inconsistency, the less inconsistent the base. The second class takes into account the proportion of the language that is affected by the inconsistency: the more propositional variables affected, the more inconsistent the base. Both approaches are sensible, but there is no proposal for combining them. We address this need in this paper: our proposal takes into account both the number of variables affected by the inconsistency and the distribution of the inconsistency among the formulae of the base. Our idea is to use existing inconsistency measures (ones that takes into account the proportion of the language affected by the inconsistency, and so allow us to look inside the formulae) in order to define a game in coalitional form, and then to use the Shapley value to obtain an inconsistency measure that indicates the responsibility/contribution of each formula to the overall inconsistency in the base. This allows us to provide a more reliable image of the belief base and of the inconsistency in it.

Introduction

There are numerous works on reasoning under inconsistency. One can quote for example paraconsistent logics, argumentation frameworks, belief revision and fusion, etc. All these approaches illustrate the fact that the dichotomy between consistent and inconsistent sets of formulae that comes from classical logics is not sufficient for describing these sets. As shown by these works two inconsistent sets of formulae are not trivially equivalent. They do not contain the same information and they do not contain the same contradictions.

Measures of information *à la* Shannon have been studied in logical frameworks (see for example (Kemeny 1953)). Roughly they involve counting the number of models of the set of formulae (the less models, the more informative the set). The problem is that these measures give a null information content to an inconsistent set of formulae, which is counter-intuitive (especially given all the proposals for paraconsistent reasoning). So generalizations of measures of information have been proposed to solve this problem (Lozinskii 1994; Wong & Besnard 2001; Knight 2003; Konieczny, Lang, & Marquis 2003; Hunter & Konieczny 2005).

In comparison, there are relatively few proposals for inconsistency measures (Grant 1978; Hunter 2002; Knight 2001; Konieczny, Lang, & Marquis 2003; Hunter 2004; Grant & Hunter 2006). However, these measures are potentially important in diverse applications in artificial intelligence, such as belief revision, belief merging, and negotiation, and more generally in computer science. Already measuring inconsistency has being seen to be a useful tool in analysing a diverse range of information types including news reports (Hunter 2006), integrity constraints (Grant & Hunter 2006), software specifications (Barragáns-Martínez, Pazos-Arias, & Fernández-Vilas 2004; 2005; Mu *et al.* 2005), and ecommerce protocols (Chen, Zhang, & Zhang 2004).

The current proposals for measuring inconsistecy can be classified in two approaches. The first approach involves "counting" the minimal number of formulae needed to produce the inconsistency. The more formulae needed to produce the inconsistency, the less inconsistent the set (Knight 2001). This idea is an interesting one, but it rejects the possibility of a more fine-grained inspection of the (content of the) formulae. In particular, if one looks to singleton sets only, one is back to the initial problem, with only two values: consistent or inconsistent.

The second approach involves looking at the proportion of the language that is touched by the inconsistency. This allows us to look *inside* the formulae (Hunter 2002; Konieczny, Lang, & Marquis 2003; Grant & Hunter 2006). This means that two formulae (singleton sets) can have different inconsistency measures. But, in these approaches one can identify the set of formulae with its conjunction (i.e. the set $\{\varphi, \varphi'\}$ has the same inconsistency measure as the set $\{\varphi \wedge \varphi'\}$). This can be sensible in several applications, but this means that the distribution of the contradiction among the formulae is not taken into account.

What we propose in this paper is a definition for inconsistency measures that allow us to take the best of the two approaches. This will allow us to build inconsistency measures that are able to look inside the formulae, but also to take into account the distribution of the contradiction among the different formulae of the set. The advantage of such a method is twofold. First, this allows us to know the degree of blame/responsability of each formula of the base in the inconsistency, and so it provides a very detailed view of the inconsistency. Second, this allows us to define measures of consistency for the whole base that are more accurate, since they take into account those two dimensions.

To this end we will use a notion that comes from coalitional game theory: the Shapley value. This value assigns to each player the payoff that this player can expect from her utility for each possible coalition. The idea is to use existing inconsistency measures (that allow us to look inside the formulae) in order to define a game in coalitional form, and then to use the Shapley value to obtain an inconsistency measure with the wanted properties. We will study these measures and show that they are more interesting than the other existing measures.

After stating some notations and definitions in the next section, we introduce inconsistency measures that count the number of formulae needed to produce an inconsistency. Then we present the approaches where the inconsistency measure is related to the number of variables touched by the inconsistency. The next section gives the definition of coalitional games and of the Shapley value. Then we introduce the inconsistency measures based on the Shapley value. The penultimate section sketches the possible applications of those measures for belief change operators. In the last section we conclude and give perspectives of this work.

Preliminaries

We will consider a propositional language $\mathcal{L}$ built from a finite set of propositional symbols $\mathcal{P}$. We will use $a, b, c, \ldots$ to denote the propositional variables, and Greek letters $\alpha, \beta, \varphi, \ldots$ to denote the formulae. An interpretation is a total function from $\mathcal{P}$ to $\{0, 1\}$. The set of all interpretations is denoted $\mathcal{W}$. An interpretation ω is a model of a formula φ, denoted $\omega \models \varphi$, if and only if it makes φ true in the usual truth-functional way. $Mod(\varphi)$ denotes the set of models of the formula φ, i.e. $Mod(\varphi) = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. We will use $\subseteq$ to denote the set inclusion, and we will use $\subset$ to denote the strict set inclusion, i.e. $A \subset B$ iff $A \subseteq B$ and $B \nsubseteq A$. We will denote the set of real numbers by $I\!R$.

A *belief base* K is a finite set of propositional formulae. More exactly, as we will need to identify the different formulae of a belief base in order to associate them with their inconsistency value, we will consider belief bases K as vectors of formulae. For logical properties we will need to use the set corresponding to each vector, so we suppose that we have a function such that for each vector $K = (\alpha_1, \ldots, \alpha_n)$, $\overline{K}$ is the set $\{\alpha_1, \ldots, \alpha_n\}$. As it will never be ambigous, in the following we will omit the $\overline{}$ and write K as both the vector and the set.

Let us note $\mathcal{K}_\mathcal{L}$ the set of belief bases definable from formulae of the language $\mathcal{L}$. A belief base is consistent if there is at least one interpretation that satisfies all its formulae.

If a belief base K is not consistent, then one can define the minimal inconsistent subsets of K as:

$$\mathsf{MI}(K) = \{K' \subseteq K \mid K' \vdash \bot \text{ and } \forall K'' \subset K', K'' \nvdash \bot\}$$

If one wants to recover consistency from an inconsistent base K by removing some formulae, then the minimal inconsistent subsets can be considered as the purest form of inconsistency. To recover consistency, one has to remove at least one formula from each minimal inconsistent subset (Reiter 1987).

A *free formula* of a belief base K is a formula of K that does not belong to any minimal inconsistent subset of the belief base K, or equivalently any formula that belongs to any maximal consistent subset of the belief base.

Inconsistency Measures based on Formulae

When a base is not consistent the classical inference relation trivializes, since one can deduce every formula of the language from the base (*ex falso quodlibet*). Otherwise in this case the use of paraconsistent reasoning techniques allows us to draw non-trivial consequences from the base. One possibility is to take maximal consistent subsets of formulae of the base (cf (Manor & Rescher 1970; Benferhat, Dubois, & Prade 1997; Nebel 1991)). This idea can also be used to define an inconsistency measure. This is the way followed in (Knight 2001; 2003).

Definition 1 *A probability function on $\mathcal{L}$ is a function $P : \mathcal{P} \rightarrow [0, 1]$ s.t.:*

- *if $\models \alpha$, then $P(\alpha) = 1$*
- *if $\models \neg(\alpha \wedge \beta)$, then $P(\alpha \vee \beta) = P(\alpha) + P(\beta)$*

See (Paris 1994) for more details on this definition. In the finite case, this definition gives a probability distribution on the interpretations, and the probability of a formula is the sum of the probability of its models.

Then the inconsistency measure defined by Knight (2001) is given by:

Definition 2 *Let K be a belief base.*

- *K is $\eta-$consistent ($0 \leq \eta \leq 1$) if there is a probability function P such that $P(\alpha) \geq \eta$ for all $\alpha \in K$.*
- *K is maximally $\eta-$consistent if η is maximal (i.e. if $\gamma > \eta$ then K is not $\gamma-$consistent).*

The notion of *maximal η-consistency* can be used as an inconsistency measure. This is the direct formulation of the idea that the more formulae are needed to produce the inconsistency, the less this inconsistency is problematic. As it is easily seen, in the finite case, a belief base is maximally 0-consistent if and only if it contains a contradictory formula. And a belief base is maximally 1-consistent if and only if it is consistent.

Example 1 *Let $K_1 = \{a, b, \neg a \vee \neg b\}$.*
K_1 is maximally $\frac{2}{3}-$consistent.
Let $K_2 = \{a \wedge b, \neg a \wedge \neg b, a \wedge \neg b\}$.
K_2 is maximally $\frac{1}{3}-$consistent, whereas each subbase of cardinality 2 is maximally $\frac{1}{2}-$consistent.

For minimal inconsistent sets of formulae, computing this inconsistency measure is easy:

Proposition 1 *If $K' \in \mathsf{MI}(K)$, then K' is maximally $\frac{|K'|-1}{|K'|}$−consistent.*

But in general this measure is harder to compute. However it is possible to compute it using the simplex method (Knight 2001).

Inconsistency Measures based on Variables

Another method to evaluate the inconsistency of a belief base is to look at the proportion of the language concerned with the inconsistency. To this end, it is clearly not possible to use classical logics, since the inconsistency contaminates the whole language. But if we look at the two bases $K_1 = \{a \wedge \neg a \wedge b \wedge c \wedge d\}$ and $K_2 = \{a \wedge \neg a \wedge b \wedge \neg b \wedge c \wedge \neg c \wedge d \wedge \neg d\}$, we can observe that in K_1 the inconsistency is mainly about the variable a, whereas in K_2 all the variables are touched by a contradiction. This is this kind of distinction that these approaches allow.

One way to circumscribe the inconsistency only to the variables directly concerned is to use multi-valued logics, and especially three-valued logics, with the third "truth value" denoting the fact that there is a conflict on the truth value (true-false) of the variable.

We do not have space here to detail the range of different measures that have been proposed. See (Grant 1978; Hunter 2002; Konieczny, Lang, & Marquis 2003; Hunter & Konieczny 2005; Grant & Hunter 2006) for more details on these approaches. We only give one such measure, that is a special case of the degrees of contradiction defined in (Konieczny, Lang, & Marquis 2003). The idea of the definition of these degrees in (Konieczny, Lang, & Marquis 2003) is, given a set of tests on the truth value of some formulae of the language (typically on the variables), the degree of contradiction is the cost of a minimum test plan that ensures recovery of consistency.

The inconsistency measure we define here is the (normalized) minimum number of inconsistent truth values in the LP_m models (Priest 1991) of the belief base. Let us first introduce the LP_m consequence relation.

- 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". $3^{\mathcal{P}}$ 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 $\omega(\alpha) = \mathsf{F}$
  - $\omega(\alpha \wedge \beta) = \min_{\leq_t}(\omega(\alpha), \omega(\beta))$
  - $\omega(\alpha \vee \beta) = \max_{\leq_t}(\omega(\alpha), \omega(\beta))$
- The set of models of a formula φ is:

$$Mod_{LP}(\varphi) = \{\omega \in 3^{\mathcal{P}} \mid \omega(\varphi) \in \{\mathsf{T}, \mathsf{B}\}\}$$

Define $\omega!$ as the set of "inconsistent" variables in an interpretation w, i.e.

$$\omega! = \{x \in \mathcal{P} \mid \omega(x) = \mathsf{B}\}$$

Then the minimum models of a formula are the "most classical" ones:

$$min(Mod_{LP}(\varphi)) = \{\omega \in Mod_{LP}(\varphi) \mid \nexists \omega' \in Mod_{LP}(\varphi) \text{ s.t. } \omega'! \subset \omega!\}$$

The LP_m consequence relation is then defined by:

$$K \models_{LP_m} \varphi \text{ iff } min(Mod_{LP}(K)) \subseteq Mod_{LP}(\varphi)$$

So φ is a consequence of K if all the "most classical" models of K are models of φ.

Then let us define the LP_m measure of inconsistency, noted I_{LP_m}, as:

Definition 3 *Let K be a belief base.*

$$I_{LP_m} = \frac{\min_{\omega \in Mod_{LP}(K)}(\mid \omega! \mid)}{\mid \mathcal{P} \mid}$$

Example 2 $K_4 = \{a \wedge \neg a, b, \neg b, c\}$. $I_{LP_m}(K_4) = \frac{2}{3}$

In this example one can see the point in these kinds of measures compared to measures based on formulae since this base is maximally 0-consistent because of the contradictory formula $a \wedge \neg a$. But there are also non-trivial formulae in the base, and this base is not very inconsistent according to I_{LP_m}.

Conversely, measures based on variables like this one are unable to take into account the distribution of the contradiction among formulae. In fact the result would be exactly the same with $K'_4 = \{a \wedge \neg a \wedge b \wedge \neg b \wedge c\}$. This can be sensible in several applications, but in some cases this can also be seen as a drawback.

Games in Coalitional Form - Shapley Value

In this section we give the definitions of games in coalitional form and of the Shapley value.

Definition 4 *Let $N = \{1, \dots, n\}$ be a set of n players. A game in coalitional form is given by a function $v : 2^N \rightarrow \mathbb{R}$, with $v(\emptyset) = 0$.*

This framework defines games in a very abstract way, focusing on the possible coalitions formations. A coalition is just a subset of N. This function gives what payoff can be achieved by each coalition in the game v when all its members act together as a unit.

There are numerous questions that are worthwhile to investigate in this framework. One of these questions is to know how much each player can expect in a given game v. This depend on her position in the game, i.e. what she brings to different coalitions.

Often the games are super-additive.

Definition 5 *A game is* super-additive *if for each $T, U \subseteq N$ with $T \cap U = \emptyset$, $v(T \cup U) \geq v(T) + v(U)$.*

In super-additive games when two coalitions join, then the joined coalition wins at least as much as (the sum of) the initials coalitions. In particular, in super-additive games, the grand coalition N is the one that brings the higher utility for the society N. The problem is how this utility can be shared among the players[1].

Example 3 *Let $N = \{1, 2, 3\}$, and let v be the following coalitional game:*

$$\begin{array}{ccc} v(\{1\}) = 1 & v(\{2\}) = 0 & v(\{3\}) = 1 \\ v(\{1,2\}) = 10 & v(\{1,3\}) = 4 & v(\{2,3\}) = 11 \\ & v(\{1,2,3\}) = 12 & \end{array}$$

This game is clearly super-additive. The grand coalition can bring 12 to the three players. This is the highest utility achievable by the group. But this is not the main aim for all the players. In particular one can note that two coalitions can bring nearly as much, namely $\{1, 2\}$ and $\{2, 3\}$ that gives respectively 10 and 11, that will have to be shared only between 2 players. So it is far from certain that the grand coalition will form in this case. Another remark on this game is that all the players do not share the same situation. In particular player 2 is always of a great value for any coalition she joins. So she seems to be able to expect more from this game than the other players. For example she can make an offer to player 3 for making the coalition $\{2, 3\}$, that brings 11, that will be split in 8 for player 2 and 3 for player 3. As it will be hard for player 3 to win more than that, 3 will certainly accept.

A solution concept has to take into account these kinds of arguments. It means that one wants to *solve* this game by stating what is the payoff that is "due" to each agent. That requires to be able to quantify the payoff that an agent can claim with respect to the power that her position in the game offers (for example if she always significantly improves the payoff of the coalitions she joins, if she can threat to form another coalition, etc.).

Definition 6 *A value is a function that assigns to each game v a vector of payoff $S(v) = (S_1, \ldots, S_n)$ in $\mathbb{R}^n$.*

This function gives the payoff that can be expected by each player i for the game v, i.e. it measures i's power in the game v.

Shapley proposes a beautiful solution to this problem. Basically the idea can be explained as follows: considering that the coalitions form according to some order (a first player enters the coalition, then another one, then a third one, etc), and that the payoff attached to a player is its marginal utility (i.e. the utility that it brings to the existing coalition), so if C is a coalition (subset of N) not containing i, player's i marginal utility is $v(C \cup \{i\}) - v(C)$. As one can not make any hypothesis on which order is the correct one, suppose that each order is equally probable. This leads to the following formula:

Let σ be a permutation on N, with σ_n denoting all the possible permutations on N. Let us note

$$p^i_\sigma = \{j \in N \mid \sigma(j) < \sigma(i)\}$$

That means that p^i_σ represents all the players that precede player i for a given order σ.

[1] One supposes the transferable utility (TU) assumption, i.e. the utility is a common unit between the players and sharable as needed (roughly, one can see this utility as some kind of money).

Definition 7 *The Shapley value of a game v is defined as:*

$$S_i(v) = \frac{1}{n!} \sum_{\sigma \in \sigma_n} v(p^i_\sigma \cup \{i\}) - v(p^i_\sigma)$$

The Shapley value can be directly computed from the possible coalitions (without looking at the permutations), with the following expression:

$$S_i(v) = \sum_{C \subseteq N} \frac{(c-1)!(n-c)!}{n!}(v(C) - v(C \setminus \{i\}))$$

where c is the cardinality of C.

Example 4 *The Shapley value of the game defined in Example 3 is $(\frac{17}{6}, \frac{35}{6}, \frac{20}{6})$.*

These values show that it is player 2 that is the best placed in this game, accordingly to what we explained when we presented Example 3.

Besides this value, Shapley proposes axiomatic properties a value should have.

- $\sum_{i \in N} S_i(v) = v(N)$ **(Efficiency)**
- If i and j are such that for all C s.t. $i, j \notin C$, $v(C \cup \{i\}) = v(C \cup \{j\})$, then $S_i(v) = S_j(v)$ **(Symmetry)**
- If i is such that $\forall C \ v(C \cup \{i\}) = v(C)$, then $S_i(v) = 0$ **(Dummy)**
- $S_i(v + w) = S_i(v) + S_i(w)$ **(Additivity)**

These four axioms seem quite sensible. Efficiency states that the payoff available to the grand coalition N must be efficiently redistributed to the players (otherwise some players could expect more that what they have). Symmetry ensures that it is the role of the player in the game in coalitional form that determines her payoff, so it is not possible to distinguish players by their name (as far as payoffs are concerned), but only by their respective merits/possibilities. So if two players always are identical for the game, i.e. if they bring the same utility to every coalitions, then they have the same value. The dummy player axiom says simply that if a player is of no use for every coalition, this player does not deserve any payoff. And additivity states that when we join two different games v and w in a whole super-game $v + w$ ($v + w$ is straightforwardly defined as the function that is the sum of the two functions v and w, that means that each coalition receive as payoff in the game $v + w$ the payoff it has in v plus the payoff it has in w), then the value of each player in the supergame is simply the sum of the values in the compound games.

These properties look quite natural, and the nice result shown by Shapley is that they characterize exactly the value he defined (Shapley 1953):

Proposition 2 *The Shapley value is the only value that satisfies all of Efficiency, Symmetry, Dummy and Additivity.*

This result supports several variations : there are other equivalent axiomatizations of the Shapley value, and there are some different values that can be defined by relaxing some of the above axioms. See (Aumann & Hart 2002).

Inconsistency Values using Shapley Value

Given an inconsistency measure, the idea is to take it as the payoff function defining a game in coalitional form, and then using the Shapley value to compute the part of the inconsistency that can be imputed to each formula of the belief base.

This allows us to combine the power of inconsistency measures based on variables and hence discriminating between singleton inconsistent belief base (like Coherence measure in (Hunter 2002), or like the test action values of (Konieczny, Lang, & Marquis 2003)), and the use of the Shapley value for knowing what is the responsibility of a given formula in the inconsistency of the belief base.

We just require some basic properties on the underlying inconsistency measure.

Definition 8 *An inconsistency measure I is called a basic inconsistency measure if it satisfies the following properties, $\forall K, K' \in \mathcal{K}_{\mathcal{L}}$, $\forall \alpha, \beta \in \mathcal{L}$:*

- *$I(K) = 0$ iff K is consistent* **(Consistency)**
- *$0 \leq I(K) \leq 1$* **(Normalization)**
- *$I(K \cup K') \geq I(K)$* **(Monotony)**
- *If α is a free formula of $K \cup \{\alpha\}$, then $I(K \cup \{\alpha\}) = I(K)$* **(Free Formula Independence)**
- *If $\alpha \vdash \beta$ and $\alpha \nvdash \bot$, then $I(K \cup \{\alpha\}) \geq I(K \cup \{\beta\})$* **(Dominance)**

We ask for few properties on the initial inconsistency measure. The consistency property states that a consistent base has a null inconsistency measure. The monotony property says that the amount of inconsistency of a belief base can only grow if one adds new formulae (defined on the same language). The free formula independence property states that adding a formula that does not cause any inconsistency cannot change the inconsistency measure of the base. The Dominance property states that logically stronger formulae bring (potentially) more conflicts. The normalization property of the inconsistency measure is not mandatory, it is asked only for simplification purposes.

Now we are able to define the Shapley inconsistency values :

Definition 9 *Let I be a basic inconsistency measure. We define the corresponding Shapley inconsistency value (SIV), noted S_I, as the Shapley value of the coalitional game defined by the function I, i.e. let $\alpha \in K$:*

$$S_I^K(\alpha) = \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!} (I(C) - I(C \setminus \{\alpha\}))$$

where n is the cardinality of K and c is the cardinality of C.

Note that this SIV gives a value for each formula of the base K, so if one considers the base K as the vector $K = (\alpha_1, \ldots, \alpha_n)$, then we will use $S_I(K)$ to denote the vector of corresponding SIVs, i.e.

$$S_I(K) = (S_I^K(\alpha_1), \ldots, S_I^K(\alpha_n))$$

This definition allows us to define to what extent a formula inside a belief base is concerned with the inconsistencies of the base. It allows us to draw a precise picture of the contradiction of the base.

From this value, one can define an inconsistency value for the whole belief base:

Definition 10 *Let K be a belief base, $\hat{S}_I(K) = \max_{\alpha \in K} S_I(\alpha)$*

One can figure out other aggregation functions to define the inconsistency measure of the belief base from the inconsistency measure of its formulae, such as the leximax for instance. Taking the maximum will be sufficient for us to have valuable results and to compare this with the existing measures from the literature. Note that taking the sum as aggregation function is not a good choice here, since as shown by the distribution property of Theorem 3 this equals $I(K)$, "erasing" the use of the Shapley value.

We think that the most interesting measure is S_I, since it describes more accurately the inconsistency of the base. But we define $\hat{S}_I$ since it is a more concise measure, that is of the same type as existing ones (it associates a real to each base), that is convenient to compare our framework with existing measures.

Let us see see now two instantiations of SIVs.

Drastic Shapley Inconsistency Value

We will start this section with the simplest inconsistency measure one can define:

Definition 11 *The drastic inconsistency value is defined as:*

$$I_d(K) = \begin{cases} 0 & \text{if } K \text{ is consistent} \\ 1 & \text{otherwise} \end{cases}$$

This measure is not of great interest by itself, since it corresponds to the usual dichotomy of classical logic. But it will be useful to illustrate the use of the Shapley inconsistency values, since, even with this over-simple measure, one will produce interesting results. Let us illustrate this on some examples.

Example 5 $K_1 = \{a, \neg a, b\}$.
Then $I_d(\{a, \neg a\}) = I_d(\{a, \neg a, b\}) = 1$, and the value is $S_{I_d}(K_1) = (\frac{1}{2}, \frac{1}{2}, 0)$. So $\hat{S}_{I_d}(K_1) = \frac{1}{2}$.

As b is a free formula, it has a value of 0, the two other formulae are equally responsible for the inconsistency.

Example 6 $K_2 = \{a, b, b \wedge c, \neg b \wedge d\}$.
Then the value is $S_{I_d}(K_2) = (0, \frac{1}{6}, \frac{1}{6}, \frac{4}{6})$.
And $\hat{S}_{I_d}(K_2) = \frac{2}{3}$.

The last three formulae are the ones that belong to some inconsistency, and the last one is the one that causes the most problems (removing only this formula restores the consistency of the base).

Example 7 $K_4 = \{a \wedge \neg a, b, \neg b, c\}$.
The value is $S_{I_d}(K_4) = (\frac{4}{6}, \frac{1}{6}, \frac{1}{6}, 0)$. So $\hat{S}_{I_d}(K_4) = \frac{2}{3}$.

$\mathbf{LP_m}$ Shapley Inconsistency Value

Let us turn now to a more elaborate value. For this we use the LP_m inconsistency measure (defined earlier) to define a SIV.

Example 8 *Let* $K_4 = \{a \wedge \neg a, b, \neg b, c\}$ *and* $K_4' = \{a \wedge \neg a \wedge b \wedge \neg b \wedge c\}$.
Then $S_{I_{LP_m}}(K_4) = (\frac{1}{3}, \frac{1}{6}, \frac{1}{6}, 0)$, *and* $\hat{S}_{I_{LP_m}}(K_4) = \frac{1}{3}$.
Whereas $S_{I_{LP_m}}(K_4') = (\frac{2}{3})$ *and* $\hat{S}_{I_{LP_m}}(K_4') = \frac{2}{3}$.

As we can see on this example, the SIV value allows us to make a distinction between K_4 and K_4', since $\hat{S}_{I_{LP_m}}(K_4') = \frac{2}{3}$ whereas $\hat{S}_{I_{LP_m}}(K_4) = \frac{1}{3}$. This illustrates the fact that the inconsistency is more distributed in K_4 than in K_4'. This distinction is not possible with the original I_{LP_m} value. Note that with Knight's coherence value, the two bases have the worst inconsistency value (maximally 0-consistent).

So this example illustrates the improvement brought by this work, compared to inconsistency measures on formulae and to inconsistency measures on variables, since none of them was able to make a distinction between K_4 and K_4', whereas for $\hat{S}_{I_{LP_m}}$ K_4 is more consistent than K_4'.

Let us see a more striking example.

Example 9 *Let* $K_5 = \{a, b, b \wedge c, \neg b \wedge \neg c\}$.
Then $S_{I_{LP_m}}(K_5) = (0, \frac{1}{18}, \frac{4}{18}, \frac{7}{18})$,
and $\hat{S}_{I_{LP_m}}(K_5) = \frac{7}{18}$.

In this example one can easily see that it is the last formula that is the more problematic, and that $b \wedge c$ brings more conflict than b alone, which is perfectly expressed in the obtained values.

Logical properties

Let us see now some properties of the defined values.

Proposition 3 *Every Shapley Inconsistency Value satisfies:*

- $\sum_{\alpha \in K} S_I(\alpha) = I(K)$ **(Distribution)**
- *If* $\exists \alpha, \beta \in K$ *s.t. for all* $K' \subseteq K$ *s.t.* $\alpha, \beta \notin K'$, $I(K' \cup \{\alpha\}) = I(K' \cup \{\beta\})$, *then* $S_I(\alpha) = S_I(\beta)$ **(Symmetry)**
- *If* α *is a free formula of* K, *then* $S_I(\alpha) = 0$ **(Free Formula)**
- *If* $\alpha \vdash \beta$ *and* $\alpha \nvdash \bot$, *then* $S_I(\alpha) \geq S_I(\beta)$ **(Dominance)**

The distribution property states that the inconsistency values of the formulae sum to the total amount of inconsistency in the base ($I(K)$). The Symmetry property ensures that only the amount of inconsistency brought by a formula matters for computing the SIV. As one could expect, a formula that is not embedded in any contradiction (i.e. does not belong to any minimal inconsistent subset) will not be blamed by the Shapley inconsistency values. This is what is expressed in the Free formula property. The Dominance property states that logically stronger formulae bring (potentially) more conflicts.

The first three properties are a restatement in this logical framework of the properties of the Shapley value. One can note that the Additivity axiom of the Shapley value is not translated here, since it makes little sense to add different inconsistency values.

Let us turn now to the properties of the measure on belief bases.

Proposition 4

- $\hat{S}_I(K) = 0$ *if and only if* K *is consistent* **(Consistency)**
- $0 \leq \hat{S}_I(K) \leq 1$ **(Normalization)**
- *If* α *is a free formula of* $K \cup \{\alpha\}$, *then* $\hat{S}_I(K \cup \{\alpha\}) = \hat{S}_I(K)$ **(Free Formula Independence)**
- $\hat{S}_I(K) \leq I(K)$ **(Upper Bound)**
- $\hat{S}_I(K) = I(K) > 0$ *if only if* $\exists \alpha \in K$ *s.t.* α *is inconsistent and* $\forall \beta \in K$, $\beta \neq \alpha$, β *is a free formula of* K **(Isolation)**

The first three properties are the ones given in Definition 8 for the basic inconsistency measures. As one can easily note an important difference is that the monotony property and the dominance property do not hold for the SIVs on belief bases. It is sensible since distribution of the inconsistencies matters for SIVs. The upper bound property shows that the use of the SIV aims at looking at the distribution of the inconsistencies of the base, so the SIV on belief bases is always less or equal to the inconsistency measure given by the underlying basic inconsistency measure. The isolation property details the case where the two measures are equals. In this case, there is only one inconsistent formula in the whole base.

Let us see, on Example 10, counter-examples to monotony and dominance for SIV on belief bases:

Example 10 *Let* $K_6 = \{a, \neg a, \neg a \wedge b\}$, $K_7 = \{a, \neg a, \neg a \wedge b, a \wedge b\}$, *and* $K_8 = \{a, \neg a, \neg a \wedge b, b\}$.
$\hat{S}_{I_d}(K_6) = \frac{2}{3}$, $\hat{S}_{I_d}(K_7) = \frac{1}{4}$, $\hat{S}_{I_d}(K_8) = \frac{2}{3}$.

On this example one can see why monotony can not be satisfied by SIV on belief bases. Clearly $K_6 \subset K_7$, but $\hat{S}_{I_d}(K_6) > \hat{S}_{I_d}(K_7)$. This is explained by the fact that the inconsistency is more *diluted* in K_7, than in K_8. In K_7 the formula a is the one that is the most blamed for the inconsistency ($S^{K_6}_{I_d}(a) = \hat{S}_{I_d}(K_6) = \frac{2}{3}$), since it appears in all inconsistent sets. Whereas in K_7 inconsistencies are equally caused by a and by $a \wedge b$, that decreases the responsability of a, and the whole inconsistency value of the base.

For a similar reason dominance is not satisfied, we clearly have $a \wedge b \vdash b$ (and $a \wedge b \nvdash \bot$), but $\hat{S}_{I_d}(K_7) < \hat{S}_{I_d}(K_8)$.

Applications for Belief Change Operators

As the measures we define allow us to associate with each formula its degree of responsibility for the inconsistency of the base, they can be used to guide any paraconsistent reasoning, or any repair of the base. Let us quote two such possible uses for belief change operators, first for belief revision and then for negotiation.

Iterated Revision and Transmutation Policies

The problem of belief revision is to incorporate a new piece of information which is more reliable than (and conflicting with) the old beliefs of the agent. This problem has received a nice answer in the work of Alchourron, Gardenfors, Makinson (Alchourrón, Gärdenfors, & Makinson 1985) in the one-step case. But when one wants to iterate revision (i.e. to generalize it to the n-steps case), there are numerous problems and no definitive answer has been reached in the purely qualitative case (Darwiche & Pearl 1997; Friedman & Halpern 1996). Using a partially quantitative framework, some proposals have given interesting results (see e.g. (Williams 1995; Spohn 1987)). Here "partially quantitative" means that the incoming piece of information needs to be labeled by a degree of confidence denoting how strongly we believe it. The problem in this framework is to justify the use of such a degree, what does it mean exactly and where does it come from. One possibility is to use an inconsistency measure (or a composite measure computed from an information measure (Lozinskii 1994; Knight 2003; Konieczny, Lang, & Marquis 2003) and an inconsistency measure) to determine this degree of confidence. Then one can define several policies for the agent (we can suppose that an agent accepts a new piece of information only if it brings more information than contradiction, etc). We can then use the partially quantitative framework to derive revision operators with a nice behaviour. In this setting, since the degree attached to the incoming information is not a given data, but computed directly from the information itself and the agent policy (behaviour with respect to information and contradiction, encoded by a composite measure) then the problem of the justification of the meaning of the degrees is avoided.

Negotiation

The problem of negotiation has been investigated recently under the scope of belief change tools (Booth 2001; 2002; 2006; Zhang *et al.* 2004; Meyer *et al.* 2004; Konieczny 2004; Gauwin, Konieczny, & Marquis 2005). The problem is to define operators that take as input belief profiles (multi-set of formulae[2]) and that produce a new belief profile that aims to be less conflicting. We call these kind of operators conciliation operators. The idea followed in (Booth 2002; 2006; Konieczny 2004) to define conciliation operators is to use an iterative process where at each step a set of formulae is selected. These selected formulae are logically weakened. The process stops when one reaches a consensus, i.e. a consistent belief profile[3]. Many interesting operators can be defined when one fixes the selection function (the function that selects the formulae that must be weaken at each round) and the weakening method. In (Konieczny 2004) the selection function is based on a notion of distance. It can be sensible if such a distance is meaningful in a particular application. If not, it is only an arbitrary choice. It would then be sensible to choose instead one of the inconsistency measures we defined in this paper. So the selection function would choose the formulae with the highest inconsistency value. These formulae are clearly the more problematic ones. More generally SIVs can be used to define new belief merging methods.

[2]More exactly belief profiles are sets of belief bases. We use this simplifying assumption just for avoiding technical details here.

[3]A belief profile is consistent if the conjunction of its formulae is consistent.

Conclusion

We have proposed in this paper a new framework for defining inconsistency values. The SIV values we introduce allow us to take into account the distribution of the inconsistency among the formulae of the belief base and the variables of the language. This is, as far as we know, the only definition that allows us to take both types of information into account, thus allowing to have a more precise picture of the inconsistency of a belief base. The perspectives of this work are numerous. First, as sketched in the previous section, the use of inconsistency measures, and especially the use of Shapley inconsistency values, can be valuable for several belief change operators, for instance for modelizations of negotiation. The Shapley value is not the only solution concept for coalitional games, so an interesting question is to know if other solutions concept can be sensible as a basis for defining other inconsistency measures. But the main way of research opened by this work is to study more closely the connections between other notions of (cooperative) game theory and the logical modelization of belief change operators.

Acknowledgments

The authors would like to thank CNRS and the Royal Society for travel funding while collaborating on this research. The second author is supported by the Région Nord/Pas-de-Calais and the European Community FEDER program.

Proofs

Proof of Proposition 3 : To show distribution, let us recall that

$$\begin{aligned} S_I^K(\alpha) &= \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!}(I(C) - I(C \setminus \{\alpha\})) \\ &= \frac{1}{n!} \sum_{\sigma \in \sigma_n} I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha) \end{aligned}$$

where σ_n is the set of possible permutations on $\overline{K}$, and $p_\sigma^\alpha = \{\beta \in K \mid \sigma(\beta) < \sigma(\alpha)\}$. Now

$$\begin{aligned} \sum_{\alpha \in K} S_I(\alpha) &= \sum_{\alpha \in K} \frac{1}{n!} \sum_{\sigma \in \sigma_n} I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha) \\ &= \frac{1}{n!} \sum_{\sigma \in \sigma_n} \sum_{\alpha \in K} I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha) \end{aligned}$$

Now note that we can order the elements of K accordingly to σ when computing the inside sum, that gives:

$$\begin{aligned} &= \frac{1}{n!} \sum_{\sigma \in \sigma_n} [I(\{\alpha_{\sigma(1)}, \ldots, \alpha_{\sigma(n)}\}) \\ &\quad - I(\{\alpha_{\sigma(1)}, \ldots, \alpha_{\sigma(n-1)}\})] \\ &\quad + [I(\{\alpha_{\sigma(1)}, \ldots, \alpha_{\sigma(n-1)}\}) \\ &\quad - I(\{\alpha_{\sigma(1)}, \ldots, \alpha_{\sigma(n-2)}\})] \\ &\quad + \ldots + [I(\{\alpha_{\sigma(1)}\}) - I(\emptyset)] \\ &= \frac{1}{n!} \sum_{\sigma \in \sigma_n} I(\{\alpha_{\sigma(1)}, \ldots, \alpha_{\sigma(n)}\}) - I(\emptyset) \\ &= \frac{1}{n!}\, n!\, I(K) \\ &= I(K) \end{aligned}$$

To show symmetry, assume that there are $\alpha, \beta \in K$ s.t. for all $K' \subseteq K$ s.t. $\alpha, \beta \notin K'$, $I(K' \cup \{\alpha\}) = I(K' \cup \{\beta\})$.

Now by definition

$$S_I^K(\alpha) = \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!}(I(C) - I(C \setminus \{\alpha\}))$$

Let us show that $S_I^K(\alpha) = S_I^K(\beta)$ by showing (by cases) that the elements of the sum are the same:

If $\alpha \notin C$ and $\beta \notin C$, then $I(C) = I(C\setminus\{\alpha\}) = I(C\setminus\{\beta\})$, so $I(C) - I(C \setminus \{\alpha\}) = I(C) - I(C \setminus \{\beta\})$.

If $\alpha \in C$ and $\beta \in C$, then note that by hypothesis, as $\alpha, \beta \notin C \setminus \{\alpha, \beta\}$, we deduce that $I(C \setminus \{\alpha\}) = I(C \setminus \{\beta\})$. So $I(C) - I(C \setminus \{\alpha\}) = I(C) - I(C \setminus \{\beta\})$.

If $\alpha \in C$ and $\beta \notin C$. Then $I(C) - I(C\setminus\{\beta\}) = 0$, and let us note $I(C) - I(C\setminus\{\alpha\}) = a$. Let us note $C = C' \cup \{\alpha\}$, and $C'' = C' \cup \{\beta\}$. Now notice that $I(C'') - I(C'' \setminus \{\alpha\}) = 0$, and as we can deduce $I(C \setminus \{\alpha\}) = I(C'' \setminus \{\beta\})$ by the hypothesis, we also have $I(C'') - I(C'' \setminus \{\beta\}) = a$.

To show the free formula property, just note that if α is a free formula of K, then for every subset C of K, by the free formula independence property of the basic inconsistency measure we have that for every C, such that $\alpha \in C$, $I(C) = I(C \setminus \alpha)$, so $I(C) - I(C \setminus \alpha) = 0$. Straightforwardly if $\alpha \notin C$, $I(C) = I(C \setminus \alpha)$. So the whole expression $S_I^K(\alpha) = \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!}(I(C) - I(C \setminus \{\alpha\}))$ sums to 0.

Finally, to show dominance we will proceed in a similar way than to show symmetry. Assume that $\alpha, \beta \in K$ are such that $\alpha \vdash \beta$ and $\alpha \nvdash \perp$. Then, by the dominance property of the underlying basic inconsistency measure, we know that for all $C \subseteq K$, $I(C \cup \{\alpha\}) \geq I(C \cup \{\beta\})$. Now by definition of the SIV $S_I^K(\alpha) = \sum_{C \subseteq K} \frac{(c-1)!(n-c)!}{n!}(I(C) - I(C \setminus \{\alpha\}))$. Let us show that $S_I^K(\alpha) \geq S_I^K(\beta)$ by showing (by cases) that the elements of the first sum are greater or equal to the corresponding elements of the second one:

If $\alpha \notin C$ and $\beta \notin C$, then $I(C) = I(C\setminus\{\alpha\}) = I(C\setminus\{\beta\})$, so $I(C) - I(C \setminus \{\alpha\}) \geq I(C) - I(C \setminus \{\beta\})$.

If $\alpha \in C$ and $\beta \in C$, then let us note $C \setminus \{\alpha\} = C' \cup \{\beta\}$. So we also have $C \setminus \{\beta\} = C' \cup \{\alpha\}$. Now note that by hypothesis $I(C' \cup \{\beta\}) \leq I(C' \cup \{\alpha\})$, so $I(C \setminus \{\alpha\}) \leq I(C\setminus\{\beta\})$. Hence $I(C) - I(C\setminus\{\alpha\}) \geq I(C) - I(C\setminus\{\beta\})$.

If $\alpha \in C$ and $\beta \notin C$. Then $I(C) - I(C\setminus\{\beta\}) = 0$, and let us note $I(C) - I(C\setminus\{\alpha\}) = a$. Let us note $C = C' \cup \{\alpha\}$, and $C'' = C' \cup \{\beta\}$. Now notice that $I(C'') - I(C'' \setminus \{\alpha\}) = 0$. So $I(C) - I(C \setminus \{\beta\}) \geq I(C'') - I(C'' \setminus \{\alpha\})$ Note that $I(C)\setminus\{\beta\} = I(C'')\setminus\{\alpha\} = C'$. As we can deduce $I(C) \geq I(C'')$ by the hypothesis, we also have $I(C) - I(C\setminus\{\alpha\}) \geq I(C'') - I(C'' \setminus \{\beta\})$.

□

Proof of Proposition 4 : To prove consistency note that if K is consistent, then for every $C \subseteq K$, $I(C) = 0$ (this is a direct consequence of the consistency property of the underlying basic inconsistency measure). Then for every $\alpha \in K$, $S_I^K(\alpha) = 0$. Hence $\hat{S}_I(K) = \max_{\alpha \in K} S_I(\alpha) = 0$. For the only if direction, by contradiction, suppose that $\hat{S}_I(K) = 0$ and that K is not consistent. As K is not consistent, then by the consistency property of the underlying basic inconsistency measure $I(K) = a \neq 0$. By the distribution property of the SIV we know that $\sum_{\alpha \in K} S_I(\alpha) = a \neq 0$, then $\exists \alpha \in K$ such that $S_I(\alpha) > 0$, so $\hat{S}_I(K) = \max_{\alpha \in K} S_I(\alpha) > 0$. Contradiction.

The normalization property is a consequence of the definition of $\hat{S}_I(K)$ as a maximum of values that are all greater than zero, that ensures $0 \leq \hat{S}_I(K)$, and that are all smaller than 1. An easy way to show $\hat{S}_I(K) \leq 1$ is as a consequence of the upper bound property (shown below) $\hat{S}_I(K) \leq I(K)$ and of $I(K) \leq 1$ obtained by the normalization property of the underlying basic inconsistency measure I.

To show the free formula independence property, just notice that for any formula β that is a free formula of $K \cup \{\beta\}$, it is also a free formula of every of its subsets. It is easy to see from the definition that for any $\alpha \in K$, $S_I^K(\alpha) = S_I^{K\cup\{\beta\}}(\alpha)$. This is easier if we consider the second form of the definition: $S_I^K(\alpha) = \frac{1}{n!}\sum_{\sigma \in \sigma_n} I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha)$ where σ_n is the set of possible permutations on $\overline{K}$. Now note that for $S_I^{K\cup\{\beta\}}(\alpha)$, the free formula does not bring any contradiction, so it does not change the marginal contribution of every other formulae. Let us call the extensions of a permutation σ on K by β, all the permutations of $K \cup \{\beta\}$ whose restriction on elements of K is identical to σ, i.e. an extension of $\sigma = (\alpha_1, \ldots, \alpha_n)$ by β is a permutation $\sigma' = (\alpha_1, \ldots, \alpha_i, \beta, \alpha_{i+1}, \ldots, \alpha_n)$. Now note that there are $n+1$ such extensions, and that if σ' is an extension of sigma, $I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha) = I(p_{\sigma'}^\alpha \cup \{\alpha\}) - I(p_{\sigma'}^\alpha)$. So $S_I^{K\cup\{\beta\}}(\alpha) = \frac{1}{(n+1)!}(n+1)\sum_{\sigma \in \sigma_n} I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha)$ $= \frac{1}{n!}\sum_{\sigma \in \sigma_n} I(p_\sigma^\alpha \cup \{\alpha\}) - I(p_\sigma^\alpha) = S_I^K(\alpha)$. Now as we have for any $\alpha \in K$, $S_I^K(\alpha) = S_I^{K\cup\{\beta\}}(\alpha)$, we have $\hat{S}_I(K \cup \{\alpha\}) = \hat{S}_I(K)$.

The upper bound property is stated by rewriting $I(K)$ as $\sum_{\alpha \in K} S_I(\alpha)$ with the distribution property of the SIV, and by recalling the definition of $\hat{S}_I(K)$ as $\max_{\alpha \in K} S_I(\alpha)$. Now by noticing that for every vector $a = (a_1, \ldots, a_n)$, $\max_{a_i \in a} a_i \leq \sum_{a_i \in a} a_i$, we conclude $\max_{\alpha \in K} S_I(\alpha) \leq \sum_{\alpha \in K} S_I(\alpha)$, i.e. $\hat{S}_I(K) \leq I(K)$.

Let us show isolation. The if direction is straightforward: As α is inconsistent, K is inconsistent, and by the consistency and normalization properties of the underlying basic inconsistency measure we know that $I(K) > 0$. By the free formula property of SIV, for every free formula β of K we have $S_I(\beta) = 0$. As by the distribution property we have $\sum_{\alpha \in K} S_I(\alpha) = I(K)$, this means that $S_I(\alpha) = I(K)$, and that $\hat{S}_I(K) = \max_{\alpha \in K} S_I(\alpha) = S_I(\alpha)$. So $\hat{S}_I(K) = I(K) > 0$. For the only if direction suppose that $\hat{S}_I(K) = I(K)$, that means that $\max_{\alpha \in K} S_I(\alpha) = I(K)$. But, by the distribution property we know that $I(K) = \sum_{\alpha \in K} S_I(\alpha)$. So it means that $\max_{\alpha \in K} S_I(\alpha) = \sum_{\alpha \in K} S_I(\alpha) = I(K)$. There exists α such that $S_I(\alpha) = I(K)$ (consequence of the definition of the max), and if there exists a $\beta \neq \alpha$ such that

$S_I(\beta) > 0$, then $\sum_{\alpha \in K} S_I(\alpha) > I(K)$. Contradiction. So it means that there is α such that $S_I(\alpha) = I(K)$ and for every $\beta \neq \alpha$, $S_I(\beta) = 0$. That means that every β is a free formula, and that α is inconsistent.

□

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.

Aumann, R., and Hart, S., eds. 2002. *Handbook of Game Theory with Economic Applications*, volume 3. North-Holland.

Barragáns-Martínez, B.; Pazos-Arias, J. J.; and Fernández-Vilas, A. 2004. On Measuring Levels of Inconsistency in Multi-Perspective Requirements Specifications. In *Proceedings of the 1st Conference on the Principles of Software Engineering (PRISE 2004)*, 21–30.

Barragáns-Martínez, B.; Pazos-Arias, J. J.; and Fernández-Vilas, A. 2005. Merging Requirements Views with Incompleteness and Inconsistency. In *Proceedings of the IEEE Australian Software Engineering Conference (ASWEC'05)*. Brisbane, Australia: IEEE Computer Society Press.

Benferhat, S.; Dubois, D.; and Prade, H. 1997. Some syntactic approaches to the handling of inconsistent knowledge bases: a comparative study, part 1: the flat case. *Studia Logica* 58:17–45.

Booth, R. 2001. A negotiation-style framework for non-prioritised revision. In *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge (TARK'01)*, 137–150.

Booth, R. 2002. Social contraction and belief negotiation. In *Proceedings of the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, 374–384.

Booth, R. 2006. Social contraction and belief negotiation. *Information Fusion* 7(1):19–34.

Chen, Q.; Zhang, C.; and Zhang, S. 2004. A verification model for electronic transaction protocols. In *Advanced Web Technologies and Applications*, volume 3007 of *LNCS*, 824–833. Springer.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29.

Friedman, N., and Halpern, J. Y. 1996. Belief revision: a critique. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR'96)*, 421–431.

Gauwin, O.; Konieczny, S.; and Marquis, P. 2005. Conciliation and consensus in iterated belief merging. In *Proceedings of the Eighth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'05)*, 514–526.

Grant, J., and Hunter, A. 2006. Measuring inconsistency in knowledgebases. *Journal of Intelligent Information Systems*. (in press).

Grant, J. 1978. Classifications for inconsistent theories. *Notre Dame Journal of Formal Logic* 19:435–444.

Hunter, A., and Konieczny, S. 2005. Approaches to measuring inconsistent information. In *Inconsistency Tolerance*, volume LNCS 3300, 189–234. Springer.

Hunter, A. 2002. Measuring inconsistency in knowledge via quasi-classical models. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI'2002)*, 68–73. MIT Press.

Hunter, A. 2004. Logical comparison of inconsistent perspectives using scoring functions. *Knowledge and Information Systems Journal* 6(5):528–543.

Hunter, A. 2006. How to act on inconsistent news: Ignore, resolve, or reject. *Data and Knowledge Engineering*. (in press).

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

Knight, K. M. 2001. Measuring inconsistency. *Journal of Philosophical Logic* 31:77–98.

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

Konieczny, S.; Lang, J.; and Marquis, P. 2003. Quantifying information and contradiction in propositional logic through epistemic tests. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI'03)*, 106–111.

Konieczny, S. 2004. Belief base merging as a game. *Journal of Applied Non-Classical Logics* 14(3):275–294.

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

Manor, R., and Rescher, N. 1970. On inferences from inconsistent information. *Theory and Decision* 1:179–219.

Meyer, T.; Foo, N.; Zhang, D.; and Kwok, R. 2004. Logical foundations of negotiation: Strategies and preferences. In *Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR'04)*, 311–318.

Mu, K.; Jin, Z.; Lu, R.; and Liu, W. 2005. Measuring inconsistency in requirements engineering. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU05)*, volume 3571 of *LNAI*, 440–451. Springer.

Nebel, B. 1991. Belief revision and default reasoning: syntax-based approaches. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR'91)*, 417–428.

Paris, J. B. 1994. *The uncertain reasoner's companion: a mathematical perspective*. Number 39 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press.

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

Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32:57–95.

Shapley, L. S. 1953. A value for n-person games. In Kuhn, H. W., and Tucker, A. W., eds., *Contributions to the Theory of Games II (Annals of Mathematics Studies 28)*, 307–317. Princeton University Press.

Spohn, W. 1987. Ordinal conditional functions: a dynamic theory of epistemic states. In Harper, W. L., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume 2. Princeton University Press. 105–134.

Williams, M. 1995. Iterated theory base change: A computational model. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, 1541–1547.

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

Zhang, D.; Foo, N.; Meyer, T.; and Kwok, R. 2004. Negotiation as mutual belief revision. In *Proceedings of the Nineteenth American National Conference on Artificial Intelligence (AAAI'04)*, 317–322.