

Measuring Disagreement among Knowledge Bases

Nico Potyka

University of Osnabrück, Germany

Abstract. When combining beliefs from different sources, often not only new knowledge but also conflicts arise. In this paper, we investigate how we can measure the disagreement among sources. We start our investigation with disagreement measures that can be induced from inconsistency measures in an automated way. After discussing some problems with this approach, we propose a new measure that is inspired by the η -inconsistency measure. Roughly speaking, it measures how well we can satisfy all sources simultaneously. We show that the new measure satisfies desirable properties, scales well with respect to the number of sources and illustrate its applicability in inconsistency-tolerant reasoning.

1 Introduction

One challenge in logical reasoning are conflicts between given pieces of information. Therefore, a considerable amount of work has been devoted to repairing inconsistent knowledge bases [1, 2] or performing paraconsistent reasoning [3–5]. Inconsistency measures [6, 7] quantify the degree of inconsistency and help analyzing and resolving conflicts. While work on measuring inconsistency was initially inspired by ideas from repairing knowledge bases and paraconsistent reasoning [8], inconsistency measures also inspired new repair [9, 10] and paraconsistent reasoning mechanisms [11, 12].

Here, we are interested in belief profiles $(\kappa_1, \dots, \kappa_n)$ rather than single knowledge bases κ . Intuitively, we can think of each κ_i as the set of beliefs of an agent. Our goal is then to measure the disagreement among the agents. A natural idea is to reduce measuring disagreement to measuring inconsistency by transforming multiple knowledge bases to a single base using multiset union or conjunction. However, both approaches have some flaws as we will discuss in the following. This observation is similar to the insight that merging belief profiles should be guided by other principles than repairing single knowledge bases [13]. We will therefore propose some new principles for measuring disagreement and introduce a new measure that complies with them.

After explaining the necessary basics in Section 2, we will discuss the relationship between inconsistency measures and disagreement measures in Section 3. To begin with, we will define disagreement measures as functions with two basic properties that seem quite indisputable. We will then show that disagreement measures induced from inconsistency measures by taking the multiset union or conjunction satisfy these basic desiderata and give us some additional guarantees. In Section 4, we will propose some stronger principles for measuring disagreement. One key idea is to allow resolving conflicts by majority decisions. We will show that many measures that are induced from inconsistency measures must necessarily violate some of these principles. In Section 5, we will then introduce a new disagreement measure that is inspired by the

η -inconsistency measure from [6]. Intuitively, it attempts to satisfy all agents' beliefs as well as possible and then measures the average dissatisfaction. We will show that the measure satisfies the principles proposed in Section 4 and some other properties that correspond to principles for measuring inconsistency. To give additional motivation for this work, we will sketch how the measure can be used for belief merging and inconsistency-tolerant reasoning at the end of Section 5.

2 Basics

We consider a propositional logical language \mathcal{L} built up over a finite set \mathcal{A} of propositional atoms using the usual connectives. Satisfaction of formulas $F \in \mathcal{L}$ by valuations $v : \mathcal{A} \rightarrow \{0, 1\}$ is defined as usual. A *knowledge base* κ is a non-empty finite multiset over \mathcal{L} . \mathcal{K} denotes the set of all knowledge bases. An n -tuple $\mathcal{B} = (\kappa_1, \dots, \kappa_n) \in \mathcal{K}^n$ is called a *belief profile*. We let $\sqcup \mathcal{B} = \sqcup_{i=1}^n \kappa_i$, where \sqcup denotes multiset union. Note that using multisets is crucial to avoid information loss when several sources contain syntactically equal beliefs. For instance, $\{\neg a\} \sqcup \{a\} \sqcup \{a\} = \{\neg a, a, a\}$. We let $\mathcal{B} \circ \kappa = (\kappa_1, \dots, \kappa_n, \kappa)$, that is, $\mathcal{B} \circ \kappa$ is obtained from \mathcal{B} by adding κ at the end of the profile. Furthermore, we let $\mathcal{B} \circ^1 \kappa = \mathcal{B} \circ \kappa$ and $\mathcal{B} \circ^k \kappa = (\mathcal{B} \circ^{k-1} \kappa) \circ \kappa$ for $k > 1$. That is, $\mathcal{B} \circ^k \kappa$ is obtained from \mathcal{B} by adding k copies of κ . We call a non-contradictory formula f *safe in* κ iff f and κ are built up over distinct variables from \mathcal{A} . Intuitively, adding a safe formula to κ cannot introduce any conflicts.

A model of κ is a valuation v that satisfies all $f \in \kappa$. We denote the set of all *models* of κ by $\text{Mod}(\kappa)$. If $\text{Mod}(\kappa) \neq \emptyset$, we call κ consistent and inconsistent otherwise. A *minimal inconsistent (maximal consistent) subset* of κ is a subset of κ that is inconsistent (consistent) and minimal (maximal) with this property. If $\text{Mod}(\kappa) \subseteq \text{Mod}(\kappa')$, we say that κ entails κ' and write $\kappa \models \kappa'$. If $\kappa \models \kappa'$ and $\kappa' \models \kappa$, we call κ and κ' equivalent and write $\kappa \equiv \kappa'$. If $\kappa = \{f\}$ and $\kappa' = \{g\}$ are singletons, we just write $f \models g$ or $f \equiv g$.

An inconsistency measure $\mathcal{I} : \mathcal{K}^n \rightarrow \mathbb{R}_0^+$ maps knowledge bases to non-negative degrees of inconsistency. The most basic example is the *drastic measure* that yields 0 if the knowledge base is consistent and 1 otherwise [14]. Hence, it basically performs a satisfiability test. There exist various other measures, see [15] for a recent overview. While there is an ongoing debate about what properties an inconsistency measure should satisfy, there is general agreement that it should be *consistent* in the sense that $\mathcal{I}(\kappa) = 0$ if and only if κ is consistent. Hence, the inconsistency value is greater than zero if and only if κ is inconsistent. Various other properties of inconsistency measures have been discussed [14, 16, 15]. We will present some of these later, when talking about corresponding properties of disagreement measures.

3 Induced Disagreement Measures

To begin with, we define disagreement measures as functions over the set of all belief profiles $\bigcup_{n=1}^{\infty} \mathcal{K}^n$ that satisfy two basic desiderata.

Definition 1 (Disagreement Measure). A disagreement measure is a function $\mathcal{D} : \bigcup_{n=1}^{\infty} \mathcal{K}^n \rightarrow \mathbb{R}_0^+$ such that for all belief profiles $\mathcal{B} = (\kappa_1, \dots, \kappa_n)$, we have

1. Consistency: $\mathcal{D}(\mathcal{B}) = 0$ iff $\bigsqcup_{i=1}^n \kappa_i$ is consistent.
2. Symmetry: $\mathcal{D}(\mathcal{B}) = \mathcal{D}(\kappa_{\sigma(1)}, \dots, \kappa_{\sigma(n)})$ for each permutation σ of $\{1, \dots, n\}$.

Consistency generalizes the corresponding property for inconsistency measures. Symmetry assures that the disagreement value is independent of the order in which the knowledge bases are presented. It is similar to *Anonymity* in social choice theory [17] and guarantees equal treatment of different sources.

Note that each disagreement measure \mathcal{D} induces a *corresponding inconsistency measure* $\mathcal{I}_{\mathcal{D}} : \mathcal{K} \rightarrow \mathbb{R}_0^+$ defined by $\mathcal{I}_{\mathcal{D}}(\kappa) = \mathcal{D}(\kappa)$. Conversely, we can induce disagreement measures from inconsistency measures as we discuss next.

3.1 \sqcup -induced Disagreement Measures

It is easy to see that each inconsistency measure induces a corresponding disagreement measure by taking the multiset union of knowledge bases in the profile.

Proposition 1 (\sqcup -induced Measure). *If \mathcal{I} is an inconsistency measure, then the function $\mathcal{D}_{\mathcal{I}}^{\sqcup} : \bigcup_{n=1}^{\infty} \mathcal{K}^n \rightarrow \mathbb{R}_0^+$ defined by $\mathcal{D}_{\mathcal{I}}^{\sqcup}(\mathcal{B}) = \mathcal{I}(\bigsqcup \mathcal{B})$ for all $\mathcal{B} \in \mathcal{K}^n$ is a disagreement measure. We call $\mathcal{D}_{\mathcal{I}}^{\sqcup}$ the measure \sqcup -induced by \mathcal{I} .*

What can we say about the properties of \sqcup -induced measures? As we explain first, many properties for inconsistency measures have a natural generalization to disagreement measures that is compatible with \sqcup -induced measures in the following sense.

Definition 2 (Corresponding Properties). *Let P be a property for inconsistency measures and let P' be a property for disagreement measures. We call (P, P') a pair of corresponding properties iff*

1. if an inconsistency measure \mathcal{I} satisfies P , then the \sqcup -induced measure $\mathcal{D}_{\mathcal{I}}^{\sqcup}$ satisfies P' ,
2. if a disagreement measure \mathcal{D} satisfies P' , then the corresponding inconsistency measure $\mathcal{I}_{\mathcal{D}}$ satisfies P .

One big class of properties for inconsistency measures gives guarantees about the relationship between inconsistency values when we extend the knowledge bases by particular formulas. We start with a general lemma and give some examples in the subsequent proposition.

Lemma 1 (Transfer Lemma). *Let R be a binary relation on \mathbb{R} and let $C \subseteq \mathcal{K}^3$ be a ternary constraint on knowledge bases. Given a property for inconsistency measures*

$$\text{For all } \kappa, S, T \in \mathcal{K}, \text{ if } C(\kappa, S, T) \text{ then } \mathcal{I}(\kappa \sqcup S) R \mathcal{I}(\kappa \sqcup T), \quad (1)$$

define a property for disagreement measures as follows:

$$\text{For all } \kappa_1, \dots, \kappa_n, S, T \in \mathcal{K}, \text{ if } C\left(\bigsqcup_{i=1}^n \kappa_i, S, T\right) \text{ then} \\ \mathcal{D}(\kappa_1 \sqcup S, \kappa_2, \dots, \kappa_n) R \mathcal{D}(\kappa_1 \sqcup T, \kappa_2, \dots, \kappa_n). \quad (2)$$

Then $((1), (2))$ is a pair of corresponding properties.

Remark 1. The reader may wonder why the corresponding property looks only at the first argument. Note that by symmetry of disagreement measures, the same is true for all other arguments. For instance, we have $\text{Inc}^*(\kappa_1, \kappa_2 \sqcup S) = \text{Inc}^*(\kappa_2 \sqcup S, \kappa_1) \text{ R } \text{Inc}^*(\kappa_2 \sqcup T, \kappa_1) = \text{Inc}^*(\kappa_1, \kappa_2 \sqcup T)$.

We now apply Lemma 1 to some basic properties for inconsistency measures from [14] and adjunction invariance from [16] that will play an important role later.

Proposition 2. *The following are pairs of corresponding properties for inconsistency and disagreement measures:*

- Monotony:
 $\mathcal{I}(\kappa) \leq \mathcal{I}(\kappa \sqcup \kappa')$
 $\mathcal{D}(\kappa_1, \kappa_2, \dots, \kappa_n) \leq \mathcal{D}(\kappa_1 \sqcup \kappa', \kappa_2, \dots, \kappa_n)$
- Dominance: For $f, g \in \mathcal{L}$ such that $f \models g$ and $f \not\models \perp$,
 $\mathcal{I}(\kappa \sqcup \{f\}) \geq \mathcal{I}(\kappa \sqcup \{g\})$
 $\mathcal{D}(\kappa \sqcup \{f\}, \kappa_2, \dots, \kappa_n) \geq \mathcal{D}(\kappa \sqcup \{g\}, \kappa_2, \dots, \kappa_n)$
- Safe Formula Independence: If $f \in \mathcal{L}$ is safe in κ , then
 $\mathcal{I}(\kappa \sqcup \{f\}) = \mathcal{I}(\kappa)$
 If $f \in \mathcal{L}$ is safe in $\bigsqcup_{i=1}^n \kappa_i$, then
 $\mathcal{D}(\kappa_1 \sqcup \{f\}, \kappa_2, \dots, \kappa_n) = \mathcal{D}(\kappa_1, \kappa_2, \dots, \kappa_n)$
- Adjunction Invariance: For all $f, g \in \mathcal{L}$,
 $\mathcal{I}(\kappa \cup \{f, g\}) = \mathcal{I}(\kappa \cup \{f \wedge g\})$
 $\mathcal{D}(\kappa_1 \cup \{f, g\}, \kappa_2, \dots) = \mathcal{D}(\kappa_1 \cup \{f \wedge g\}, \kappa_2, \dots)$

Monotony demands that adding knowledge can never decrease the disagreement value. Dominance says that replacing a claim with a (possibly weaker) implication of the original claim can never increase the disagreement value. Safe Formula Independence demands that a safe formula does not affect the disagreement value. Adjunction invariance says that it makes no difference whether two pieces of information are presented independently or as a single formula.

Example 1. The inconsistency measure \mathcal{I}_{LP_m} that was discussed in [18] satisfies Monotony, Dominance, Safe Formula Independence and Adjunction Invariance. From Proposition 2, we can conclude that the \sqcup -induced disagreement measure $\text{Inc}_{LP_m}^{\sqcup}$ satisfies the corresponding properties for disagreement measures.

What we can take from our discussion so far is that each inconsistency measure induces a disagreement measure with similar properties. As it turns out, each \sqcup -induced disagreement measure satisfies an additional property and, in fact, only the \sqcup -induced measures do. We call this property *partition invariance*. Intuitively, partition invariance means that the disagreement value depends only on the pieces of information in the belief profile and is independent of the distribution of these pieces. In the following proposition, a partition of a multiset M is a sequence of non-empty multisets M_1, \dots, M_k such that $\bigsqcup_{i=1}^k M_i = M$.

Proposition 3 (Characterizations of Induced Families). *The following statements are equivalent:*

1. \mathcal{D} is \sqcup -induced by an inconsistency measure.
2. \mathcal{D} is \sqcup -induced by $\mathcal{I}_{\mathcal{D}}$.
3. \mathcal{D} is partition invariant, that is, for all $\kappa \in \mathcal{K}$ and for all partitions $\bigsqcup_{i=1}^{n_1} P_i = \bigsqcup_{i=1}^{n_2} P'_i = \kappa$ of κ , we have that $\mathcal{D}(P_1, \dots, P_{n_1}) = \mathcal{D}(P'_1, \dots, P'_{n_2})$.

So the \sqcup -induced disagreement measures are exactly the partition invariant measures. However, partition variance can be undesirable in some scenarios.

Example 2. Consider the political goals 'increase wealth of households' (h), 'increase wealth of firms' (f), 'increase wages' (w). Suppose there are three political parties whose positions we represent in the profile

$$\mathcal{B} = (\{f, w, f \rightarrow w\}, \{w, h, w \rightarrow h\}, \{f, \neg w, w \rightarrow \neg f\}).$$

In this scenario, the parties only disagree about w . We modify \mathcal{B} by moving $w \rightarrow \neg f$ from the third to the second party:

$$\mathcal{B}' = (\{f, w, f \rightarrow w\}, \{w, h, w \rightarrow h, w \rightarrow \neg f\}, \{f, \neg w\}).$$

The conflict with respect to w remains, but party 2's positions now imply $\neg f$. Since we now have an additional conflict with respect to f , we would expect $\mathcal{D}(\mathcal{B}) < \mathcal{D}(\mathcal{B}')$.

Partition invariant measures are unable to detect the difference in Example 2. Since partition invariance is an inherent property of \sqcup -induced measures, we should also investigate non- \sqcup -induced measures.

3.2 \wedge -induced disagreement Measures

Instead of taking the multiset union of all knowledge bases in the profile, we can also just replace each knowledge base with the conjunction of the formulas that it contains in order to induce a disagreement measure.

Proposition 4 (\wedge -induced Measure). *If \mathcal{I} is an inconsistency measure, then $\mathcal{D}_{\mathcal{I}}^{\wedge} : \bigcup_{n=1}^{\infty} \mathcal{K}^n \rightarrow \mathbb{R}_0^+$ defined by $\mathcal{D}_{\mathcal{I}}^{\wedge}(\mathcal{B}) = \mathcal{I}(\bigsqcup_{\kappa \in \mathcal{B}} \{\bigwedge_{F \in \kappa} F\})$ for $\mathcal{B} \in \mathcal{K}^n$ is a disagreement measure. We call $\mathcal{D}_{\mathcal{I}}^{\wedge}$ the measure \wedge -induced by \mathcal{I} .*

By repeated application of adjunction invariance (c.f. Proposition 2), one can show that each adjunction invariant inconsistency measure satisfies $\mathcal{I}(\kappa) = \mathcal{I}(\{\bigwedge_{f \in \kappa} f\})$, see [16], Proposition 9. We can use this result to show that for adjunction invariant inconsistency measures, the \wedge -induced and the \sqcup -induced measures are equal.

Corollary 1. *If \mathcal{I} is an adjunction invariant inconsistency measure, then $\mathcal{D}_{\mathcal{I}}^{\wedge} = \mathcal{D}_{\mathcal{I}}^{\sqcup}$.*

This is actually the only case in which the \wedge -induced measure can be \sqcup -induced.

Proposition 5. *Let \mathcal{I} be an inconsistency measure. $\mathcal{D}_{\mathcal{I}}^{\wedge}$ is \sqcup -induced if and only if \mathcal{I} is adjunction invariant.*

The \sqcup -induced disagreement measures are characterized by partition invariance. Adjunction invariance plays a similar role for \wedge -induced measures.

Proposition 6. *For each inconsistency measure \mathcal{I} , $\mathcal{D}_{\mathcal{I}}^{\wedge}$ satisfies adjunction invariance.*

Note that the inconsistency measure $\mathcal{I}_{\mathcal{D}_{\mathcal{I}}^{\wedge}}$ induced by $\mathcal{D}_{\mathcal{I}}^{\wedge}$ will also be adjunction invariant. Therefore, $\mathcal{I}_{\mathcal{D}_{\mathcal{I}}^{\wedge}} \neq \mathcal{I}$ if \mathcal{I} is not adjunction invariant. In particular, $\mathcal{D}_{\mathcal{I}}^{\wedge}$ can be a rather coarse measure if \mathcal{I} is not adjunction invariant.

Example 3. The inconsistency measure \mathcal{I}_{MI} from [18] counts the number of minimal inconsistent sets of a knowledge base. \mathcal{I}_{MI} is not adjunction invariant. For instance, $\mathcal{I}_{MI}(\{a, \neg a, a \wedge b\}) = 2$ because $\{a, \neg a\}$ and $\{\neg a, a \wedge b\}$ are the only minimal inconsistent sets. However, $\mathcal{I}_{MI}(\{a \wedge \neg a \wedge a \wedge b\}) = 1$ because the only minimal inconsistent set is the knowledge base itself. Furthermore, we will have $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}(\kappa) = 1$ whenever $\bigwedge_{f \in \kappa} f$ is inconsistent and $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}(\kappa) = 0$ otherwise. Hence, the inconsistency measure corresponding to $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}$ is the drastic measure.

Proposition 6 tells us that \wedge -induced measures are necessarily adjunction invariant. Whether or not each adjunction invariant disagreement measure is \wedge -induced is currently an open question. However, we have the following result.

Proposition 7. *If \mathcal{D} satisfies adjunction invariance and*

$$\mathcal{D}(\{f_1\}, \dots, \{f_n\}) = \mathcal{D}\left(\bigsqcup_{i=1}^n \{f_i\}\right), \quad (3)$$

then \mathcal{D} is \wedge -induced by an inconsistency measure.

We call property (3) *singleton union invariance* in the following. While adjunction invariance and singleton union invariance are sufficient for being \wedge -induced, they are no longer necessary as the following example illustrates.

Example 4. Consider again the inconsistency measure \mathcal{I}_{MI} from [18] that was explained in Example 3. We have $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}(\{a \wedge b\}, \{\neg a \wedge b\}, \{a \wedge \neg b\}) = \mathcal{I}_{MI}(\{a \wedge b, \neg a \wedge b, a \wedge \neg b\}) = 3$ by definition of the \wedge -induced measure. However, $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}(\{a \wedge b, \neg a \wedge b, a \wedge \neg b\}) = \mathcal{I}_{MI}(\{a \wedge b \wedge \neg a \wedge b \wedge a \wedge \neg b\}) = 1$. Hence, $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}$ is not singleton union invariant.

We close this section by showing that the set of disagreement measures \sqcup -induced and \wedge -induced from inconsistency measures are neither equal nor disjoint.

To begin with, the \mathcal{I}_{LP_m} inconsistency measure that was discussed in [18] is adjunction invariant. Therefore, $\mathcal{D}_{\mathcal{I}_{LP_m}}^{\sqcup} = \mathcal{D}_{\mathcal{I}_{LP_m}}^{\wedge}$ according to Corollary 1. Hence, the intersection of \sqcup -induced and \wedge -induced disagreement measures is non-empty.

In order to show that there are partition invariant measures that are not adjunction invariant and vice versa, we use the minimal inconsistent set measure \mathcal{I}_{MI} from [18]. As demonstrated in Example 3, \mathcal{I}_{MI} is not adjunction invariant. Therefore, the Transfer Lemma implies that $\mathcal{D}_{\mathcal{I}_{MI}}^{\sqcup}$ is not adjunction invariant either. Hence, $\mathcal{D}_{\mathcal{I}_{MI}}^{\sqcup}$ cannot be \wedge -induced according to Proposition 6.

On the other hand, $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}$ is adjunction invariant because each \wedge -induced measure is. However, since \mathcal{I}_{MI} is not adjunction invariant, we know from Proposition 5 that $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}$ is not \sqcup -induced. Hence, $\mathcal{D}_{\mathcal{I}_{MI}}^{\wedge}$ is an example of a disagreement measure that is \wedge -induced, but not \sqcup -induced.




Fig. 1. Venn diagram illustrating induced measures in the space of all disagreement measures.

We illustrate our findings in Figure 1. The \wedge -induced incompatibility measures are a subset of the adjunction invariant measures (Proposition 6). The fact that all measures in the intersection of partition invariant and adjunction invariant measures are \wedge -induced follows from observing that partition invariance implies singleton union invariance (3) and Proposition 7.

4 Principles for Measuring Disagreement

As illustrated in Figure 1, induced measures correspond to disagreement measures with very specific properties. \perp -induced measures are necessarily partition invariant. This may be undesirable in certain applications as illustrated in Example 2. If an inconsistency measure is adjunction invariant, the \wedge -induced measure will also be partition invariant. If it is not adjunction invariant, the \wedge -induced measure will not be partition invariant, but the measure may become rather coarse as illustrated in Example 3. This is some evidence that it is worth investigating non-induced measures. To further distinguish inconsistency from disagreement measures, we will now propose some stronger principles that go beyond our basic desiderata from Definition 1.

To guide our intuition, we think of each knowledge base as the belief set of an agent. We say that κ_i *contradicts* κ_j if $\kappa_i \cup \kappa_j$ is inconsistent. To begin with, let us consider an agent whose beliefs do not contradict any consistent position (its knowledge base is tautological). When adding such an agent to a belief profile, the disagreement value should not increase. Dually, if we add an agent that contradicts every position (its knowledge base is inconsistent), the disagreement value should not decrease. This intuition is captured by the following principles.

Tautology Let $\mathcal{B} \in \mathcal{K}^n$ and let $\kappa_{\top} \in \mathcal{K}$ be tautological. Then $\mathcal{D}(\mathcal{B} \circ \kappa_{\top}) \leq \mathcal{D}(\mathcal{B})$.

Contradiction Let $\mathcal{B} \in \mathcal{K}^n$ and let $\kappa_{\perp} \in \mathcal{K}$ be contradictory. Then $\mathcal{D}(\mathcal{B} \circ \kappa_{\perp}) \geq \mathcal{D}(\mathcal{B})$.

Inconsistency measures focus mainly on the existence of conflicts. However, in a multiagent setting, conflicts can often be resolved by majority decisions. Given a belief profile $\mathcal{B} = (\kappa_1, \dots, \kappa_n) \in \mathcal{K}^n$, we call a subset $C \subseteq \{1, \dots, n\}$ a *consistent coalition*

iff $\bigcup_{i \in C} \kappa_i$ is consistent. We say that κ_j is *involved in a conflict in \mathcal{B}* iff there is a consistent coalition C such that $\kappa_j \cup \bigcup_{i \in C} \kappa_i$ is inconsistent. Our next principle demands that conflicts can be eased by majority decisions.

Majority Let $\mathcal{B} = (\kappa_1, \dots, \kappa_n) \in \mathcal{K}^n$. If κ_j is consistent and involved in a conflict, then there is a $k \in \mathbb{N}$ such that $\mathcal{D}(\mathcal{B} \circ^k \kappa_j) < \mathcal{D}(\mathcal{B})$.

Intuitively, Majority says that we can decrease the severity of a conflict by giving sufficient support for one of the conflicting positions. It does not matter what position we choose as long as this position is consistent. In future work, one may look at alternative principles based on other methods to make group decisions [17], but Majority seems to be a natural starting point.

Majority implies that we can strictly decrease the disagreement value by adding copies of one consistent position. However, this does not imply that the disagreement value will vanish. If we keep adding copies, the disagreement value will necessarily decrease but it may converge to a value strictly greater than 0. While one may argue that the limit should be 0 if almost all agents agree, one may also argue that the limit should be bounded from below by a positive constant if an unresolved conflict remains. We therefore do not strengthen majority. Instead, we consider an additional principle that demands that the limit is indeed 0 if the majority agrees on all non-contradictory positions. This intuition is captured by the next principle.

Majority Agreement in the Limit Let $\mathcal{B} \in \mathcal{K}^n$. If M is a \subseteq -maximal consistent subset of $\bigsqcup \mathcal{B}$, then $\lim_{k \rightarrow \infty} \mathcal{D}(\mathcal{B} \circ^k M) = 0$.

We close this section with an impossibility result: Monotony and Partition Invariance cannot be satisfied jointly with our majority principles. The reason is that such measures can never decrease when receiving new information as explained in the following proposition.

Proposition 8. *If \mathcal{D} satisfies Monotony and Partition Invariance, then $\mathcal{D}(\mathcal{B} \circ^k \kappa) \geq \mathcal{D}(\mathcal{B})$ for all $\mathcal{B} \in \mathcal{K}^n, \kappa \in \mathcal{K}, k \in \mathbb{N}$.*

The conditions of Proposition 8 are in particular met by several induced measures.

Corollary 2. *Every disagreement measure that is*

- *partition invariant and monotone or*
- *\sqcup -induced from a monotone inconsistency measure or*
- *\wedge -induced from a monotone and adjunction invariant inconsistency measure*

violates Majority and Majority Agreement in the Limit.

5 The η -disagreement Measure

We now consider a novel disagreement measures inspired by the η -inconsistency measure from [6]. Roughly speaking, the *η -inconsistency measure* attempts to maximize the probability of all formulas within a knowledge base. By subtracting this probability from 1, we get an inconsistency value. In order to assign probabilities to formulas, we consider probability distributions over the set of all valuations $\Omega = \{v \mid$

$v : \mathcal{A} \rightarrow \{0, 1\}$ of our language. Given a probability distribution $\pi : \Omega \rightarrow [0, 1]$ ($\sum_{v \in \Omega} \pi(v) = 1$) and a formula $F \in \mathcal{L}$, we let

$$P_\pi(F) = \sum_{v \models F} \pi(v).$$

Intuitively, $P_\pi(F)$ is the probability that F is true with respect to π . The η -inconsistency measure from [6] is defined by

$$\mathcal{I}_\eta(\kappa) = 1 - \max\{p \mid \exists \pi : \forall F \in \kappa : P_\pi(F) \geq p\}.$$

This formula describes the intuition that we explained in the beginning. $p^* = \max\{p \mid \exists \pi : \forall F \in \kappa : P_\pi(F) \geq p\}$ is the maximum probability that all formulas in κ can simultaneously take. We will have $p^* = 1$ if and only if κ is consistent [6].

Let us first look at the disagreement measures induced by \mathcal{I}_η . \mathcal{I}_η satisfies Monotony [15]. Therefore, $\mathcal{D}_{\mathcal{I}_\eta}^\sqcup$ will violate our majority principles as explained in Corollary 2. However, \mathcal{I}_η is not adjunction invariant [15]. Therefore, Proposition 5 implies that $\mathcal{D}_{\mathcal{I}_\eta}^\sqcup \neq \mathcal{D}_{\mathcal{I}_\eta}^\wedge$. Still, $\mathcal{D}_{\mathcal{I}_\eta}^\wedge$ does not satisfy our majority principles either.

Example 5. Let $\mathcal{B} = (\{a\}, \{\neg a\})$. Since $P_\pi(a) = 1 - P_\pi(\neg a)$, we have for all $n \in \mathbb{N}$

$$\begin{aligned} \mathcal{D}_{\mathcal{I}_\eta}^\wedge(\{a\}, \{\neg a\}) &= \mathcal{I}_\eta(\{a, \neg a\}) = 0.5 \\ &= \mathcal{I}_\eta(\{a, \neg a\} \sqcup \bigsqcup_{i=1}^n \{a\}) = \mathcal{D}_{\mathcal{I}_\eta}^\wedge((\{a\}, \{\neg a\}) \circ^n \{a\}). \end{aligned}$$

However, we can modify the definition of the η -inconsistency measure in order to get a disagreement measure that satisfies our desiderata. If we think of $P_\pi(F)$ as the degree of belief in F , then we should try to find a π such that the beliefs of all agents are satisfied as well as possible. To do so, we can first look at how well π satisfies the beliefs of each agent and then look at how well π satisfies the agents' beliefs overall. To measure satisfaction of one agent's beliefs, we take the minimum of all probabilities assigned to the formulas in the agent's knowledge base. Formally, for all probability distributions π and knowledge bases κ over our language, we let

$$s_\pi(\kappa) = \min\{P_\pi(F) \mid F \in \kappa\}.$$

and call $s_\pi(\kappa)$ the *degree of satisfaction* of κ . In order to measure satisfaction of a belief profile, we take the average degree of satisfaction of the knowledge bases in the profile. Formally, we let for all probability distributions π and belief profiles \mathcal{B}

$$S_\pi(\mathcal{B}) = \frac{1}{|\mathcal{B}|} \sum_{\kappa \in \mathcal{B}} s_\pi(\kappa)$$

and call $S(\mathcal{B})$ the *degree of satisfaction* of \mathcal{B} . We now define a new disagreement measure. Intuitively, it attempts to maximize the degree of satisfaction of the profile. By subtracting the maximum degree of satisfaction from 1, we get a disagreement value.

Definition 3 (η -Disagreement Measure). *The η -Disagreement Measure is defined by*

$$\mathcal{D}_\eta(\mathcal{B}) = 1 - \max\{p \mid \exists \pi : S_\pi(\mathcal{B}) = p\}.$$

To begin with, we note that \mathcal{D}_η is a disagreement measures as defined in Definition 1 and can be computed by linear programming techniques.

Proposition 9. *\mathcal{D}_η is a disagreement measures and can be computed by solving a linear optimization problem.*

As we show next, \mathcal{D}_η is neither \sqcup - nor \wedge -induced from any inconsistency measure. According to Proposition 3 and Proposition 6, it suffices to show that it is neither partition invariant nor adjunction invariant.

Example 6. Consider again the belief profiles \mathcal{B} and \mathcal{B}' from Example 2. We have $\mathcal{D}_\eta(\mathcal{B}) \approx 0.33$ and $\mathcal{D}_\eta(\mathcal{B}') \approx 0.44$. As desired, \mathcal{D}_η recognizes the increased disagreement in the profile. In particular, \mathcal{D}_η is not partition invariant.

Example 7. To see that \mathcal{D}_η is not adjunction invariant, note that $\mathcal{D}_\eta(\{a, \neg a\}) = 0.5$, whereas $\mathcal{D}_\eta(\{a \wedge \neg a\}) = 1$ (contradictory formulas have probability 0 with respect to each π). Hence, \mathcal{D}_η is also not adjunction invariant.

\mathcal{D}_η satisfies our four principles for measuring disagreement as we show next. To begin with, we note that the disagreement value necessarily decreases as the proportion of agreeing agents increases.

Proposition 10. *Let $\mathcal{B} \in \mathcal{K}^n$. If \mathcal{B} contains a consistent coalition of size k , then $\mathcal{D}_\eta(\mathcal{B}) \leq 1 - \frac{k}{n}$.*

Proposition 10 implies, in particular, that the disagreement value goes to 0 as the proportion of agreeing agents $\frac{k}{n}$ goes to 1. Therefore, \mathcal{D}_η satisfies our majority principles.

Corollary 3. *\mathcal{D}_η satisfies Majority and Majority Agreement in the Limit.*

Tautology and Contradiction are also satisfied and can be strengthened slightly.

Proposition 11. *\mathcal{D}_η satisfies Tautology and Contradiction. Furthermore,*

- If $\mathcal{D}_\eta(\mathcal{B}) > 0$, then $\mathcal{D}_\eta(\mathcal{B} \circ \kappa_\top) < \mathcal{D}_\eta(\mathcal{B})$.
- If $\mathcal{D}_\eta(\mathcal{B}) < 1$, then $\mathcal{D}_\eta(\mathcal{B} \circ \kappa_\perp) > \mathcal{D}_\eta(\mathcal{B})$.

Regarding the properties corresponding to principles for measuring inconsistency from Proposition 2, \mathcal{D}_η satisfies all except Adjunction Invariance (Example 7).

Proposition 12. *\mathcal{D}_η satisfies Monotony, Dominance and Safe Formula Independence.*

We already know that \mathcal{D}_η yields 0 if and only if all knowledge bases in the profile are consistent with each other. In the following proposition, we explain in what cases it takes the maximum value 1.

Proposition 13. *Let $\mathcal{B} \in \mathcal{K}^n$. We have $\mathcal{D}_\eta(\mathcal{B}) = 1$ iff all κ_i contain at least one contradictory formula.*

Intuitively, if there is a knowledge base that does not contain any contradictory formulas, then all beliefs of one agent can be partially satisfied and the disagreement value with respect to \mathcal{D}_η cannot be 1. So the degree of disagreement can only be maximal if each agent has contradictory beliefs.

In some applications, we may want to restrict to belief profiles with consistent knowledge bases. We can rescale \mathcal{D}_η for this purpose. Proposition 10 gives us the following upper bounds on the disagreement value.

Corollary 4. *Let $\mathcal{B} = (\kappa_1, \dots, \kappa_n)$. If some κ_i is consistent, then $\mathcal{D}_\eta(\mathcal{B}) \leq 1 - \frac{1}{n}$.*

The bound in Corollary 4 is actually tight even if all knowledge bases in the profile are individually consistent as we explain in the following example.

Example 8. For $n = 2$ agents, we have $\mathcal{D}_\eta(\{a\}, \{-a\}) = \frac{1}{2}$. For $n = 3$, we have $\mathcal{D}_\eta(\{a \wedge b\}, \{-a \wedge b\}, \{-b\}) = \frac{2}{3}$. In general, if we have n satisfiable but pairwise inconsistent ($F_i \wedge F_j \equiv \perp$) formulas F_1, \dots, F_n , then $\mathcal{D}_\eta(\{F_1\}, \dots, \{F_n\}) = 1 - \frac{1}{n}$.

Hence, if we want to restrict to consistent knowledge bases, we can renormalize \mathcal{D}_η by multiplying by $\frac{n}{n-1}$. The disagreement value will then be maximal whenever all agents have pairwise inconsistent beliefs.

As explained in Proposition 9, computing $\mathcal{D}_\eta(\mathcal{B})$ is a linear optimization problem. Interior-point methods can solve these problems in polynomial time in the number of optimization variables and constraints [19]. While the number of optimization variables is exponential in the number of atoms $|\mathcal{A}|$ of our language, the number of constraints is linear in the number of formulas in all knowledge bases in the profile. Roughly speaking, computing $\mathcal{D}_\eta(\mathcal{B})$ is very sensitive to the number of atoms, but scales well with respect to the number of agents. In the language of parameterized complexity theory [20], computing $\mathcal{D}_\eta(\mathcal{B})$ is fixed-parameter tractable (that is, polynomial if we fix the number of atoms).

Proposition 14. *Computing $\mathcal{D}_\eta(\mathcal{B})$ is fixed-parameter tractable with parameter $|\mathcal{A}|$.*

While interior-point methods give us a polynomial worst-case guarantee, they are often outperformed in practice by the simplex algorithm. The simplex algorithm has exponential runtime for some artificial examples, but empirically runs in time linear in the number of optimization variables (exponential in $|\mathcal{A}|$) and quadratic in the number of constraints (quadratic in the overall number of formulas in the belief profile) [19].

In the long-term, our goal is to reason over belief profiles that contain conflicts among agents. While we must leave a detailed discussion for future work, we will now sketch how the η -disagreement measure can be used for this purpose. The optimal solutions of the linear optimization problem corresponding to \mathcal{D}_η form a topologically closed and convex set of probability distributions. This allows us to compute lower and upper bound on the probability (or more intuitively, the degree of belief) of formulas with respect to the optimal solutions that minimize disagreement. This is similar to the probabilistic entailment problem [21], where we compute lower and upper bounds with respect to probability distributions that satisfy probabilistic knowledge bases. If, for a belief profile \mathcal{B} , the lower bound of the formula F is l and the upper bound is u , we write $P_{\mathcal{B}}(F) = [l, u]$. If $l = u$, we just write $P_{\mathcal{B}}(F) = l$. We call $P_{\mathcal{B}}$ the *aggregated group belief*.

Example 9. Suppose we have 100 reviews about a restaurant. While most reviewers agree that the food (f) and the service (s) are good, two reviewers disagree about the interior design (d) of the restaurant. Let us assume that $\mathcal{B} = (((\{d, f, s\}, \{\neg d, f, s\}) \circ^{0.95} \{f, s\}) \circ^3 \{\neg f, \neg s\})$. We have $\mathcal{D}_\eta(\mathcal{B}) \approx 0.03$. Intuitively, the degree of disagreement among agents is low because the majority of agents seem not to care about the interior design. The aggregated group beliefs for the atoms in this example are $P_{\mathcal{B}}(d) = 0.5$, $P_{\mathcal{B}}(f) = 1$, $P_{\mathcal{B}}(s) = 1$.

We can use $P_{\mathcal{B}}$ to define an entailment relation. For instance, we could say that \mathcal{B} entails F iff the lower bound is strictly greater than 0.5. Then, in Example 9, $P_{\mathcal{B}}$ entails f and s , but neither d nor $\neg d$.

6 Related Work

The authors in [22] considered the problem of measuring disagreement in limited choice problems, where each agent can choose from a finite set of alternatives. The measures are basically defined by counting the decisions and relating the counts. The authors give intuitive justification for their measures, but do not consider general principles. In order to transfer their approach to our setting, one may identify atomic formulas with alternatives in their framework, but it is not clear how this approach could be extended to knowledge bases that contain complex formulas.

Some other conflict measures have been considered in non-classical frameworks. These measures are often closer to distance measures because they mainly compare how close two quantitative belief representations like probability functions, belief functions or fuzzy membership functions are [23–25]. In [26], some compatibility measures for Markov logic networks have been proposed. The measures are normalized and the maximum degree of compatibility can be related to a notion of coherence of Markov logic networks. However, this notion cannot be transferred to classical knowledge bases easily.

As we discussed, measuring disagreement is closely related to measuring inconsistency [6, 14, 27] and merging knowledge bases [28–30]. The principles *Majority* and *Majority Agreement in the Limit* from Section 4 are inspired by *Majority merging operators* that allow that a sufficiently large interest group can determine the merging outcome. The η -disagreement measure is perhaps most closely related to *model-based operators* and *DA² operators*, which attempt to minimize some notion of distance between interpretations and the models of the knowledge bases in the profile. In contrast, the η -disagreement measure minimizes a probabilistic degree of dissatisfaction of the belief profile.

[31] introduced some entailment relations based on consensus in belief profiles. We will investigate relationships to entailment relations derived from the η -disagreement measure in future work.

7 Conclusions and Future Work

In this paper, we investigated approaches to measuring disagreement among knowledge bases. In principle, inconsistency measures can be applied for this purpose by trans-

forming belief profiles to single knowledge bases. However, we noticed some problems with this approach. For instance, many measures that are naively induced from inconsistency measures violate Majority and Agreement in the Limit as explained in Corollary 2. Even though this problem does not apply to measures \wedge -induced from inconsistency measures that violate adjunction invariance, these induced measures show another problem: they may be unable to notice that a conflict can be resolved by giving up parts of agents' beliefs. For instance, the measures $\mathcal{D}_{\mathcal{I}_{MI}}^\wedge$ and $\mathcal{D}_{\mathcal{I}_\eta}^\wedge$ cannot distinguish the profiles $(\{a, b\}, \{\neg a, b\})$ and $(\{a\}, \{\neg a\})$ because \mathcal{I}_{MI} and \mathcal{I}_η cannot distinguish the knowledge bases $\{a \wedge b, \neg a \wedge b\}$ and $\{a, \neg a\}$.

The η -inconsistency measure \mathcal{D}_η satisfies our principles for measuring disagreement and some other basic properties that correspond to principles for measuring inconsistency. Since \mathcal{D}_η can perform satisfiability tests, we cannot expect to compute disagreement values in polynomial time with respect to the number of atoms. However, if our agents argue only about a moderate number of statements (we fix the number of atoms), the worst-case runtime is polynomial with respect to the number of agents.

In the long-term, we are in particular interested in reasoning over belief profiles that contain conflicts. We can use the η -inconsistency measure for this purpose as we sketched at the end of Section 5. However, the aggregated group belief P_B does not behave continuously. For instance, if we gradually increase the support for $\neg s$ in Example 9, $P_B(s)$ will not gradually go to 0, but will jump to an undecided state like 0.5 or will jump to 0 at some point. This is not a principal problem for defining an entailment relation that either says that a formula is entailed or not entailed by a profile. However, a continuous notion of group beliefs would allow us to shift the focus from measuring disagreement among agents to measuring disagreement about statements (logical formulas). We could do so by measuring how well we can bound the aggregated beliefs about the formulas in the profile away from 0.5. However, if P_B does not behave continuously, this approach will give us a rather coarse measure (basically three-valued). Therefore, an interesting question for future research is whether we can modify \mathcal{D}_η or design other measures that give us an aggregated group belief with a more continuous behavior.

References

1. Kalyanpur, A., Parsia, B., Sirin, E., Cuenca-Grau, B.: Repairing unsatisfiable concepts in owl ontologies. In: European Semantic Web Conference, Springer (2006) 170–184
2. Lehmann, J., Bühmann, L.: Ore—a tool for repairing and enriching knowledge bases. In: International Semantic Web Conference, Springer (2010) 177–193
3. Benferhat, S., Dubois, D., Prade, H.: Some syntactic approaches to the handling of inconsistent knowledge bases: A comparative study part 1: The flat case. *Studia Logica* **58**(1) (1997) 17–45
4. Arieli, O., Avron, A., Zamansky, A.: What is an ideal logic for reasoning with inconsistency?. In: IJCAI 2011. (2011) 706–711
5. Priest, G.: Paraconsistent logic. In: Handbook of philosophical logic. Springer (2002) 287–393
6. Knight, K.: Measuring inconsistency. *Journal of Philosophical Logic* **31**(1) (2002) 77–98
7. Grant, J., Hunter, A.: Measuring inconsistency in knowledgebases. *Journal of Intelligent Information Systems* **27**(2) (2006) 159–184

8. Hunter, A., Konieczny, S.: Approaches to measuring inconsistent information. In: Inconsistency tolerance. Springer (2005) 191–236
9. Thimm, M.: Measuring inconsistency in probabilistic knowledge bases. In: UAI 2009, AUA Press (2009) 530–537
10. Muiño, D.P.: Measuring and repairing inconsistency in probabilistic knowledge bases. International Journal of Approximate Reasoning **52**(6) (2011) 828–840
11. Potyka, N.: Linear programs for measuring inconsistency in probabilistic logics. In: KR 2014. (2014) 568–577
12. Potyka, N., Thimm, M.: Probabilistic reasoning with inconsistent beliefs using inconsistency measures. In: IJCAI 2015. (2015) 3156–3163
13. Konieczny, S.: On the difference between merging knowledge bases and combining them. In: KR. (2000) 135–144
14. Hunter, A., Konieczny, S.: Measuring inconsistency through minimal inconsistent sets. KR 2008 **8** (2008) 358–366
15. Thimm, M.: On the compliance of rationality postulates for inconsistency measures: A more or less complete picture. KI-Künstliche Intelligenz (2016) 1–9
16. Besnard, P.: Revisiting postulates for inconsistency measures. In: European Workshop on Logics in Artificial Intelligence, Springer (2014) 383–396
17. Zwicker, W.: Introduction to the theory of voting. Handbook of computational social choice (2016) 23–56
18. Hunter, A., Konieczny, S.: On the measure of conflicts: Shapley inconsistency values. Artificial Intelligence **174**(14) (2010) 1007–1026
19. Matousek, J., Gärtner, B.: Understanding and Using Linear Programming. Universitext (1979). Springer (2007)
20. Flum, J., Grohe, M.: Parameterized complexity theory. Springer Science & Business Media (2006)
21. Hansen, P., Jaumard, B.: Probabilistic satisfiability. In: Handbook of Defeasible Reasoning and Uncertainty Management Systems. Springer (2000) 321–367
22. Whitworth, B., Felton, R.: Measuring disagreement in groups facing limited-choice problems. ACM SIGMIS Database **30**(3-4) (1999) 22–33
23. Liu, W.: Analyzing the degree of conflict among belief functions. Artificial Intelligence **170**(11) (2006) 909–924
24. Castiñeira, E.E., Cubillo, S., Montilla, W.: Measuring incompatibility between atanassovs intuitionistic fuzzy sets. Information Sciences **180**(6) (2010) 820–833
25. Jusselme, A.L., Maupin, P.: Distances in evidence theory: Comprehensive survey and generalizations. International Journal of Approximate Reasoning **53**(2) (2012) 118–145
26. Thimm, M.: Coherence and compatibility of markov logic networks. In: ECAI 2014, IOS Press (2014) 891–896
27. Grant, J., Hunter, A.: Distance-based measures of inconsistency. In: European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, Springer (2013) 230–241
28. Baral, C., Kraus, S., Minker, J.: Combining multiple knowledge bases. IEEE transactions on knowledge and data engineering **3**(2) (1991) 208–220
29. Liberatore, P., Schaerf, M.: Arbitration (or how to merge knowledge bases). IEEE Transactions on Knowledge and Data Engineering **10**(1) (1998) 76–90
30. Konieczny, S., Pérez, R.P.: Logic based merging. Journal of Philosophical Logic **40**(2) (2011) 239–270
31. Grégoire, É., Konieczny, S., Lagniez, J.: On consensus extraction. In: IJCAI 2016, New York, AAAI Press (2016) 1095–1101