

Consolidation via Tacit Culpability Measures: Between Explicit and Implicit Degrees of Culpability

Jandson S. Ribeiro, Matthias Thimm

University of Koblenz-Landau, Germany
{jandson, thimm}@uni-koblenz.de

Abstract

Restoring consistency of a knowledge base, known as consolidation, should preserve as much information as possible of the original knowledge base. On the one hand, the field of *belief change* captures this principle of minimal change via rationality postulates. On the other hand, within the field of inconsistency measurement, *culpability measures* have been developed to assess how much a formula participates in making a knowledge base inconsistent. We look at culpability measures as a tool to disclose epistemic preference relations and build rational consolidation functions. We introduce *tacit culpability measures* that consider semantic counterparts between conflicting formulae, and we define a special class of these culpability measures based on a fixed-point characterisation: *the stable tacit culpability measures*. We show that the stable tacit culpability measures yield rational consolidation functions and that these are also the only culpability measures that yield rational consolidation functions.

1 Introduction

Handling inconsistency is a problem that appears in many scenarios within Computer Science and Artificial Intelligence (AI), such as in integrating distributed data-sources, ontology evolution, formal specification and verification of systems, automated reasoning and commonsense reasoning. Two main fields within AI that study how to handle inconsistencies are *belief change* (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988; Hansson 1999) and *inconsistency/culpability measures* (Grant 1978; Hunter and Konieczny 2005; Thimm 2019).

Belief change studies how to keep consistency of an agent's epistemic state as it evolves. One of the most studied belief change operations is *contraction*, which consists in relinquishing an unwanted (or obsolete) piece of information. Operations such as contraction are governed by sets of rationality postulates that prescribe suitable behaviours for belief change operators. These rationality postulates are motivated by the principle of minimal change: when removing an undesirable/obsolete piece of information φ , an agent should remove only those pieces of information that contribute to entailing φ . In addition, classes of belief change

operators that are characterised by such rationality postulates, called rational belief change operators, were proposed, such as the smooth kernel contraction operators (Hansson 1994) and partial meet contraction operators (Alchourrón, Gärdenfors, and Makinson 1985). A main difference between these two classes of contraction functions is that while smooth kernel functions choose what to remove, partial meet functions choose what to keep (Falappa, Fermé, and Kern-Isberner 2006). In both approaches, the actual choice of what to remove/keep is postponed to an extra-logical mechanism that realises the agent's epistemic preferences.

A special kind of belief change operation is *consolidation* that purges inconsistencies from a knowledge base. Consolidation corresponds to contracting by *falsum* (Hansson 1997). In this paper, we focus on the consolidation operation from the perspective of rational belief change, specifically the operations based on the what-to-remove strategy (smooth kernel contraction).

While belief change has a qualitative perspective to handle inconsistencies, the field of *inconsistency measurement* (Grant 1978; Hunter and Konieczny 2005; Thimm 2019) handles inconsistencies from a quantitative perspective: it proposes functions, called inconsistency measures, that tag a knowledge base with a real number that indicates its degree of inconsistency. Inconsistency measures are helpful in pinpointing causes of inconsistencies in a knowledge base, and assist in semi-automatic approaches for inconsistency removal, see e. g. (Grant and Martinez 2018) for some concrete applications. A related family of measures are *culpability measures* (Daniel 2009), also called *inconsistency values* (Hunter and Konieczny 2006), that assess how much a formula participates in making a knowledge base inconsistent. The higher the culpability value of a formula, the more responsible for producing inconsistency is such a formula.

The values that a culpability measure assigns to the formulae of a knowledge base could be interpreted as the epistemic preference relation of an agent, when its knowledge base becomes inconsistent. The higher the culpability value of a formula, the less reliable/entrenched it is. From this perspective, a culpability measure induces a choice function that points to the most undesirable pieces of information that should be removed. Thus, a culpability measure can be used to construct a consolidation function by preferably removing those formulae with high culpability values.

Unfortunately, as we show in Section 4, the few culpability measures existing in the literature do not induce rational consolidation functions. We argue that this failure occurs because these measures rely significantly on syntactic aspects of knowledge bases and disregard the semantic counterparts. To overcome this problem, we propose a new class of culpability measures that consider both syntactic and semantic elements: the *tacit culpability measures*. In addition, we construct a special kind of tacit culpability measures based on a fixed-point characterisation: the *stable tacit culpability measures*. We show that stable tacit culpability measures not only yield rational consolidation functions, but they are also the only class of culpability measures capable of generating such rational consolidation functions.

In Section 2, we introduce some notation that will be used throughout this paper. In Section 3, we briefly recap belief contraction and consolidation. In Section 4 we give a brief introduction to culpability measures, and we show their limitation in inducing rational consolidation functions. In Section 5, we introduce the notions of explicit and implicit roles in producing inconsistency. From the latter, we define tacit culpability measures. This notion of implicit role also helps us in identifying a sufficient condition for a culpability measure to yield a rational consolidation function: *tacit dominance*. We then exploit tacit dominance in order to devise a novel class of culpability measures that yield rational consolidation function by combining both explicit and implicit roles of culpabilities: the *stable tacit culpability measures*. Finally, in Section 6, we make some final consideration and discuss some future works. Full proofs are presented in the appendix at http://mthimm.de/misc/proofs_kr21_jsmt.pdf.

2 Notation and Technical Background

Let At be some fixed finite set of propositional symbols, and \mathcal{L} be the propositional language built upon At and the usual Boolean connectives \wedge (conjunction), \vee (disjunction), \rightarrow (implication), \leftrightarrow (equivalence) and the classical negation operator \neg . Let \models be the usual semantic consequence relation of propositional logic. A knowledge base K is a finite set of formulæ. Given two formulæ φ and ψ , we write $\varphi \equiv \psi$ to indicate that φ and ψ are logically equivalent.

Given a set D , and a binary relation $<: D \times D$, we define $max_{<}(X) = \{a \in X \mid a \not< b \text{ or } b < a, \text{ for all } b \in X\}$, and $min_{<}(X) = \{a \in X \mid b \not< a \text{ or } a < b, \text{ for all } b \in X\}$, for all $X \subseteq D$. Let $\mathbb{R}_{\geq 0}$ denote the set of all non-negative real numbers. Let $f: D \rightarrow \mathbb{R}_{\geq 0}$ be a function that maps some elements from a domain D to a non-negative real number. For a finite set $X \subseteq D$, we define $max_f(X) = y$, such that: (1) $y = f(a)$ for some $a \in X$ and $y \geq f(b)$ for all $b \in X$, if $X \neq \emptyset$; (2) $y = 0$; if $X = \emptyset$. Analogously, $min_f(X) = y$, where: (1) $y = f(a)$ for some $a \in X$ and $y \leq f(b)$ for all $b \in X$, if $X \neq \emptyset$; (2) $y = 0$; if $X = \emptyset$. Given two functions $g, g': D \rightarrow D'$, we write $g \equiv g'$ to denote they are equivalent, that is, $g(a) = g'(a)$, for all $a \in D$.

3 Belief Contraction and Consolidation

A major approach to representing an agent's beliefs is via a set of sentences known as *belief base*. The term belief base

has been used in the literature with two specific purposes: (i) as a finite representation of an agent's beliefs (Nebel 1990; Dixon 1994; Dalal 1988), and (ii) as a way of differentiating the explicit beliefs of an agent from its implicit beliefs (Fuhrmann 1991; Hansson 1999). In this work, we look at a belief base as a finite representation of an agent's beliefs. We will use the terms belief base and knowledge base interchangeably.

Removal of an undesirable belief is known as contraction. Let K be a belief base, a contraction function for K is a function $\dot{-}: \mathcal{L} \rightarrow 2^{\mathcal{L}}$ that given an unwanted piece of information outputs a subset of K . A contraction function is subject to the following basic rationality postulates (Hansson 1991, 1994):

(success): if $\not\models \varphi$ then $K \dot{-} \varphi \not\models \varphi$;

(inclusion): $K \dot{-} \varphi \subseteq K$;

(vacuity): if $K \not\models \varphi$ then $K \dot{-} \varphi = K$;

(uniformity): if for all $K' \subseteq K$ it holds that $K' \models \varphi$ iff $K' \models \psi$, then $K \dot{-} \varphi = K \dot{-} \psi$;

(core-retainment): if $\psi \in K \setminus (K \dot{-} \varphi)$ then there exists a $K' \subseteq K$ s.t $K \dot{-} \varphi \subseteq K'$, $K' \not\models \psi$, $K' \cup \{\varphi\} \models \psi$.

A contraction function that satisfies all the rationality postulates above is called a *rational contraction function*. For a discussion about the rationale of these postulates see (Hansson 1999). The five rationality postulates above characterise the class of *kernel contraction functions* (Definition 3 below), which is founded on α -kernels and incision functions:

Definition 1. An α -kernel of a knowledge base K is a set X such that (1) $X \subseteq K$; (2) $X \models \alpha$; and (3) if $X' \subset X$ then $X' \not\models \alpha$.

An α -kernel of a knowledge base K is a minimal subset of K that does entail α . The set of all α -kernels of a knowledge base K is denoted by $K \perp\!\!\!\perp \alpha$.

Definition 2. An incision function σ is a function such that

- $\sigma(K \perp\!\!\!\perp \varphi) \subseteq \bigcup K \perp\!\!\!\perp \varphi$;
- if $X \in K \perp\!\!\!\perp \varphi$ and $X \neq \emptyset$, then $X \cap \sigma(K \perp\!\!\!\perp \varphi) \neq \emptyset$.

Intuitively, given a formula α to be relinquished, an incision function σ for K selects at least one formula from each α -kernel within K . The rationale is that σ behaves as an epistemic choice function that realises the epistemic preferences of the underlying agent, and it picks the least entrenched pieces of information to be removed.

Definition 3. Given a belief base K and an incision function σ for K , the kernel contraction operator $\dot{-}_\sigma$ is defined as: $K \dot{-}_\sigma \varphi = K \setminus \sigma(K \perp\!\!\!\perp \varphi)$.

Theorem 1. (Hansson 1999) A contraction function satisfies success, inclusion, vacuity, uniformity, and core-retainment iff it is a kernel contraction function.

Hansson (1999) has identified that the basic rationality postulates are not strong enough to forbid some irrational behaviours. Consider the following example:

Example 1. Let $K = \{p, p \vee q, p \vee \neg q, \neg p\}$. We want to contract the contradiction \perp . Note that $K \perp\!\!\!\perp \perp = \{A_1, A_2\}$, where $A_1 = \{p, \neg p\}$ and $A_2 = \{p \vee q, p \vee \neg q, \neg p\}$. Let σ be an incision function such that $\sigma(K \perp\!\!\!\perp \perp) = \{p, \neg p\}$. Thus, $K \dot{-}_\sigma \perp = \{p \vee q, p \vee \neg q\}$.

The example above illustrates that kernel contraction operators do not necessarily guarantee a minimal change, since p was removed unnecessarily. To see this, note that $p \vee q$ and $p \vee \neg q$, which were retained, jointly restore p . Hansson (1999) proposed an additional rationality postulate to forbid this spurious behaviour:

(relative closure): $\alpha \in K \dot{-} \varphi$, if $\alpha \in K$ and $K \dot{-} \varphi \models \alpha$.

Relative closure states that if a belief α is removed when contracting a formula φ , then the new epistemic state of the agent cannot entail α . In other words, *relative closure* forbids spurious removal of beliefs. To capture *relative closure*, one can constrain the way that incision functions behave by imposing the following condition:

(smoothness): if $K' \subseteq K$, and $K' \models \psi$ and $\psi \in \sigma(K \perp \perp \varphi)$ then $K' \cap \sigma(K \perp \perp \varphi) \neq \emptyset$.

An incision function that satisfies *smoothness* is called a *smooth incision function*, and a kernel contraction built upon a smooth incision function is called a *smooth kernel contraction*.

Theorem 2. (Hansson 1999) *A kernel contraction $\dot{-}_\sigma$ is smooth iff it satisfies all basic rationality postulates and relative closure.*

A specific kind of contraction is *consolidation* whose purpose is to remove the inconsistencies within a knowledge base. Consolidation corresponds to the special case of contracting the contradiction. Therefore, consolidation can be defined via (smooth) kernel contraction:

Definition 4. *Given a kernel contraction function $\dot{-}_\sigma$ for a knowledge base K , a kernel consolidation function for K is a function $!_\sigma$ such that $K!_\sigma = K \dot{-}_\sigma \perp$. If σ is smooth then $!_\sigma$ is called a smooth consolidation function.*

As consolidation focuses only on removing inconsistency, the original rationality postulates can be simplified:

Theorem 3. (Hansson 1997) *A function $K! = K \dot{-} \perp$ is a kernel consolidation function iff $!_\sigma$ satisfies:*

(consistency), $K! \not\models \perp$; **(inclusion),** $K! \subseteq K$; and **(core-retainment),** if $\psi \in K \setminus K!$, then there is a $K' \subseteq K$ such that $K' \not\models \perp$ and $K' \cup \{\psi\} \models \perp$.

An incision function behaves as an extra-logical mechanism that reveals the epistemic choices that an agent makes in order to remove a piece of information α . As an agent should not decide arbitrarily, its choices should be rationalised by some epistemic preference between its beliefs. This epistemic preference could be defined via a binary relation $<$ on an agent's belief base. Intuitively, $\varphi < \psi$ indicates that φ is at least as entrenched as ψ . Towards this end, an incision function can be "induced" by some binary preference relation that captures the epistemic preferences of an agent among its beliefs.

Definition 5. *Given a knowledge base K , and a binary relation $<: K \times K$. The function $\sigma_<$ is a relational incision function iff (1) $\sigma_<(K \perp \perp \alpha) = \bigcup_{X \in K \perp \perp \alpha} \text{max}_<(X)$, and (2) $\sigma_<$ is an incision function.*

Although not every kernel contraction function can be obtained via relational incision functions (Hansson 1999), we

show that kernel consolidation functions can be obtained via relational incision functions:

Theorem 4. *For every kernel consolidation function $!_\sigma$, there is some binary relation $<$ such that $\sigma(K \perp \perp \perp) = \sigma_<(K \perp \perp \perp)$.*

In the next section, we propose to use culpability measures to unveil an agent's epistemic preference relation.

4 Culpability Measures

A culpability measure (Daniel 2009; Hunter and Konieczny 2006) for a knowledge base K is a function $\mathcal{C}_K: K \rightarrow \mathbb{R}_{\geq 0}$ such that:

(consistency): if $K \not\models \perp$ then $\mathcal{C}_K(\varphi) = 0$ for all $\varphi \in K$;
(blame): if $\varphi \in \bigcup K \perp \perp \perp$ then $\mathcal{C}_K(\varphi) > 0$ for all $\varphi \in K$.

When it is clear from the context, we will omit the subscript K from a culpability measure \mathcal{C}_K . The value 0 is reserved to indicate absence of inconsistency. If a knowledge base is consistent, none of its formulae should be blamed (*consistency*). If a formula participates in yielding inconsistency, its degree of culpability is higher than zero (*blame*).

Let $\text{core}(K, \varphi) = \{X \in K \perp \perp \perp \mid \varphi \in X\}$ be the set of all inconsistent kernels within a knowledge base K that contains a formula φ . We present below some culpability measures from the literature (Hunter and Konieczny 2008):

$$\mathcal{C}_K^D(\varphi) = \begin{cases} 1 & \text{if } \varphi \in \bigcup K \perp \perp \perp \\ 0 & \text{otherwise.} \end{cases}$$

$$\mathcal{C}_K^\#(\varphi) = |\text{core}(K, \varphi)| \quad \mathcal{C}_K^C(\varphi) = \sum_{X \in \text{core}(K, \varphi)} \frac{1}{|X|}$$

The drastic culpability measure \mathcal{C}_K^D distinguishes only two levels: 1, if culpable; 0, if not culpable. The culpability measure $\mathcal{C}_K^\#$ simply counts the number of inconsistent kernels that a formula φ participates in, while \mathcal{C}_K^C considers the size of each inconsistent kernel that a formula is in: directly proportional to the number of inconsistent kernels, and inverse proportional to the size of the inconsistent kernels.

Kernel consolidation functions, as presented in Section 3, depend on an incision function that pinpoints which formulae should be removed, based on a relation that reveals an agent's epistemic preferences. It is not clear, however, from where such an epistemic preference would come, let alone how to construct it. To address this issue, we propose to look at culpability measures as a way to grade an agent's beliefs according to their participation degree in producing inconsistency. Towards this end, a culpability measure works as a tool that unveils the underlying epistemic preference relation of an agent, and a relational incision function could be constructed to perform consolidation. Definition 5 can be adapted in the following way to construct incision functions based on culpability measures:

$$\sigma_{\mathcal{C}_K}(K \perp \perp \perp) = \bigcup_{X \in K \perp \perp \perp} \{\varphi \in X \mid \text{max}_{\mathcal{C}_K}(X) = \mathcal{C}_K(\varphi)\}.$$

Example 2 illustrates how consolidation functions can be constructed via the culpability measures \mathcal{C}_K^C and $\mathcal{C}_K^\#$, and why these measures do not yield smooth incision functions.

Example 2. Consider the inconsistent knowledge base $K = \{p, p \vee q, p \vee \neg q, \neg p \wedge q, \neg p \wedge \neg q, \neg p \wedge r, \neg r\}$ with 8 inconsistent kernels:

$$\begin{aligned} X_1 &= \{p, \neg p \wedge q\} & X_2 &= \{p, \neg p \wedge \neg q\} \\ X_3 &= \{p, \neg p \wedge r\} & X_4 &= \{p \vee q, \neg p \wedge \neg q\} \\ X_5 &= \{p \vee \neg q, \neg p \wedge q\} & X_6 &= \{\neg p \wedge q, \neg p \wedge \neg q\} \\ X_7 &= \{\neg p \wedge r, \neg r\} & X_8 &= \{\neg p \wedge r, p \vee q, p \vee \neg q\} \end{aligned}$$

The culpability measures $\mathcal{C}_K^\#$ and \mathcal{C}_K^C assigns the following values to the formulæ in K :

	p	$p \vee q$	$p \vee \neg q$	$\neg p \wedge q$	$\neg p \wedge \neg q$	$\neg p \wedge r$	$\neg r$
$\mathcal{C}_K^\#$	3	2	2	3	3	3	1
\mathcal{C}_K^C	3/2	5/6	5/6	3/2	3/2	4/3	1/2

Thus, $\sigma_{\mathcal{C}_K^\#}(K \perp \perp) = \sigma_{\mathcal{C}_K^C}(K \perp \perp) = \{p, \neg p \wedge q, \neg p \wedge \neg q, \neg p \wedge r\}$. Note that $\sigma_{\mathcal{C}_K^\#}$ and $\sigma_{\mathcal{C}_K^C}$ are not smooth, because $p \in \sigma_{\mathcal{C}_K^\#}(K \perp \perp) = \sigma_{\mathcal{C}_K^C}(K \perp \perp)$, the set $X = \{p \vee q, p \vee \neg q\} \models p$, but $X \cap \sigma_{\mathcal{C}_K^\#}(K \perp \perp) = X \cap \sigma_{\mathcal{C}_K^C}(K \perp \perp) = \emptyset$.

Observation 5. The culpability measures $\mathcal{C}_K^\#$ and \mathcal{C}_K^C do not induce smooth incision functions for every knowledge base.

Proposition 6. The culpability measure \mathcal{C}_K^D produces only the severe consolidation: $K! = K \setminus \bigcup K \perp \perp$.

Although the drastic culpability measure induces a smooth incision function, the corresponding consolidation function is too severe, as it removes all formulæ that produce inconsistency. This happens because \mathcal{C}_K^D is unable to distinguish different levels of culpabilities. In order to devise more robust consolidation functions from culpability measures, we define in the next section the concept of tacit functions, which will be used to devise new culpability measures and establish a representation theorem between smooth consolidation functions and culpability measures. This supports the idea that culpability measures can work as a general framework to provide epistemic preference relations.

5 Tacit Culpability Measures

Most of the culpability measures discussed in Section 4 judge the degree of culpability of a formula by quantifying its explicit participation in entailing inconsistency. For instance, the knowledge base $\{p, p \vee q, p \vee \neg q, \neg p\}$ has only two sources of inconsistencies: kernels $A_1 = \{p, \neg p\}$ and $A_2 = \{p \vee q, p \vee \neg q, \neg p\}$. The culpability measure $\mathcal{C}^\#$ assigns value 1 to p because it appears only in the kernel A_1 . A closer look shows us that the formula p contributes to inconsistency both explicitly and implicitly. Explicitly, because it conflicts with $\neg p$ in A_1 , and implicitly because in A_2 the reason for inconsistency is that the formulæ $p \vee q$ and $p \vee \neg q$ jointly imply p which in turn conflicts with $\neg p$. As in both cases, p appears as the central reason of inconsistency by conflicting with $\neg p$, it would be plausible to assign to p the same inconsistency value of $\neg p$, or at least take into account the implicit role of p in yielding inconsistency in A_2 . If one

can measure such an implicit role of p by a value x , then a more accurate culpability value for p would be to merge both explicit and implicit degrees of inconsistency, that is, $1 + x$.

The above recipe can be generalised to any culpability measure. Assume we agree upon a culpability measure \mathcal{C} that considers how explicit the formulæ in a knowledge base contributes to making it inconsistent, such as $\mathcal{C}^\#$ and \mathcal{C}^C . If we define a function τ that discloses the implicit degree of inconsistencies of the formulæ in a knowledge base, then the culpability measure that takes explicit and implicit degree of inconsistency can be defined via $\mathcal{C}^\circ(\varphi) = \mathcal{C}(\varphi) + \tau(\varphi)$. A function τ that discloses the implicit degree of culpability will be called a *tacit function*, while a culpability measure \mathcal{C}° that assembles both explicit and implicit degrees of culpability will be called a *tacit culpability measure*.

Tacit functions will be used to construct more accurate culpability measures that combine both explicit and implicit degrees of culpability. But first we need to formally define what it means for a formula to *implicitly* contribute in making a knowledge base inconsistent. For this, we first properly define what it means in general for a formula to participate in making a knowledge base inconsistent, and then we define the special cases of explicit versus implicit contributions of inconsistency.

Definition 6. Let $X_\varphi^+ = \{A \in X \perp \perp \mid A \not\models \perp\}$ denote the set of all consistent φ -kernels of X .

1. A φ -witness within K is an inconsistent set $X \subseteq K$ such that (i) $(X \setminus \bigcup X_\varphi^+) \not\models \perp$ and (ii) $(X \setminus \bigcup X_\varphi^+) \cup \{\varphi\} \models \perp$.
2. A formula φ participates in making K inconsistent iff φ has some witness within K .
3. A formula φ **explicitly** makes K inconsistent iff φ has some witness X such that $\varphi \in X$.
4. A formula φ **implicitly** makes K inconsistent iff φ has some witness X such that $\varphi \notin X$.

Intuitively φ contributes in making K inconsistent if in some subset X of K , the information that a formula φ carries is essential in producing inconsistency. This is formalised by removing from X all formulæ that consistently imply φ and replacing them by φ itself. If such a removal of formulæ produces a consistent set, condition 1.(i), but the insertion of φ brings back the inconsistency, condition 1.(ii), then φ is indeed involved in producing inconsistency. In our previous example, the inconsistent kernel $A_2 = \{p \vee q, p \vee \neg q, \neg p\}$ is a p -witness, because the only consistent p -kernel within A_2 is the set $A' = \{p \vee q, p \vee \neg q\}$ and (1.i) $(A_2 \setminus A') \not\models \perp$, and (1.ii) replacing A' by φ results in $A_1 = \{p, \neg p\}$ which indeed produces inconsistency, that is, $(A_2 \setminus A') \cup \{p\} = A_1 \models \perp$. Moreover, p implicitly contributes to the inconsistency because $p \notin A_2$. Note that the notion of an implicit contribution to inconsistency, unlike the explicit counterpart, does not require a formula φ to be in the knowledge base K . Rather, it focus on the consistent information entailed by φ that conflicts with consistent information within K .

Tacit functions will be explored in the next section as a means to produce culpability measures capable of inducing smooth consolidation function.

5.1 Constructing Rational Consolidation via Tacit Culpability Measures

In this section, we look at culpability measures as a way of disclosing an epistemic preference of an agent in order to perform consolidation. We identify a sufficient condition that guarantees a culpability measure to yield rational consolidation functions:

tacit dominance If $A \in \text{core}(K, \psi)$, $A' \in K \perp \perp \psi$, and $X \in ((A \setminus \{\psi\}) \cup A') \perp \perp$ then there is a $\varphi \in X \cap A'$ such that $\mathcal{C}_K(\varphi) \geq \mathcal{C}_K(\psi)$.

The postulate *tacit dominance* addresses the concept of implicit contributions towards inconsistency, and relates tacit and explicit degrees of culpability. It states that if a formula ψ participates in yielding inconsistency implicitly, then some formula that contributes to the implicit role of ψ in an inconsistent kernel X should be at least as culpable as ψ .

Theorem 7. *Given a culpability measure \mathcal{C}_K and an incision function $\sigma_{\mathcal{C}_K}$. If \mathcal{C}_K satisfies tacit dominance then the consolidation function $!_{\sigma_{\mathcal{C}_K}}$ is smooth.*

Proof sketch. Let $K' \subseteq K$ and $\psi \in \sigma_{\mathcal{C}_K}(K \perp \perp)$ s.t. $K' \models \psi$. Thus, there are $A \in \text{core}(K, \psi)$ and $A' \in K \perp \perp \psi$ s.t. $\max_{\mathcal{C}_K}(A) = \mathcal{C}_K(\psi)$. Let $Y = (A \setminus \{\psi\}) \cup A'$, and $X \in Y \perp \perp$ (note that $Y \models \psi$). From *tacit-dominance*, there is a $\varphi \in X \cap A'$ s.t. $\mathcal{C}_K(\varphi) \geq \mathcal{C}_K(\psi)$. Thus, $\max_{\mathcal{C}_K}(A) \leq \mathcal{C}_K(\varphi) \leq \max_{\mathcal{C}_K}(X)$. This implies $\max_{\mathcal{C}_K}(X) = \mathcal{C}_K(\alpha)$, for some $\alpha \in A'$. Hence, σ is smooth: $K' \cap \sigma(K \perp \perp) \neq \emptyset$. \square

Theorem 7 states that it is enough for a culpability measure to satisfy *tacit dominance* in order to induce a smooth consolidation function. In Subsection 5.2, we explore *tacit dominance* as a means to define a new class of culpability measures that yield smooth incision functions: the *stable tacit culpability measures* (Definition 8). This new class exploits the concept of tacit functions introduced in the previous section. In the remainder of this section, we construct the auxiliary tools necessary for defining the stable tacit culpability measures. For sake of clarity, we provide examples to support the intuition of how these auxiliary tools work.

We give a brief overview of how these auxiliary tools will be used to define stable tacit culpability measures. The auxiliary tools assume an underlying culpability measure \mathcal{C}_K that identifies explicit culpability degrees. The core ingredient is the function $\nu_{\mathcal{C}_K}$ that identifies some of the tacit values within a knowledge base K . In Section 5.2, $\nu_{\mathcal{C}_K}$ is used in an iterative approach to identify all the tacit values within K . These tacit values are realised by exploiting a strategy of maximising (via a function $\mu_{\mathcal{C}_K}$) and minimising (via a function $\delta_{\mathcal{C}_K}$) distances between the explicit culpability values provided by \mathcal{C}_K . The notion of distance we conceptualise here depends on other elements we introduce. The first ingredient is the function $\lambda_{\mathcal{C}_K}$ that, for each formula φ involved in producing inconsistency, identifies the most significant culpability value related to φ . This significant value is called φ 's *pole of inconsistency*. The intuition is that the closer a formula φ is to its pole of inconsistency, the more influential is φ in producing inconsistency. From this concept of influence degree, we define a function Δ that establishes a

notion of distance between the influence degrees of two formulae in producing inconsistency. This notion of distance is then further generalised by two other functions: $d_{\mathcal{C}_K}(\varphi, \psi)$ that localises the distance between φ and ψ within the inconsistent kernels that φ is; and $\delta_{\mathcal{C}_K}(X, \psi)$ that minimises the distances between a single formula ψ and the formulae within a set X of formulae. A function $\mu_{\mathcal{C}_K}(\varphi, \psi)$ is defined to maximise the distances between ψ and all φ -kernels that contributes to the implicit role of ψ in yielding inconsistency. The intuition is that the more φ -kernels support ψ in yielding inconsistency, the higher is the tacit value of ψ .

We start by defining the notion of a *pole of inconsistency*, which will be helpful in inducing tacit functions from explicit degrees of culpability. Not all sources of inconsistencies contribute equally towards inconsistency, and the higher the culpability values of the formulae in an inconsistent kernel, the more influential are such formulae in producing inconsistencies. Within an inconsistent kernel X , a formula ψ with the highest culpability value is deemed the most culpable one in producing the inconsistency. The closer the culpability value of a formula $\varphi \in X$ approaches to the culpability value of ψ , the more influential is φ within X in producing inconsistency. Let us call ψ a pivot of φ . Intuitively, a pivot of a formula φ can be understood as a significant formula in producing inconsistency in the presence of φ . The pivots of each formula is given by

$$\text{piv}_{\mathcal{C}_K}(\varphi) = \bigcup_{X \in \text{core}(\varphi, K)} \{\psi \in X \mid \mathcal{C}_K(\psi) = \max_{\mathcal{C}_K}(X)\}.$$

A formula φ might have more than one pivot. The closer the culpability value of φ is to the culpability value of a pivot, more influential is φ in producing inconsistency. Therefore, φ is more influential exactly in the inconsistent kernels where φ and its least culpable pivots are. The pole of inconsistency of a formula φ describes the significant culpability value where φ maximises its influence in yielding inconsistency. If φ is itself one of its pivots, then the culpability value of φ is its own pole of inconsistency. The pole of inconsistency of φ is therefore the minimal culpability value of all its pivots. Given a culpability measure \mathcal{C}_K , the pole of inconsistency of each formula involved in producing inconsistency within K is given by the function $\lambda_{\mathcal{C}_K} : (\bigcup K \perp \perp) \rightarrow \mathbb{R}_{\geq 0}$,

$$\lambda_{\mathcal{C}_K}(\varphi) = \min_{\mathcal{C}_K}(\text{piv}_{\mathcal{C}_K}(\varphi)).$$

Observation 8. *Given a culpability measure \mathcal{C}_K , if $\text{core}(K, \varphi) \neq \emptyset$ then $\lambda_{\mathcal{C}_K}(\varphi) \geq \mathcal{C}_K(\varphi)$.*

Example 3 below illustrates the intuition of the pole of inconsistency.

Example 3. Let $K_1 = \{p, p \vee r, p \vee \neg r, \neg p, r, (p \vee r) \rightarrow z, z \leftrightarrow r, z, \neg z\}$, and a culpability measure \mathcal{C}' that ranks the culpability of these formulae between 1 and 5 as depicted in Figure 1. The formula r has exactly 3 inconsistent kernels: $A_1 = \{(p \vee r) \rightarrow z, r, \neg z\}$, $A_2 = \{r, z \leftrightarrow r, \neg z\}$ and $A_3 = \{p \vee \neg r, r, \neg p\}$ which are highlighted in Figure 1. Note that the $\neg z$ is the most culpable formula in both A_1 and A_2 , while $\neg p$ is the most culpable formula in A_3 . The culpability value of r is much closer to the culpability value of $\neg z$ rather than $\neg p$. This means that the culpability value of r is more

influential in the presence of $\neg z$, which makes the culpability value of $\neg z$ the pole of inconsistency of r . Table 1 below presents the poles of inconsistency of the formulae in K_1 :

	p	$p \vee r$	$p \vee \neg r$	$\neg p$	r	$(p \vee r) \rightarrow z$	$z \leftrightarrow r$	z	$\neg z$
λ_{C_K}	5	4	5	5	4	4	4	4	4

Table 1: Poles of inconsistency for K_1 .




Figure 1: A graphical overview of how the culpability values of the formulae in Example 3 are ranked. The inconsistent kernels of r are outlined.

Lemma 9. Given a culpability measure C_K , if there exists a set $A \in K \perp \perp$, and two formulae $\varphi, \psi \in A$ such that $C_K(\psi) = \max_{C_K}(A)$ then: (1) $\lambda_{C_K}(\psi) = C_K(\psi)$, and (2) $\lambda_{C_K}(\varphi) \leq \lambda_{C_K}(\psi)$.

The pole of inconsistency helps us to grade the formulae based on how influential they are in producing inconsistencies: the closer the culpability value of a formula is to its pole of inconsistency, the more influential it is. Towards this end, the *influence degree* of a formula φ can be measured as the ratio between its culpability value and its pole of inconsistency, that is, $C_K(\varphi)/\lambda_{C_K}(\varphi)$. Figure 2 depicts the formulae of Example 3 graded by its influence degrees. The distance between the influence degrees of two formulae can be measured by the signed distance function¹ $\Delta_{C_K} : (\bigcup K \perp \perp) \times (\bigcup K \perp \perp) \rightarrow \mathbb{R}$ such that

$$\Delta_{C_K}(\varphi, \psi) = \frac{C_K(\psi)}{\lambda_{C_K}(\psi)} - \frac{C_K(\varphi)}{\lambda_{C_K}(\varphi)}.$$

The function Δ_{C_K} not only measures how far the influence degrees of two formulae are from each other, but also indicates which one is more influential than the other: $\Delta_{C_K}(\varphi, \psi) < 0$ indicates that φ is more influential than ψ , while $\Delta_{C_K}(\varphi, \psi) \geq 0$ indicates that ψ is at least as influential as φ . For instance, in Example 3, the formulae r and $p \vee \neg r$ have the same culpability value 3. However, as shown in Figure 2, r is more influential than $p \vee \neg r$ because the culpability value of r is much closer to its pole of inconsistency than $p \vee \neg r$ is to its pole of inconsistency (see Table 1 for the pole of inconsistencies). On the other hand, $z \leftrightarrow r$ and z are equally influential, because their culpability values are

¹The term distance is often used a synonym for metric which should, among other properties, be a non-negative real function. Here, instead, the term distance is used to capture the notion of fairness between culpability values of formulae and negative values are used to give a notion of oriented distance.




Figure 2: Formulae from Example 3 ranked according to their influence degrees.

not only equal, but they also present the same pole of inconsistency 4.

A formula produces inconsistency only within its inconsistent kernels, or core for short. The distance, given by Δ_{C_K} , between two formulae should then be localised within their cores.

If one wants to localise the distance between the influence degrees of two formulae φ and ψ within one of their cores, say φ , then it suffices to multiply the distance by the pole of inconsistency of φ . This localisation is realised by the function $d_{C_K} : (\bigcup K \perp \perp) \times (\bigcup K \perp \perp) \rightarrow \mathbb{R}_{\geq 0}$,

$$d_{C_K}(\varphi, \psi) = \Delta(\varphi, \psi) \cdot \lambda_{C_K}(\varphi).$$

A formula φ depends on other formulae in order to implicitly contribute to entail inconsistency. For instance, the formula p in the knowledge base $\{p, p \vee q, p \vee \neg q, \neg p\}$ implicitly entails inconsistency only because in the kernel $\{p \vee q, p \vee \neg q, \neg p\}$, the formulae $p \vee q$ and $p \vee \neg q$ jointly entail p . We say that both $p \vee q$ and $p \vee \neg q$ contribute to the implicit role of p in yielding inconsistency; or simply that they contribute to the implicit role of p towards inconsistency. It will be helpful to have a function, that given a formula φ , identifies which formulae ψ the formula φ potentially contributes to the implicit role of ψ towards inconsistency:

$$\text{imp}(K, \varphi) = \{\psi \in \bigcup K \perp \perp \mid \varphi \in \bigcup K \perp \perp \psi\}.$$

We will define a tacit function that uses the influence degrees to determine the tacit value of a formula.

Given a formula φ that contributes to entail inconsistency within a knowledge base K , its tacit value could be determined by how strongly connected its influence degree is to the influential degrees of the formulae in $\text{imp}(\varphi, K)$. The function d_{C_K} will be essential in this task. Let us start by defining the function $\delta_{C_K} : 2^K \times (\bigcup K \perp \perp) \rightarrow \mathbb{R}_{\geq 0}$ such that

$$\delta_{C_K}(X, \psi) = \min(\{d_{C_K}(\varphi, \psi) \mid \varphi \in X \cap (\bigcup K \perp \perp)\}).$$

The function borrows intuition from the pole of inconsistency. Back to Example 3, the set $X_1 = \{r, p \vee \neg r\}$ entails p . The distance from the influence degrees of r and $p \vee \neg r$ to the influential degree of p can be visualised in Figure 2. While both r and $p \vee \neg r$ have the same culpability value, the influence degree of r is closer than the influence degree of $p \vee \neg r$ to the influence degree of p . This means that r contributes more than $p \vee \neg r$ to the implicit role of p in

producing inconsistency. Therefore, the closer the influence degree of a formula φ is to the influence degree of a formula ψ , the more φ contributes to the implicit role of ψ in producing inconsistency. The distance Δ provides the raw distance between the influence degrees of two formulae, but does not provide any context of how this distance is interpreted in their cores. The distance $\Delta(\psi, \varphi)$ can be projected towards φ to weight the connection between φ and ψ , which can be done by using d_{C_K} instead of Δ .

Intuitively, $\delta_{C_K}(X, \psi)$ minimises the distances between ψ and the formulae in X that participate in producing inconsistency. Back to Example 3, for $X_1 = \{r, p \vee \neg r\}$, we have $\delta_{C_K}(X_1, p) = \min(\{d_{C_K}(r, p), d_{C_K}(p \vee \neg r, p)\}) = \min(\{1, 2\}) = 1$. The concept of influence degree is applied only to formulae that contribute in yielding inconsistency, therefore the condition $\varphi \in X \cap (\bigcup K \perp \perp)$ is used to guarantee that only such formulae will be considered.

If a formula φ contributes to the implicit role of a formula ψ in yielding inconsistency via only one ψ -kernel, then δ_{C_K} can be used to assess how much φ contributes to the implicit role of ψ . However, φ may contribute to the implicit role of ψ in producing inconsistency via more than one ψ -kernel. The more ψ -kernel φ participates in, the farther φ potentially becomes from ψ . Therefore, the tacit value of φ should maximise the value of δ_{C_K} to each of the ψ -kernels that φ participates in. This is accomplished by the function:

$$\mu_{C_K}(\varphi, \psi) = \max(\{\delta_{C_K}(X, \psi) \mid X \in K \perp \perp \psi, \varphi \in X\}).$$

Example 4. (Continued from Example 3). The formula $p \vee \neg r$ contributes to entail p via two kernels: $X_1 = \{r, p \vee \neg r\}$ and $X_2 = \{p \vee r, p \vee \neg r\}$. The implicit connection between $p \vee \neg r$ and p can be measured by $\mu_{C_K}(p \vee \neg r, p)$. Note that $\delta_{C_K}(X_2, p) = 2$, and $\delta_{C_K}(X_1, p) = 1$. Thus,

$$\mu_{C_K}(p \vee \neg r, p) = \max(\{\delta_{C_K}(X_1, p), \delta_{C_K}(X_2, p)\}) = 2.$$

The function μ is used to assess how much a formula φ contributes to the implicit role towards inconsistency of a formula ψ . However, a formula φ might contribute to the implicit role towards inconsistency of more than one formula. Let us define the function $\nu_{C_K} : K \rightarrow \mathbb{R}_{\geq 0}$ that considers all the formulae that a formula φ contributes to the implicit role towards inconsistency:

- $\nu_{C_K}(\varphi) = 0$, if $\text{core}(K, \varphi) = \emptyset$;
- $\nu_{C_K}(\varphi) = \max(\{\mu_{C_K}(\varphi, \psi) \mid \psi \in \text{imp}(\varphi, K)\})$, otherwise.

Intuitively, ν_{C_K} approximates the culpability value of φ to its own pole of inconsistency taking the distances between the culpability value of φ and each formula ψ it contributes to the implicitly role towards inconsistency. If a formula does not contribute to make K inconsistent, then $\nu_{C_K}(\varphi) = 0$.

Example 5. (Continued from Example 4) Let us look at how implicitly connected is the formula $p \vee r$ in producing inconsistencies according to $\nu_{C'}$. The formula $p \vee r$ participates in entailing only three formulae that contributes in making K_1 inconsistent: p , z and r via the kernels $X_2 = \{p \vee r, p \vee \neg r\}$, $X_3 = \{(p \vee r) \rightarrow z, p \vee r\}$ and $X_4 = \{p \vee r, p \vee r \rightarrow z, z \leftrightarrow r\}$, respectively. Thus, $\text{imp}(K_1, p \vee r) = \{p, z, r, p \vee r\}$,

	p	$p \vee r$	$p \vee \neg r$	$\neg p$	r	$(p \vee r) \rightarrow z$	$r \leftrightarrow z$	z	$\neg z$
$\nu_{C'}$	0	2	2	0	1	0	2	2	0
C'°	5	3	5	5	4	2	4	4	4

Table 2: Functions $\nu_{C'}$ and C'° on formulae from Example 3.

and $\nu_{C'}(p \vee r) = \max(\{\mu_{C'}(p \vee r, p), \mu_{C'}(p \vee r, z), \mu_{C'}(p \vee r, r), \mu_{C'}(p \vee r, p \vee r)\})$. From the previous example, we have $\delta_{C'}(X_2, p) = 2$. Note that $\delta_{C'}(X_3, z) = \min(\{d_{C'}(p \vee r \rightarrow z, z), d_{C'}(p \vee r, z)\}) = \min(\{0, 1\}) = 0$; and $\delta_{C'}(X_4, r) = \min(\{d_{C'}(p \vee r, r), d_{C'}(p \vee r \rightarrow z, r), d_{C'}(z \leftrightarrow r, r)\}) = \min(\{2, 1, 1\}) = 1$. Thus,

$$\mu_{C'}(p \vee r, p) = \max(\{\delta(X_2, p)\}) = 2$$

$$\mu_{C'}(p \vee r, z) = \max(\{\delta(X_3, z)\}) = 0$$

$$\mu_{C'}(p \vee r, r) = \max(\{\delta(X_4, r)\}) = 1$$

$$\mu_{C'}(p \vee r, p \vee r) = \max(\{\delta(\{p \vee r\}, p \vee r)\}) = 0.$$

Thus, $\nu_{C'}(p \vee r) = \max(\{0, 1, 2\}) = 2$. Table 2 summarises $\nu_{C'}$.

Proposition 10. For every culpability measure C_K , $0 \leq \nu_{C_K}(\varphi) \leq \lambda_{C_K}(\varphi) - C_K(\varphi)$, for all $\varphi \in \bigcup K \perp \perp$.

According to Proposition 10, the values identified by ν_{C_K} are always non-negative and bounded by the poles of inconsistency. We can use ν_{C_K} to construct a culpability measure that combines both the explicitly and implicitly counterparts. Let C_K be a culpability measure, we define

$$C_K^{\circ}(\varphi) = C_K(\varphi) + \nu_{C_K}(\varphi).$$

The function C_K° is indeed a culpability measure.

Observation 11. For every culpability measure C_K , the function C_K° built upon C_K is also a culpability measure.

Table 2 shows the culpability measure C'° built upon the culpability measure C' from Example 3.

5.2 A Tacit Function Based on Fixed Points

Let us see how ν_{C_K} and C_K° connect with tacit dominance:

Example 6. (Continued from Example 5) Let us compare C' and C'° . Note that C' is not tacit dominant because $C'(p) = 5$, while the formulae in the p -kernels $X_1 = \{p \vee \neg r, r\}$ and $X_2 = \{p \vee r, p \vee \neg r\}$ have culpability value strictly less than 5. Using $\nu_{C'}$ to identify the tacit values within K_1 and combining them with C' , produces the culpability measure C'° that grades at least one formula in both X_1 and X_2 as culpable as p . Precisely, due to the tacit values identified by $\nu_{C'}$, the formulae $p \vee \neg r$ has the same culpability value of p , that is, $C'(p) = C'(p \vee \neg r) = 5$. At a first glance, this shows the potential of $\nu_{C'}$ in capturing tacit dominance. On the other hand, C' is not tacit dominant, because $C'(z) = 4$, while all the formulae in the z -kernel $X_5 = \{p \vee r, p \vee r \rightarrow z\}$ have culpability value strictly less than 4.

Example 6 shows us that ν_{C_K} is not strong enough to capture tacit dominance. A closer look, however, will show us why ν_{C_K} fails tacit dominance and how it could be modified in order to achieve tacit dominance. Precisely, ν_{C_K} focuses

on tacit values of formulæ directly connected to the most influential formulæ. For instance, the formula z jointly with $z \leftrightarrow r$ and $p \vee \neg r$ implies p whose culpability value coincides with its pole of inconsistency 5. On the other hand, the formula $\alpha = p \vee r \rightarrow z$ had no tacit value identified, because the only formulæ it consistently contributes to entail are r and z which are originally not among the most influential formulæ, as their influence degree are respectively $3/4$ and $1/2$ (the influence degrees are depicted in Figure 2). However, z has a tacit value of 2, due do its close connection with p , putting z among the most influential formulæ in K_1 : $C^{\circ}(z) = 2+2 = 4$ and $\lambda_{C^{\circ}}(\alpha) = 4$ which implies in an influence degree of 1. But since $\{p \vee r, \alpha\}$ implies z , which is indeed among the most influential ones due to its tacit value, then α should indeed have some tacit value as well, which was not identified by $\nu_{C^{\circ}}$. The reason why $\nu_{C^{\circ}}$ was unable to identify the tacit value of α is because $\nu_{C^{\circ}}$ is not sensitive to the realisations that the tacit value of a formula, say z , would interfere with the tacit value of another formula, say α .

We can then extend ν_{C_K} to become sensitive to these shifting of values as the tacit values of other formulæ are identified. We can achieve this by resorting to an iterative strategy: ν_{C_K} is used to identify the tacit values of some formulæ, and assemble them with the original culpability value, resulting in C_K° ; we then apply $\nu_{C_K^{\circ}}$ to identify remaining tacit values. We put this on the loop until no more tacit values are identified. To avoid summing up overlapping tacit values in the iterative strategy, we can restrict ν_{C_K} to focus only on some specific formulæ whose tacit values will be fully identified in each iteration. In the beginning of our discussion, we have argued that ν_{C_K} focuses mainly on the formulæ directly connected to the most influential formulæ. Thus, we can prioritise these formulæ in each iteration. For convenience, let us slightly modify ν_{C_K} to certify that in each iteration it will focus only on the formulæ closest to the most influential ones:

- $\nu_{C_K}^+(\varphi) = \nu_{C_K}(\varphi)$, if there exists $\psi \in \text{imp}(\varphi, K)$ such that $C_K(\psi) = \lambda_{C_K}(\psi)$ and $\mu_{C_K}(\varphi, \psi) > 0$;
- $\nu_{C_K}^+(\varphi) = 0$, otherwise.

For the iterative approach, we define

$$\begin{aligned} \mathcal{T}_{C_K}^0(\varphi) &= C_K(\varphi) + \nu_{C_K}^+(\varphi) \\ \mathcal{T}_{C_K}^{i+1}(\varphi) &= \mathcal{T}_{C_K}^i(\varphi) + \nu_{\mathcal{T}_{C_K}^i}^+(\varphi), \text{ for } 0 \leq i \in \mathbb{Z}. \end{aligned}$$

The base case $\mathcal{T}_{C_K}^0$ applies $\nu_{C_K}^+$ to identify the tacit values of the formulæ in the most influential subsets and combine them with C_K . The function ν^+ is applied then on the values obtained by $\mathcal{T}_{C_K}^0$ to identify the tacit values of the remaining formulæ, in the knowledge base. This process repeats until all tacit values are found. The first issue is to show that this iterative approach indeed converges to a fixed point, that is, after a number m of iterations no tacit values are produced. This can be formalised in the following way.

Definition 7. *Given a culpability measure C_K , a fixed point of \mathcal{T}_{C_K} is a $m \in \mathbb{Z}^+$ such that $\mathcal{T}_{C_K}^m \equiv \mathcal{T}_{C_K}^{m+1}$. If $0 \leq m - 1$ is not a fixed point, or $m = 0$ then m is the least fixed point.*

We will first show that \mathcal{T}_{C_K} has at least one fixed point. For this we will need the following lemma.

Lemma 12. *If $\mathcal{T}_{C_K}^{i+1} \neq \mathcal{T}_{C_K}^i$ then there is some $\alpha \in K$ such that $\mathcal{T}_{C_K}^{i+1}(\alpha) > \mathcal{T}_{C_K}^i(\alpha)$ and $\mathcal{T}_{C_K}^j(\alpha) = \mathcal{T}_{C_K}^{i+1}(\alpha)$, for all $j \geq i + 1$.*

Lemma 12 informs that in each iteration step, the tacit value of some formula α is fully identified, in the sense that in the further iterations its tacit value is not sensitive to the tacit value of other formulæ anymore. From this it is easy to see that after a number of iterations, a fixed point is reached.

Theorem 13. *For every culpability measure C_K , the function \mathcal{T}_{C_K} has a least fixed point.*

We can then define a culpability measure based on the fixed point of \mathcal{T}_{C_K} .

Definition 8. *Given a culpability measure C_K , the tacit culpability measure ν_{C_K} built upon C_K and let m be the least fixed point of \mathcal{T}_{C_K} . The function $\eta_{C_K}(\varphi) = \mathcal{T}_{C_K}^m(\varphi)$ is the stable tacit function of C_K .*

Example 7. (Continued from Example 6). Note that $\mathcal{T}_{C'}^0 \equiv C^{\circ}$, which we already have presented on Table 2. Thus,

$$\mathcal{T}_{C'}^1(\varphi) = \mathcal{T}_{C'}^0(\varphi) + \nu_{\mathcal{T}_{C'}^0}^+(\varphi) = C^{\circ}(\varphi) + \nu_{C^{\circ}}^+(\varphi).$$

One of the formulæ that still presents some tacit value is $\alpha = (p \vee r) \rightarrow z$. Solving $\nu_{\mathcal{T}_{C'}^0}^+(\alpha)$ and simplifying some equations will give us

$$\nu_{\mathcal{T}_{C'}^0}^+(\alpha) = \mu_{\mathcal{T}_{C'}^0}(\alpha, r) = \delta_{\mathcal{T}_{C'}^0}(\{\alpha, p \vee r\}, r) = 1.$$

Thus, $\mathcal{T}_{C'}^1(\alpha) = C^{\circ}(\alpha) + \nu_{\mathcal{T}_{C'}^0}^+(\alpha) = 2 + 1 = 3$. The table below summarises $\mathcal{T}_{C'}^1$. One can check that $\mathcal{T}_{C'}^2 \equiv \mathcal{T}_{C'}^1$, that is, the least fixed point of $\mathcal{T}_{C'}$ is 1.

	p	$p \vee r$	$p \vee \neg r$	$\neg p$	r	$p \vee r \rightarrow z$	$r \leftrightarrow z$	z	$\neg z$
$\nu_{\mathcal{T}_{C'}^1}^+$	0	1	0	0	0	1	0	0	0
$\mathcal{T}_{C'}^1$	5	4	5	5	4	3	4	4	4

Observation 14. *For every culpability measure, the stable tacit function η_{C_K} is a culpability measure.*

The function η_{C_K} indeed yields smooth consolidation functions:

Theorem 15. *Given a stable tacit culpability measure η , the consolidation function $!_{\sigma_\eta}$ is smooth.*

It turns out that stable tacit culpability measures not only induce rational consolidation functions, but also every smooth consolidation function can be constructed via a stable tacit culpability measure

Theorem 16. *If a consolidation function $!_{\sigma}$ is smooth, then there is some stable tacit culpability measure η such that $!_{\sigma_\eta} \equiv !_{\sigma}$.*

Actually, the representation theorem between stable tacit culpability measures and consolidation functions, obtained jointly by Theorems 15 and 16, is even stronger: the stable tacit culpability measures form the only class of culpability measures capable of inducing smooth consolidation functions.

Theorem 17. *If C is a culpability measure and $!_{\sigma_C}$ is smooth, then C is a stable tacit culpability measure.*

6 Discussion and Future Works

In this paper, we proposed to use culpability measures as a tool to disclose an agent’s epistemic preference relation and use it to perform consolidation. The existing culpability measures in the literature, however, fail to yield consolidation functions that satisfy the relative closure postulate, known as smooth consolidation functions. To address this problem, we have introduced the concept of tacit functions that take the semantic counterparts to define more precise culpability measures. In this direction, we have introduced a special class of tacit culpability functions, the *stable tacit culpability measures*, that use an initial culpability measure that assess the explicit degree of culpability and then realises the tacit values of the formulæ. We showed then a representation theorem between stable tacit culpability measures and smooth consolidation functions. In addition, the representation theorem is strengthened by showing that the stable tacit culpability measures are the only class of culpability measures that yield smooth consolidation functions. In the remainder of this section, we discuss about some future research direction worth to take.

Belief Change and Culpability Measures A future research path worth to explore is to use culpability measures, in particular the stable tacit culpability measures, to support other kinds of belief change operations:

(1) **Belief contraction and safe contraction:** assume that a formula α should be relinquished from a knowledge base K . The formulæ within K could be graded according to their contribution degree in entailing α . Respectively, one could grade the formulæ within K according to their conflicting degree with $\neg\alpha$. The conflicting degree between the formulæ within K and $\neg\alpha$ could be assessed by applying a culpability measure $\mathcal{C}_{K'}$ at $K' = K \cup \{\neg\alpha\}$. The contraction could then be performed by a (smooth) kernel contraction $K \dot{-}_{\mathcal{C}_{K'}} \alpha$ whose incision function is induced by \mathcal{C}_K . It is still not clear if, in this strategy, tacit-dominance would be enough to guarantee relative closure and/or which extra properties a culpability measure would need to satisfy to capture relative-closure. In addition, it would be necessary to investigate how well behaved are these contraction functions based on this strategy. Kernel contraction functions whose incision functions are based on epistemic preference relations are known as safe contraction (Alchourrón and Makinson 1985; Rott and Hansson 2014). Therefore, the use of culpability measures, as proposed, to construct kernel contraction functions, would bring to light a new kind of safe contraction whose rationality would be worth to explore, as well as the role of tacit dominance in this setting.

(2) **Belief revision:** In a knowledge base K , revision by a formula α , denoted $K * \alpha$, can be performed via Levi’s external identity (Hansson 1993b): $(K \cup \{\alpha\}) - \neg\alpha$, that consists in first adding α to the knowledge base, potentially making it inconsistent, and then removing the information conflicting with α . In this case, a culpability measure could be applied in $(K \cup \{\alpha\})$ to identify the most culpable formulæ and use it as an epistemic preference relation to construct a (smooth) contraction function as discussed in the topic (1) above.

(3) **Transmutation:** when a belief change operation is applied on a knowledge base, like contraction and revision, the agent’s epistemic preference relation also changes. This process of modifying the epistemic preference relation is known as transmutation (Williams 1994, 1995). Assume that the epistemic preference relation of an agent could be given by a function f that informs the entrenched degree of its beliefs. In our approach, the tacit function ν_f would inform how strongly connected are the formulæ in producing inconsistency. The minimise/maximise approach we have proposed to unveil these hidden values could be adapted to identify how strongly connected are the formulæ in entailing a formula α to be contracted/revise. In this case, the obtained tacit values could be used to define the new epistemic preference relation in the new knowledge base.

Pseudo-Consolidation and “Conflicts Under the Water”

De Bona and Hunter (2017) identified cases where sub-formulæ of a knowledge base might not participate in any of the minimal inconsistent subsets, and yet be responsible in entailing inconsistencies. To deal with this problem, they have introduced the concept of “conflicts under the water” which is similar to our concept of implicit contribution to inconsistency. De Bona and Hunter’s strategy consists in defining some consequence operators Cn^* (which we shall call here a pseudo consequence operator) that augment a knowledge base with some formulæ in order to reveal hidden sub-formulæ that indeed participate in making the knowledge base inconsistent. From the pseudo consequence operators Cn^* , De Bona and Hunter define consolidation functions based on pseudo-contraction (Hansson 1993a), which consists in consolidating the superset $Cn^*(K)$, instead of K , in order to minimise loss of information. Towards measuring inconsistencies, they have introduced some inconsistency measures (which, unlike culpability measures, evaluate the inconsistency degree of a knowledge base as a whole) based on their concept of “conflicts under the water”. Although similar in spirit, their strategy to identify hidden conflicts differs from ours. While the “under the water” formulæ responsible for causing inconsistencies depend on the adopted pseudo consequence operator, our concept of implicit contribution towards inconsistency is applied on the whole language, and are not confined to sub-formulæ of the original knowledge base. On the other hand, we focus on measuring the implicit culpability degree of the formulæ in a given knowledge base, as we are not interested in pseudo-contraction.

It would be worth to explore the precise connection between our concept of implicit contribution towards inconsistency and “conflicts under the water”. A potential research direction, towards consolidation, is to define tacit culpability measures in the supersets generated by the pseudo consequence operators proposed by De Bona and Hunter (2017) and other pseudo consequence operators in the literature (Santos et al. 2018). This includes investigating how tacit-dominance would behave in this setting. Secondly, our concept of implicit contribution towards inconsistency, and tacit culpability measure could also be used to define new pseudo-contraction/consolidation operators.

Acknowledgements

This research is supported by the German Research Association (DFG), project number 424710479.

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2): 510–530.
- Alchourrón, C. E.; and Makinson, D. 1985. On the logic of theory change: Safe contraction. *Stud Logica* 44(4): 405–422.
- Dalal, M. 1988. Investigations into a Theory of Knowledge Base Revision. In *Proceedings of the 7th National Conference on Artificial Intelligence*, 475–479. AAAI Press / The MIT Press.
- Daniel, L. 2009. *Paraconsistent Probabilistic Reasoning*. Ph.D. thesis, L'École Nationale Supérieure des Mines de Paris.
- De Bona, G.; and Hunter, A. 2017. Localising Iceberg Inconsistencies. *Artificial Intelligence* 246: 118–151.
- Dixon, S. E. 1994. *Belief revision: A computational approach*. Ph.D. thesis, University of Sydney.
- Falappa, M. A.; Fermé, E. L.; and Kern-Isberner, G. 2006. On the Logic of Theory Change: Relations Between Incision and Selection Functions. In *ECAI 2006, 17th European Conference on Artificial Intelligence*, volume 141 of *Frontiers in Artificial Intelligence and Applications*, 402–406. IOS Press.
- Fuhrmann, A. 1991. Theory contraction through base contraction. *J. Philos. Log.* 20(2): 175–203.
- Gärdenfors, P. 1988. *Knowledge in flux: Modeling the dynamics of epistemic states*. The MIT press.
- Grant, J. 1978. Classifications for Inconsistent Theories. *Notre Dame Journal of Formal Logic* 19(3): 435–444.
- Grant, J.; and Martinez, M. V., eds. 2018. *Measuring Inconsistency in Information*, volume 73 of *Studies in Logic*. College Publications.
- Hansson, S. O. 1991. Belief contraction without recovery. *Stud Logica* 50(2): 251–260.
- Hansson, S. O. 1993a. Changes of disjunctively closed bases. *J. Log. Lang. Inf.* 2(4): 255–284.
- Hansson, S. O. 1993b. Reversing the Levi identity. *J. Philos. Log.* 22(6): 637–669.
- Hansson, S. O. 1994. Kernel Contraction. *J. Symb. Log.* 59(3): 845–859.
- Hansson, S. O. 1997. Semi-Revision. *J. Appl. Non Class. Logics* 7(1-2): 151–175.
- Hansson, S. O. 1999. *A textbook of belief dynamics - theory change and database updating*, volume 11 of *Applied logic series*. Kluwer. ISBN 978-0-7923-5324-9.
- Hunter, A.; and Konieczny, S. 2005. *Approaches to Measuring Inconsistent Information*, 191–236. Berlin, Heidelberg: Springer Berlin Heidelberg.
- Hunter, A.; and Konieczny, S. 2006. Shapley inconsistency values. In *Proceedings of the 10th International Conference on Knowledge Representation (KR'06)*, 249–259. AAAI Press.
- Hunter, A.; and Konieczny, S. 2008. Measuring Inconsistency through Minimal Inconsistent Sets. In *Proceedings of the 11th International Conference on Knowledge Representation and Reasoning, KR 2008*, 358–366. AAAI Press.
- Nebel, B. 1990. *Reasoning and Revision in Hybrid Representation Systems*, volume 422 of *Lecture Notes in Computer Science*. Springer.
- Rott, H.; and Hansson, S. O. 2014. *Safe Contraction Revisited*, 35–70. Springer.
- Santos, Y. D.; Matos, V. B.; Ribeiro, M. M.; and Wassermann, R. 2018. Partial meet pseudo-contractions. *Int. J. Approx. Reason.* 103: 11–27.
- Thimm, M. 2019. Inconsistency Measurement. In *Proceedings of the 13th International Conference on Scalable Uncertainty Management (SUM'19)*, volume 11940 of *Lecture Notes in Artificial Intelligence*, 9–23. Springer International Publishing.
- Williams, M. 1994. Transmutations of Knowledge Systems. In *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning (KR'94)*, 619–629. Morgan Kaufmann.
- Williams, M. 1995. Iterated Theory Base Change: A Computational Model. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence, IJCAI 95*, 1541–1549. Morgan Kaufmann.