Kernel Contraction and the Order of Relevance

Jandson S. Ribeiro
University of Hagen, Germany
jandson.ribeiro@fernuni-hagen.de

Abstract

The postulate of *relevance* provides a suitable and general notion of minimal change for belief contraction. Relevance is tightly connected to smooth kernel contractions when an agent's epistemic state is represented as a logically closed set of formulae. This connection, however, breaks down when an agent's epistemic state is represented as a set of formulae not necessarily logically closed. We investigate the cause behind this schism, and we reconnect relevance with smooth kernel contractions by constraining the behaviour of their choice mechanisms and epistemic preference relations. Our first representation theorem connects smooth kernel contractions with a novel class of epistemic preference relations. For our second representation theorem, we introduce the principle of *symmetry of removal* that relates relevance to epistemic choices. For the last theorem, we devise a novel class of smooth kernel contractions, that satisfy relevance, which are based on epistemic preference relations that capture the principle of *symmetry of removal*.

1 Introduction

The field of *belief change* (Alchourrón, Gärdenfors, and Makinson, 1985; Hansson, 1999) studies how an agent should rationally maintain its body of beliefs as it evolves. The most interesting situations emerge when incoming information conflicts with the agent's current beliefs. In this case, for the sake of consistency, the agent must remove beliefs that conflict with the incoming information. This is known as *belief revision*, and a challenging task consists in deciding which information should be discarded. In essence, the agent should preserve most of its original beliefs, which is known as the principle of minimal change. Removal of information is by itself studied under the name of *belief contraction* and can be seen as a central problem in belief change, because it forms the cornerstone to properly define other kinds of changes. For instance, *belief revision* is in essence defined in terms of belief contraction: first remove any conflict with the incoming information by performing belief contraction, and thereafter accommodate the incoming information. When the underlying logic used to specify an agent's beliefs is closed under classical negation, this recipe for defining revision from contraction is formalised via the (external and internal) Levi Identity.

Belief contraction, and other forms of belief change, have been studied via two different, but equivalent, perspectives: (i) defining rationality postulates that conceptualise the principle of minimal change; and (ii) constructing classes of operations that explicitly capture such postulates. The first perspective provides a principled way for understanding what is a rational change, as well as identifying and forbidding counter-intuitive behaviours of belief change. The second perspective provides constructive apparatuses to understand how belief change behaves. The contrast between the two perspectives helps to understand how principles of minimal change (via rationality postulates) outline and constrain the epistemic choices an agent might make, and on which classes of epistemic preferences such choices can be based. Toward the first perspective, a major rationality postulate in belief contraction is *relevance* which intuitively states that a belief can only be removed for a good reason (see Section 2 for the formal definition). Hansson (1991) argues that *relevance* provides a general and suitable notion of minimal change, while Ribeiro et al. (2013) have shown that *relevance* is consistent with several non-classical logics.

As for the second perspective, the two dominating classes of belief contraction operations are: partial-meet functions (Alchourrón, Gärdenfors, and Makinson, 1985) which are based on a "what to keep" strategy; and (smooth) kernel contraction functions (Hansson, 1994) which are based on a "what to remove" strategy. These two classes of functions are dual, and are usually seen as two sides of the same coin (Falappa, Fermé, and Kern-Isberner, 2006). This duality permits to understand the rationality of relevance through two complementary constructive views. Indeed, in the most fundamental case, when an agent's epistemic state is represented as a logically closed set of formulae, called a theory, both *partial meet* and *smooth kernel contraction* satisfy relevance, and they collapse to the same class of functions (Hansson, 1999; Hansson and Wassermann, 2002).

Theories, however, are very restrictive, because they do not distinguish between implicit beliefs versus explicit beliefs. This distinction can be achieved by dropping the logical closure requirement, and simply representing an agent's epistemic state as a set of formulae, called a *belief base*. This alternative is more general and expressive than theories. However, for belief bases, the connection between relevance and *smooth kernel contraction* breaks down, though the connection between partial meet and relevance is not harmed. The reasons for this rupture are still unclear. Restoring the connection between relevance and smooth kernel contraction is extremely important to better support belief contraction for several reasons.

First, smooth kernel contraction and partial meet work as two complementary perspectives of belief contraction. At a first glance, it might not seem harmful to solely rely on partial meet as a canonical form for constructing belief contraction operations that satisfy relevance. It turns out that such a strategy severely limits the reach of belief contraction approaches in more expressive logics: in several non-classical logics, partial meet functions cannot even be defined (Ribeiro et al., 2013; Guerra and Wassermann, 2019), although belief contraction operations satisfying relevance do exist in such logics (Ribeiro, Nayak, and Wassermann, 2018). Kernel contraction operations could be studied in such logics as a natural alternative. Secondly, according to Hansson (2017) the way that partial meet functions conceptualise choices is counter-intuitive. Precisely, partial meet functions decide what to keep by forcing the agent to choose among the maximal conserving options, and intersecting the choices. This is counter-intuitive, because epistemically an agent's choice should not be based on maximising and maintaining the best of two or more scenarios. Therefore, partial-meet should not be taken as a canonical form to construct belief contraction operations. On the other hand, kernel contraction simply chooses beliefs to be removed, and therefore is not subject to this criticism. Moreover, kernel contraction appears as the dominating implementation strategy for handling inconsistencies in ontology repair (Horridge, 2011; Kalyanpur et al., 2007), whereas partial meet is almost unexplored (Cóbe and Wassermann, 2015).

Lastly, kernel contraction functions are closely related to other forms of belief dynamics, such as argumentation systems (Simari and Loui, 1992) and consistency maintenance approaches, such as culpability and inconsistency measures (Hunter and Konieczny, 2008, 2005).

Therefore, identifying the precise connection between kernel contractions and the postulate of relevance is of extreme importance to achieve a complete constructive overview of processes of rational belief contraction.

In this work, we restore the connection between smooth kernel contraction and relevance, by constraining the way that kernel contraction functions choose what to remove. We identify a principle relating relevance and epistemic choices, which we call *symmetry of removal*. We then examine how such a principle can be translated into epistemic preference relations. We start by defining a class of epistemic preference relations on the subsets of the agent's belief base, and we construct a novel class of kernel contraction functions based on such relations: the *spalled kernel contractions*. Thereafter, we devise a condition, *concordant-mirroring*, that embeds this principle in the epistemic preferences of an agent. This yields a new class of smooth kernel contraction functions which we call *mirrored kernel contraction*, and we show a representation theorem relating this novel class of kernel contraction functions and the relevance postulate.

Road map: In Section 2, we briefly review belief base contraction, its rationality postulates, including relevance, and smooth kernel contraction functions. In Section 3, we introduce the novel class of spalled kernel contraction functions. We show that spalled and smooth kernel contraction correspond to the same class of functions, with the difference that spalled kernel contraction explicitly realises the epistemic preference relations of an agent. In Section 4, we introduce the principle of *symmetry of removal* and the *concordant-mirroring* condition, and their respective representation theorems with relevance. In Section 5, we discuss some related works. Finally, in Section 6, we make some final considerations, and we discuss some future research lines worth to explore.

2 Belief Base Contraction

We assume that an agent's corpus of beliefs is represented as a belief base, which will be denoted by the letter $\mathcal{K}$. The term belief base has been used in the literature with two main purposes: (i) as a finite representation of an agent's beliefs (Nebel, 1990; Dixon, 1994; Dalal, 1988), and (ii) as a more general and expressive approach that distinguishes explicit from implicit beliefs (Fuhrmann, 1991; Hansson, 1999). We follow the latter approach, and therefore a belief base can be infinite. The power set of a set A is denoted by $\mathcal{P}(A)$. A belief base is specified in the language of an underlying logic. We treat a logic as a pair $\langle \mathcal{L}, Cn \rangle$, where $\mathcal{L}$ is a language, and $Cn : \mathcal{P}(\mathcal{L}) \rightarrow \mathcal{P}(\mathcal{L})$ is a logical consequence operator that indicates all the formulae that are entailed from a set of formulae in $\mathcal{L}$. We limit ourselves to logics whose consequence operator Cn satisfies:

***monotonicity**:* if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$;

***inclusion*:** $A \subseteq Cn(A)$;

***idempotency*:** $Cn(Cn(A)) = Cn(A)$;

***compactness**:* if $\varphi \in Cn(A)$ then there is some finite set $A' \subseteq A$ such that $\varphi \in Cn(A')$.

Consequence operators that satisfy the first three conditions above are called Tarskian. Some times we say that the logic itself is Tarskian. Throughout this work, unless otherwise stated, all the presented results regard logics whose consequence operators are both Tarskian and satisfy compactness. Let $\mathcal{K}$ be a belief base, a contraction function for $\mathcal{K}$ is a function $\dot{-} : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{L})$ that given an unwanted piece of information α outputs a subset of $\mathcal{K}$ which does not entail α. A contraction function is subject to the following basic rationality postulates (Hansson, 1991, 1994):

(success): if $\alpha \notin Cn(\emptyset)$ then $\alpha \notin Cn(\mathcal{K} \dot{-} \alpha)$;
(inclusion): $\mathcal{K} \dot{-} \alpha \subseteq \mathcal{K}$;
(vacuity): if $\alpha \notin Cn(\mathcal{K})$ then $\mathcal{K} \dot{-} \alpha = \mathcal{K}$;
(uniformity): if for all $\mathcal{K}' \subseteq \mathcal{K}$ it holds that $\alpha \in Cn(\mathcal{K}')$ iff $\beta \in Cn(\mathcal{K}')$, then $\mathcal{K} \dot{-} \alpha = \mathcal{K} \dot{-} \beta$;
(core-retainment): if $\beta \in \mathcal{K} \setminus (\mathcal{K} \dot{-} \alpha)$ then there exists a $\mathcal{K}' \subseteq K$ s.t $\alpha \notin Cn(\mathcal{K}')$ but $\alpha \in Cn(\mathcal{K}' \cup \{\beta\})$;
(**relative closure**): if $\beta \in \mathcal{K} \cap Cn(\mathcal{K} \dot{-} \alpha)$ then $\beta \in \mathcal{K} \dot{-} \alpha$;

(relevance): if $\beta \in \mathcal{K} \setminus (\mathcal{K} \dot{-} \alpha)$ then there is some $\mathcal{K}'$ such that $\mathcal{K} \dot{-} \alpha \subseteq \mathcal{K}' \subseteq \mathcal{K}$, $\alpha \notin Cn(\mathcal{K}')$ but $\alpha \in Cn(\mathcal{K}' \cup \{\beta\})$

For a discussion on the rationale of this postulates, see (Hansson, 1999). It worth highlighting that relevance implies core-retainment. A contraction function that satisfies the six first rationality postulates above will be dubbed a *rational contraction function*. Additionally if a rational contraction function satisfies relevance, then we will say that it is *fully rational*. One of the most influential classes of rational contraction functions is the class of (smooth) kernel contraction functions (see Definition 5), which are defined on kernels and incision functions:

Definition 1. *An α-kernel of a belief base $\mathcal{K}$ is a set X such that (1) $X \subseteq \mathcal{K}$; (2) $\alpha \in Cn(X)$; and (3) if $X' \subset X$ then $\alpha \notin Cn(X')$.*

An α-kernel of a belief base $\mathcal{K}$ is a minimal subset of $\mathcal{K}$ that does entail α. Due to compactness, every kernel is finite. The set of all α-kernels of a belief base $\mathcal{K}$ is denoted by $\mathcal{K} \perp\!\!\!\perp \alpha$. Formulae that do not appear in any α-kernel are called α-free. As α-free formulae have no connection with the formula α to be contracted, they should be kept intact, while only not α-free formulae should be picked for removal. This choice of removal is realised by an incision function:

Definition 2. *Let $\mathcal{C}(\mathcal{K}) = \{\mathcal{K} \perp\!\!\!\perp \alpha \mid \alpha \in \mathcal{L}\}$ be the set of all kernel sets of $\mathcal{K}$. A standard incision function on a belief base $\mathcal{K}$ is a function $\sigma : \mathcal{C}(\mathcal{K}) \to \mathcal{P}(\mathcal{L})$ such that*

(1) $\sigma(\mathcal{K} \perp\!\!\!\perp \alpha) \subseteq \bigcup \mathcal{K} \perp\!\!\!\perp \alpha$;

(2) if $X \in \mathcal{K} \perp\!\!\!\perp \alpha$ and $X \neq \emptyset$, then $X \cap \sigma(\mathcal{K} \perp\!\!\!\perp \alpha) \neq \emptyset$.

Intuitively, in order to contract a formula α, an agent chooses at least one formula from each α-kernel, and only formulae from such kernels. An incision function works as an extra-logical device that realises an agent's epistemic preferences, and it chooses the least preferable formulae in each α-kernel to be removed. Incision functions, as per Definition 2, do not clearly reference such preferences, and the connection between epistemic preferences and choices is obscure. In Section 3, we devise a class of incision functions that are explicitly built on the agent's epistemic preferences. As such preferences depend on the formula α to be contracted, it will be convenient to explicitly reference the formula α to be contracted, rather than its α-kernels. Towards this end, to facilitate presentation, we slightly reformulate the signature of an incision function to receive as parameter a single formula instead of all α-kernels:

Definition 3. *A formula-based incision function, centred on a belief base $\mathcal{K}$, is a mapping $\sigma : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ such that*

(1) $\sigma(\alpha) \subseteq \bigcup \mathcal{K} \perp\!\!\!\perp \alpha$;

(2) if $X \in \mathcal{K} \perp\!\!\!\perp \alpha$ and $X \neq \emptyset$, then $X \cap \sigma(\alpha) \neq \emptyset$;

(3) if $\mathcal{K} \perp\!\!\!\perp \alpha = \mathcal{K} \perp\!\!\!\perp \beta$ then $\sigma(\alpha) = \sigma(\beta)$.

In Definition 3, conditions (1) and (2) correspond respectively to the translation of conditions (1) and (2) of Definition 2. It is important to stress that standard incision functions are defined on kernel-sets with the only purpose of trivially capturing the behaviour imposed by the uniformity postulate: if two formulae, say α and β, are entailed exactly by the same subsets of $\mathcal{K}$ (we say α and β are $\mathcal{K}$-uniform), then α and β must present the same contraction result. This is equivalent to say that α and β present the same set of kernels. In this way, by working on kernel sets, the standard incision functions trivially captures *uniformity* behaviour. Condition (3) guarantees that this will also be the case for formula-based incision functions.

Proposition 4. *Formula-based incision functions and standard incision functions are interchangeable, that is,*

1. *for all standard incision function σ there is a formula-based incision function σ' such that $\sigma(\mathcal{K} \perp\!\!\!\perp \alpha) = \sigma'(\alpha)$, for all α;*
2. *for all formula-based incision function σ' there is a standard incision function σ such that $\sigma(\mathcal{K} \perp\!\!\!\perp \alpha) = \sigma'(\alpha)$, for all α;*

As shown in Proposition 4, standard and formula-based incision functions produce the same class of choices, and therefore we can use them interchangeably. From now on, we will employ solely the formula-based incision functions, and we will refer to them simply as incision functions.

A contraction operation can be constructed by removing the formulae picked by an incision function. Contraction functions that follow this recipe are called kernel contraction functions:

Definition 5. *(Hansson, 1994) Given a belief base $\mathcal{K}$ and an incision function σ for $\mathcal{K}$, the kernel contraction function $\dot{-}_\sigma$ is defined as: $\mathcal{K} \dot{-}_\sigma \varphi = \mathcal{K} \setminus \sigma(\varphi)$.*

Kernel contractions functions, however, are characterised only by the first five postulates:

Theorem 6. *(Hansson and Wassermann, 2002) A contraction function satisfies* success*,* inclusion*,* vacuity*,* uniformity*, and* core-retainment *iff it is a kernel contraction.*

To capture *relative-closure*, Hansson (1994) has proposed the *smoothness* property to be put upon incision functions:

smoothness if $\mathcal{K}' \subseteq \mathcal{K}$, and $\varphi \in Cn(\mathcal{K}')$ and $\varphi \in \sigma(\alpha)$ then $\mathcal{K}' \cap \sigma(\alpha) \neq \emptyset$.

Incision functions satisfying *smoothness* are called *smooth incision functions*, and kernel contraction functions built upon smooth incision functions are called *smooth kernel contraction functions*. Smoothness states that if a formula φ is picked for removal, then each subset of $\mathcal{K}$ that entails φ must also have some formula picked for removal. In propositional logics, smoothness captures relative-closure (Hansson, 1994). We extend this result for more expressive logics:

Theorem 7. *If Cn is Tarskian and satisfies compactness, then a contraction function satisfies* success*,* inclusion*,* vacuity*,* uniformity*,* core-retainment*, and* relative-closure *iff it is a smooth kernel contraction function.*

Though *smoothness* captures relative-closure, smoothness alone is not capable of connecting with *relevance* (Hansson, 1994). To fill this gap, at Section 4, we will first unveil in the next section the precise connection between smooth incision functions and epistemic preference relations.

3 Relational Incision Functions: Smoothness

Some beliefs might be deemed more reliable than others, and when performing a contraction, an agent should remove only the least reliable beliefs. Towards this end, an incision function realises such epistemic choices which are founded on the agent's epistemic preferences. However, incision functions, as presented in Section 2 do not make reference to the agent's epistemic preferences. Moreover, *smoothness* simply constrains what is rational to be chosen, but does not give any hint of how such choices relates to epistemic preferences. In this section, we devise a class of incision functions whose choices are explicitly founded on the epistemic preference relations of the agent, called *effacings*. We show that *effacings* and smooth incision functions coincide. This will serve as a basis to construct, in Section 4, smooth kernel contraction functions satisfying relevance.

An agent's epistemic preference relation could be specified via a single binary relation between its beliefs. Although this has been the dominating form of representing epistemic preferences on belief bases, such as *safe contraction* (Alchourrón and Makinson, 1985), and *ensconcements* (Williams, 1994), this strategy is very limiting because of the two main impossibility results for belief bases:

1. There are fully rational contractions that cannot be based on a single epistemic preference relation (Hansson, 1999);
2. There are natural epistemic preferences that can be specified between sets of beliefs, but are impossible to be translated in terms of formulae in the belief base (Hansson, 1999).

These impossibility results severely limit our alternatives for specifying epistemic preferences relations. In order to work around these issues, we will consider preference relations between the subsets of belief bases, and that an agent might present several epistemic preference relations. We move to give a brief overview of how our strategy will work, and thereafter we proceed to the formal definitions. Intuitively, when an agent wishes to relinquish a belief α from its current belief base $\mathcal{K}$, the agent shall consult an epistemic preference relation, say $\leqslant_\alpha \subseteq \mathcal{P}(\mathcal{K}) \times \mathcal{P}(\mathcal{K})$, that indicates the degree of reliability between its "blocks" of knowledge, or in other words, which clusters of information are deemed more reliable (or trustworthy) than others. A pair $A \leqslant_\alpha B$ means that A is at least as reliable as B, in the sense that the beliefs within A are jointly equally or more reliable than the joint beliefs within B. This means that the more reliable a set is, the more protected all the beliefs that it carries are. Towards this end, if a set is not among the least reliable ones, all the beliefs it carries are protected against contraction. We call such sets *resistant sets*.

Therefore, in order to contract α, the agent shall retain the beliefs within such resistant sets, whereas only the beliefs that do not fall within the resistant sets are *susceptible* for contraction. An incision function based on such epistemic preference relations picks the *susceptible* beliefs that appear in each α-kernel.

We move now to formalise this recipe. We assume that for each formula $\alpha \in \mathcal{L}$, an agent presents an epistemic preference relation $\leqslant_\alpha$ on the subsets of its current belief base $\mathcal{K}$. Such an assignment is given by a function τ, called a *spalling*. Clearly, not every binary relation can be regarded as suitable for specifying epistemic preferences, let alone yield rational contraction functions. Each of the relations assigned by τ shall then satisfy some minimal requirements,

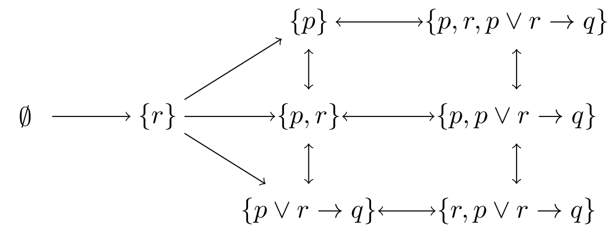

Figure 1: A q-shard on a belief base $\mathcal{K} = \{p, r, p \vee r \to q\}$. The relation is transitive, but to avoid visual pollution we omit edges obtained by transitivity.

which we present in Definition 8.

Definition 8. *Given a belief base $\mathcal{K}$, a spalling is a function $\tau : \mathcal{L} \to \mathcal{P}(\mathcal{P}(\mathcal{K}) \times \mathcal{P}(\mathcal{K}))$ that maps each formula $\alpha \in \mathcal{L}$ to a relation $\leqslant^\tau_\alpha$ on the subsets of $\mathcal{K}$ such that each $\leqslant^\tau_\alpha$ satisfies all the following conditions:*

transitivity: *if $A \leqslant^\tau_\alpha B$ and $B \leqslant^\tau_\alpha C$ then $A \leqslant^\tau_\alpha C$.*

isotonicity: *if $A \subseteq Cn(B)$ then $A \leqslant^\tau_\alpha B$.*

α-maximality: *if $\alpha \in Cn(A)$ then $B \leqslant^\tau_\alpha A$.*

α-discernment: *if $\{\psi\} \leqslant^\tau_\alpha \{\varphi\}$ and φ is α-free then ψ is α-free.*

relational uniformity: *if for all $\mathcal{K}' \subseteq \mathcal{K}$ it holds $\alpha \in Cn(\mathcal{K}')$ iff $\beta \in Cn(\mathcal{K}')$, then $\leqslant^\tau_\alpha = \leqslant^\tau_\beta$, where $\leqslant^\tau_\beta = \tau(\beta)$.*

conjunctiveness: *$A \cup B \leqslant^\tau_\alpha A$ or $A \cup B \leqslant^\tau_\alpha B$.*

Giving a spalling τ, we will write $\leqslant^\tau_\alpha$ as a shorthand for $\tau(\alpha)$. We call each $\leqslant^\tau_\alpha$, an α-shard. $A \leqslant^\tau_\alpha B$ means that A is at least as preferable as B. When it is clear from context, we will omit the superscript $^\tau$, and simply write $\leqslant_\alpha$. A spalling is similar in spirit to the concept of *faithful assignments* of Katsuno and Mendelzon (2003) used to specify preference relations for *belief update*. There are, however, many differences between our approach to the faithful assignments of Katsuno and Mendelzon. First, a spalling considers relations on the subsets of a belief base, while assignments are designed for relations between interpretations. Second, the properties put upon assignments and spallings are utterly different. Last, assignments are designed for belief update on theories finitely represented by a single formula, while spallings realise epistemic preferences for (possibly infinite) belief bases.

We proceed to explain each of the properties put upon the relations $\leqslant^\tau_\alpha$ assigned by spallings. To support the explanation of such properties, we will consider as an example the belief base $\mathcal{K} = \{p, r, p \vee r \to q\}$ and the q-shard depicted in Fig. 1. *Transitivity* is a widely explored property used to specify epistemic preferences: if one prefers A rather B and B rather C than one should also prefer A rather C. *Isotonicity* states that adding information does not make a set more reliable. For example, in Fig. 1, both $\{p\}$ and $\{r\}$ are at least as reliable as $\{p, r\}$: $\{p\}$ is as reliable as $\{p, r\}$, while $\{r\}$ is strictly more reliable than $\{p, r\}$. *Isotonicity* is the reverse of the *dominance* property used in both *Epistemic Entrenchment* and *Safe-contraction* (Hansson, 1999; Gärdenfors, 1988).

An agent should remove only the least reliable pieces of information. The property *α-maximality* imposes that a set A which entails the formula α to be contracted cannot be trusted, and therefore A must be among the least reliable ones. In our example, there are only three sets that entail the formula q, which are depicted at the rightmost column in Fig. 1: the belief base $\mathcal{K}$ itself and the two q-kernels. All of them are put among the least reliable sets. On the other hand, *α-discernment* states that α-free formulae are strictly more reliable than not α-free formulae. The property *α-discernment* jointly with *transitivity, isotonic,* and *conjunctiveness* states that finite α-free sets are strictly more reliable than not α-free sets. *Relational uniformity*, as the name suggests, is related to the *uniformity* postulate. It states that the epistemic preferences of an agent should not be syntax sensitive: if two formulae, even though not necessarily logically equivalent, are entailed exactly by the same subsets of the belief base $\mathcal{K}$, then they should present exactly the same epistemic preference relation.

For the rationale of *conjunctiveness*, consider a set $X = A \cup B$, and that a belief $\beta \in X$ is chosen to be removed during the contraction of a formula α. Intuitively, only information among the least reliable sets should be picked for removal, which means that X is among the least reliable sets. However, as β is in X, β appears either in A or in B. Therefore, it is plausible to assume that the subset in which β appears (either A or B), must also be among the least reliable sets, that is, either A or B must be as (un-)reliable as X (*conjunctiveness*). Consider, for instance, the q-shard in Fig. 1 and the set $\{p, r\}$ among the least preferable sets. According to *conjunctiveness*, at least one between $\{p\}$ or $\{r\}$ must be as (un-)reliable as $\{p, r\}$, which is the case of $\{p\}$ while $\{r\}$ is strictly more preferable than them.

We extend the notion of α-free from formulae to sets. We say that a set A is α-free, if and only if every formula in A is α-free. The maximal elements of a set A w.r.t a transitive relation $\leqslant$ is $\max_{\leqslant}(A) = \{a \in A \mid \text{for all } b \in A, \text{ if } a \leqslant b \text{ then } b \leqslant a\}$. An α-shard $\leqslant_\alpha \subseteq \mathcal{P}(\mathcal{K}) \times \mathcal{P}(\mathcal{K})$ realises the epistemic preference relation of an agent during the contraction of a formula α. Intuitively, $A \leqslant_\alpha B$ means that whole information in A is at least as reliable than the whole information in B. Additionally, if $A \leqslant_\alpha B$ and $B \nleqslant_\alpha A$, then information within A as a whole is strictly more reliable than the whole information in B. This means that B is strictly more vulnerable to contraction than A. For instance, in the relation $\leqslant_q$ depicted in Fig. 1, the set $\{p \vee r \to q\}$ is strictly more vulnerable to contraction than $\{r\}$. Ideally, an agent should remove information only from the most vulnerable sets (the sets at the two rightmost columns at Fig. 1) while keeping all other sets, like $\{r\}$, intact. Sets whose information are not subject to removal are called resistant:

Definition 9. *A set $A \subseteq \mathcal{K}$ is a resistant set w.r.t an α-shard $\leqslant_\alpha$ iff A is α-free or there is some $B \subseteq \mathcal{K}$ such that $A \leqslant_\alpha B$ and $B \nleqslant_\alpha A$. The resistant sets from $\mathcal{K}$ w.r.t $\leqslant_\alpha$ is given by*

$$\mathsf{resist}_{\leqslant_\alpha}(\mathcal{K}) = \mathcal{P}(\mathcal{K}) \setminus \mathsf{max}_{\leqslant_\alpha}(\{A \subseteq \mathcal{K} \mid A \textit{ is not } \alpha\textit{-free}\})$$

Intuitively, a resistant set carries information that shall be protected during contraction. As only not α-free formulae should be removed, every α-free set is by default resistant. *Resistant sets* are similar in spirit to the notion of safe-formulae used in safe-contraction. The main difference is that to be safe, a formula φ must be covered by some other formula in every α-kernel in which φ appears, whereas resistant sets only need to be covered by some other set. For instance, in Fig. 1, $\{r\}$ is resistant because it is strictly more preferable than $\{p\}$. As all the formulae within a resistant set shall be protected, the only information that can be removed are those that do not appear in any of the resistant sets modulo $\leqslant_\alpha$. We call such formulae α-susceptible.

Definition 10. *Within a belief base $\mathcal{K}$, a formula φ is α-susceptible w.r.t an α-shard $\leqslant_\alpha$ iff $\varphi \notin A$, for all $A \in \mathsf{resist}_{\leqslant_\alpha}(\mathcal{K})$.*

For conciseness, when the α-shard is clear from context, we will simply say that a formula is α-susceptible. For instance, in the q-shard illustrated at Fig. 1, we have only two resistant sets: $\emptyset$ and $\{r\}$. Therefore, the formulae p and $p \vee r \to q$ are q-susceptible, because they do not appear in any of the resistant sets. We highlight the following interesting properties regarding α-susceptible formulae, which will be useful for defining a new class of incision functions:

Proposition 11. *If $\leqslant_\alpha$ is an α-shard on a belief base $\mathcal{K}$,*

1. *every α-susceptible formula w.r.t $\leqslant_\alpha$ is not α-free;*
2. *α is not tautological and $\alpha \in Cn(\mathcal{K})$ iff there is an α-susceptible formula in $\mathcal{K}$.*

Proof sketch. For item 1, let φ be an α-susceptible formula w.r.t an α-shard $\leqslant_\alpha$. Thus, φ does not appear in any of the resistant sets. By definition, the set of all α-free formulae is resistant. Thus, φ does not appear in such a set, which means φ is not α-free. For item 2, the direction "$\Leftarrow$" follows from item 1, because an α-susceptible formula necessarily is not α-free which implies that $\mathcal{K}$ must entail α. For the direction "$\Rightarrow$", from $\alpha \in Cn(\mathcal{K})$ we get there is at least one α-kernel $X \in \mathcal{K} \perp\!\!\!\perp \alpha$, and from compactness we know that all of them are finite. Let us fix an α-kernel $X \in \mathcal{K} \perp\!\!\!\perp \alpha$. As α is not tautological, $X \neq \emptyset$. From α-maximality, we get that X is maximal, and from conjunctiviness, we can prove by induction on size of X that there is a $\varphi \in X$ such that $X \leqslant_\alpha \{\varphi\}$. Let us fix such a φ. Therefore, as X is maximal, we get that $\{\varphi\}$ is also maximal. This jointly with isotonicity, implies that every set in which φ appears is also maximal. Therefore, every set that φ appears is not resistant. This means that φ is α-susceptible. □

As every α-susceptible formula is not α-free, and their existence is guaranteed when the underlying belief base entails α, we can define an incision function that selects all α-susceptible formulae, called an *effacing*:

Definition 12. *Given a belief base $\mathcal{K}$, and a spalling τ on $\mathcal{K}$. An effacing is a function $\delta_\tau : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ such that*

$$\delta_\tau(\alpha) = \{\varphi \in \mathcal{K} \mid \varphi \textit{ is } \alpha\textit{-susceptible w.r.t. } \leqslant^\tau_\alpha\}.$$

Although an effacing does not look inside each α-kernel to pick a formula for removal, an effacing does hit every α-kernel and it is therefore indeed an incision function:

Proposition 13. *Every effacing is an incision function.*

Proof sketch. We have to show that every effacing δ_τ satisfies conditions (1), (2) and (3) from Definition 3. Condition (3) follows from relational uniformity. For condition (1), note that from Proposition 11, we have that every α-susceptible formula is not α-free. Thus, $\delta_\tau(\alpha) \subseteq \bigcup \mathcal{K} \perp\!\!\!\perp \alpha$. For condition (2), let $X \in \mathcal{K} \perp\!\!\!\perp \alpha$ such that $X \neq \emptyset$, we will show that there is some $\varphi \in X$ such that $\varphi \in \delta_\tau(\alpha)$. As X is an α-kernel, we get from compactness that X is finite. From α-maximality we get $X \in \max_{\leqslant_\alpha}(\mathcal{P}(\mathcal{K}))$, which means that X is not resistant. Thus, as $\leqslant_\alpha$ satisfies conjunctiveness, we can show by induction on the size of X that there is some $\varphi \in X$ such that $\{\varphi\}$ is also maximal. This jointly with isotonicity, implies that any set that contains φ will also be maximal, and therefore it is not resistant. Thus, φ is α-susceptible, which means that $\varphi \in \delta_\tau(\alpha)$. □

As effacings are incision functions, we can use them to construct a new class of kernel contraction functions:

Definition 14. *Let τ be a spalling on a belief base $\mathcal{K}$. A kernel contraction founded on τ is defined as $\mathcal{K} \dot{-}_\tau \alpha = \mathcal{K} \setminus \delta_\tau(\alpha)$. We say that $\dot{-}_\tau$ is a spalled kernel contraction.*

Example I illustrates a spalled kernel contraction function, based on the α-shard depicted in Fig. 1.

Example I. *Let $\mathcal{K} = \{p, r, p \vee r \to q\}$, and τ be a spalling such that $\leqslant^\tau_\alpha q$ corresponds to the relation depicted at Fig. 1. The only q-susceptible formulae are p and $p \vee r \to q$. Thus, $\delta_\tau(q) = \{p, p \vee r \to q\}$, and $\mathcal{K} \dot{-}_\tau q = \mathcal{K} \setminus \delta_\tau(q) = \{r\}$. This contraction satisfies relative-closure, because $\{r\}$ does not imply either of the removed formulae.*

The class of kernel contractions based on effacings matches exactly the class of smooth kernel contractions. We present the first direction of this representation theorem:

Theorem 15. *Every spalled kernel contraction is smooth.*

Proof sketch. Let $X \subseteq \mathcal{K}$ and $\varphi \in \delta_\tau(\alpha)$ such that $\varphi \in Cn(X)$. We will show that there is some $\psi \in X$ such that $\psi \in \delta_\tau(\alpha)$. From $\varphi \in Cn(X)$, we get that there is a $X' \in X \perp\!\!\!\perp \varphi$. Let us fix such a X'. Due to conjunctiveness, isotonicity and transitivity, we get that there is some $\psi \in X'$ such that $X' \leqslant_\alpha \{\psi\}$. From isotonicity, transitivity and α-discernment, we have that $\{\varphi\} \leqslant_\alpha \{\psi\}$, and ψ is not α-free. Moreover, as $\varphi \in \delta_\tau(\alpha)$, it follows that $\{\varphi\}$ is maximal among all not α-free sets. This implies from $\{\varphi\} \leqslant_\alpha \{\psi\}$ that $\{\psi\}$ is also maximal among all not α-free sets. This jointly with isotonicity implies that every set that has ψ is not resistant, and therefore ψ is α-susceptible. □

To conclude the representation theorem between effacings and smooth kernel contractions, we will need to show that each smooth incision σ function is indeed an effacing, that is, that there is a spalling τ such that $\sigma(\alpha) = \delta_\tau(\alpha)$. The main strategy is to construct, for each formulae α, a relation $\leqslant^\sigma_\alpha$ such that the formulae in $\sigma(\alpha)$ corresponds exactly to the α-susceptible formulae modulo $\leqslant^\sigma_\alpha$. We call $\leqslant^\sigma_\alpha$ the α-projection of σ. The spalling can then be achieved by mapping each formula α to its α-projection. We call this mapping between formulae and projections the shadowing of σ. There are two essential conditions we need to capture in constructing such shadowing τ: (C1) $\delta_\tau(\alpha)$ matches $\sigma(\alpha)$, and (C2) that τ is indeed a spalling (all conditions at Definition 8 are satisfied). We start by presenting how α-projections are constructed:

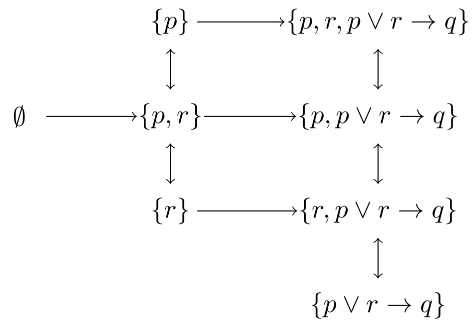

Figure 2: An α-projection of the incision function at Example II.

Definition 16. *Given an incision function σ for a belief base $\mathcal{K}$, and a formula α. The α-projection of σ is the smallest relation $\leqslant^\sigma_\alpha \subseteq \mathcal{P}(\mathcal{K}) \times \mathcal{P}(\mathcal{K})$ such that*

1. *if $A \cap \sigma(\alpha) \neq \emptyset$ then $B \leqslant^\sigma_\alpha A$.*
2. *if $A \subseteq Cn(B)$ then $A \leqslant^\sigma_\alpha B$*
3. *if A is α-free then $A \leqslant^\sigma_\alpha B$.*
4. *if both A and B are not α-free, and $(A \cup B) \cap \sigma(\alpha) = \emptyset$ then $A \leqslant^\sigma_\alpha B$*

If a formula φ is picked by $\sigma(\alpha)$, then every set A that has φ cannot be a resistant set. This can be done by putting such sets as the least preferable ones (Item 1). This not only captures α-maximality, but also guarantees that the α-susceptible formulae modulo $\leqslant^\sigma_\alpha$ match the formulae in $\sigma(\alpha)$, see Proposition 22 below. Item 2 is the *isotonicity* condition, while Item 3 captures α-discernment by putting each α-free set among the most preferable ones. Item 4 makes equally preferable any two sets that, though not α-free, have no formulae picked by $\sigma(\alpha)$. Item 4 jointly with Item 1 captures conjunctiveness. An α-projection also satisfies transitivity. Example II illustrates how an α-projection is constructed.

Example II. *Let $\mathcal{K} = \{p, r, p \vee r \to q\}$, and consider that we want to contract the formula q using the smooth incision function σ such that $\sigma(q) = \{p \vee r \to q\}$. The q-projection $\leqslant^\sigma_q$ of σ is depicted in Fig. 2. The relation $\leqslant^\sigma_q$, as we explain below, is transitive, and to avoid visual pollution, we have not drawn the transitive edges of $\leqslant^\sigma_q$ in Fig. 2. Let us explain how $\leqslant^\sigma_q$ is constructed according to Definition 16. First, every set that contains the formula $p \vee r \to q$ (sets at the rightmost side) is put as the most vulnerable ones, as imposed by condition 1. Condition 2 imposes the relation to satisfy isotonicity. For instance, $\{p\} \leqslant^\sigma_q \{p, r\}$. The empty set is the only α-free set, and therefore it is put as the most reliable one (condition 3). According to condition 4, each pair of non q-free sets that do not contain the formula $p \vee r \to q$ is put as equally preferable (sets in the middle column). To facilitate visualisation, we have split the relation into three columns. Sets in the same column are equally preferable (resp. equally vulnerable), while the sets towards* *the right are more vulnerable than all the sets towards the left end. From all this, we see that $\leqslant_q^\sigma$ is transitive.*

Lemma 17 below will be helpful in proving the properties of an α-projection. It states that any set that has no formulae in common with $\sigma(\alpha)$ is a resistant set.

Lemma 17. *For every smooth incision function and α-projection $\leqslant_\alpha^\sigma$, if $B \cap \sigma(\alpha) = \emptyset$ and $A \leqslant_\alpha^\sigma B$ then $A \cap \sigma(\alpha) = \emptyset$.*

Proposition 18. *If an incision function σ is smooth, then every α-projection of σ satisfies: isotonicity, α-maximality, α-discernment, conjunctiveness, and transitivity.*

According to Proposition 18 above, every α-projection from a smooth incision function satisfies almost all conditions put upon spallings. Another interesting property is related to *relational uniformity*: if two formulae are entailed exactly by the same subsets of a belief base $\mathcal{K}$, then their projections coincide.

Observation 19. *Let σ be an incision function on a belief base $\mathcal{K}$, and let α and β be two formulae. If for all $\mathcal{K}' \subseteq \mathcal{K}$, it holds that $\alpha \in Cn(\mathcal{K}')$ iff $\beta \in Cn(\mathcal{K}')$, then $\leqslant_\alpha^\sigma = \leqslant_\beta^\sigma$.*

We are ready to define the shadowing of an incision function σ, which basically maps each formula α to the respective α-projection of σ:

Definition 20. *Given an incision function σ, defined on a belief base $\mathcal{K}$. The shadowing of σ is the function $\mathcal{T}_\sigma : \mathcal{L} \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{K}) \times \mathcal{P}(\mathcal{K}))$ such that $\mathcal{T}_\sigma(\alpha) = \leqslant_\alpha^\sigma$, where $\leqslant_\alpha^\sigma$ is the α-projection of σ.*

We can see, from Proposition 18 and Observation 19, that shadowings are spallings:

Corollary 21. *If an incision function is smooth, then its shadowing is a spalling.*

So far, we have shown how to construct a spalling from a smooth incision function σ. The only piece missing to complete the representation theorem is to show that the effacing based on the spalling $\mathcal{T}_\sigma$ actually matches σ. This follows directly from the construction of the α-projection:

Proposition 22. *If σ is a smooth incision function on a belief base $\mathcal{K}$, then for all formula α:*

$$\sigma(\alpha) = \{\varphi \in \mathcal{K} \mid \varphi \text{ is } \alpha\text{-susceptible modulo } \leqslant_\alpha^\sigma\}.$$

Proof sketch. We need to show that a formula $\varphi \in \mathcal{K}$ is α-susceptible modulo $\leqslant_\alpha^\sigma$ iff $\varphi \in \sigma(\alpha)$. The direction "$\Leftarrow$" follows directly from condition (1) at Definition 16. For direction "$\Rightarrow$", the proof is by its contrapositive. Let us suppose that $\varphi \notin \sigma(\alpha)$, and we will show that φ is not α-susceptible. The case that φ is α-free is trivial. Let suppose then that φ is not α-free. Thus, there is some $A \in \mathcal{K} \perp\!\!\!\perp \alpha$ such that $\varphi \in A$. Let $A' = A \setminus \sigma(\alpha)$. Note that $A' \cap \sigma(\alpha) \neq \emptyset$, as $\varphi \in A$ and $\varphi \notin \sigma(\alpha)$. As σ is an incision function, we have that $A \cap \sigma(\alpha) \neq \emptyset$. Thus, from condition (1) at Definition 16, we get that (i) $A' \leqslant_\alpha^\sigma A$; and from the contrapositive of Lemma 17 (on that lemma let B stand for A'), we get (ii) $A' \not\leqslant_\alpha^\sigma A$. Therefore, A' is a resistant set, which implies that φ is not α susceptible. □

The representation theorem between smooth kernel contraction and spalled kernel contraction easily follows from Theorem 15 and Proposition 22:

Theorem 23. *A kernel contraction is smooth iff its a spalled kernel contraction.*

4 Concordance and Relevance

In this section, we present two representation theorems between kernel contraction functions and the postulate of relevance. In our first representation theorem, we identify a principle that relates incision functions and the postulate of relevance, which we call the *symmetry of removal*. This principle allows us to frame exactly the class of all incision functions that yield fully rational kernel contraction functions. The principle of *symmetry of removal*, similar to smoothness, states which properties an incision function must obey in order to capture relevance, but it does not tell us how to construct such incision functions nor how the epistemic preference relations of an agent should look. In our second representation theorem, we fill this gap by putting an extra constraint, called *concordant-mirroring*, upon the spalling of an effacing. The *concordant-mirroring* condition is inspired by the principle of *symmetry of removal*. Before we present this principle and *concordant-mirroring* we shall first introduce the notion of *completion* and *concordant sets*.

Definition 24. *Let $\mathcal{K}$ be a belief base, and α be a formula. An α-completion of a set $A \subseteq \mathcal{K}$ is a set X such that (a) $X \subseteq \mathcal{K}$, and (b) $\alpha \in Cn(A \cup X)$. The set of all α-completions of A is given by*

$$com_{\mathcal{K}}(A, \alpha) = \{X \subseteq \mathcal{K} \mid \alpha \in Cn(A \cup X)\}.$$

Intuitively, an α-completion of a set $A \subseteq \mathcal{K}$ is a set X that collaborates with A to entail α. Unlike kernels, which are required to entail α, an α-completion only carries information that "completes" the information of A towards entailing α. The following example illustrates the notion of α-completion.

Example III. *Let $\mathcal{K} = \{p, r, p \vee r \rightarrow q\}$. The set $X = \{p \vee r \rightarrow q\}$ is a q-completion of both $\{r\}$ and $\{p\}$, because $X \subseteq \mathcal{K}$, $q \in Cn(X \cup \{r\})$ and $q \in Cn(X \cup \{p\})$. The belief base $\mathcal{K}$ itself and the q-kernels $\{p, p \vee r \rightarrow q\}$ and $\{r, p \vee r \rightarrow q\}$ are trivially q-completions of both $\{p\}$ and $\{r\}$. There are no other q-completions for either $\{p\}$ or $\{r\}$. Thus, $com_{\mathcal{K}}(\{p\}, q) = com_{\mathcal{K}}(\{r\}, q)$.*

In Example III above, the sets $\{p\}$ and $\{r\}$, though not logically equivalent, present the same q-completions. Sets that agree upon the same collection of α-completions will be called α-concordant sets:

Definition 25. *Let α be a formula, and $\mathcal{K}$ a belief base $\mathcal{K}$. Two sets $A, B \subseteq \mathcal{K}$ are α-concordant within $\mathcal{K}$ iff $com_{\mathcal{K}}(A, \alpha) = com_{\mathcal{K}}(B, \alpha)$.*

Intuitively, α-concordant sets are indistinguishable w.r.t to their α-completions. When $\mathcal{K}$ is clear from context, we simply say that two sets are α-concordant. As we show in Proposition 26 below, α-concordant sets are tightly connected to the relevance postulate via the following principle:

***symmetry of removal*:** if A and B are α-concordant within $\mathcal{K}$ then $A \cap \sigma(\alpha) \neq \emptyset$ iff $B \cap \sigma(\alpha) \neq \emptyset$.

The principle of *symmetry of removal* states that formulae are removed uniformly from α-concordant sets, that is, if two sets are α-concordant, then an incision function will either remove formulae from both of them or it will keep both of them intact. Back to Example III, the sets $\{r\}$ and $\{p\}$ are α-concordant. Thus, according to *symmetry of removal*, an incision function must remove both r and p, or remove neither. The contraction in Example I violates *relevance*, because the incision function removes p but does not remove r. On the other hand, the contraction illustrated in Example II does satisfy *relevance*, because the incision function removes neither p nor r.

Indeed, symmetry of removal is a necessary condition for incision functions to yield kernel contractions that satisfy relevance:

Proposition 26. *If a kernel contraction function $\dot{-}_\sigma$ satisfies relevance then σ satisfies symmetry of removal.*

Proposition 26 exhibits that relevance induces a symmetry of removal between α-concordant sets. Although symmetry of removal is a necessary condition for incision functions to yield kernel contraction functions that do satisfy relevance, it is important to stress that symmetry of removal alone is not strong enough to capture relevance:

Observation 27. *If an incision function satisfies symmetry of removal then its kernel contraction* ***does not necessarily*** *satisfy relevance.*

Proof. Consider the belief base $\mathcal{K} = \{p, p \vee q, p \rightarrow m, p \vee q \rightarrow m\}$, and the following incision function σ satisfying symmetry of removal: $\sigma(\alpha) = \{p \vee q, p \rightarrow m, p \vee q \rightarrow m\}$, if $\mathcal{K} \perp\!\!\!\perp \alpha = \mathcal{K} \perp\!\!\!\perp m$; and $\sigma(\alpha) = \bigcup \mathcal{K} \perp\!\!\!\perp \alpha$, otherwise. For violation of relevance, note that $\mathcal{K} \dot{-}_\sigma m = \{p\}$, while $p \vee q \in \sigma(m)$. $\square$

We need both *symmetry of removal* and *smoothness* in order to capture relevance:

Proposition 28. *If an incision function σ satisfies smoothness and symmetry of removal then the smooth kernel contraction function $\dot{-}_\sigma$ satisfies relevance.*

Our first representation theorem relating relevance and kernel contraction functions follows from Proposition 26 and Proposition 28 :

Theorem 29. *A kernel contraction function satisfies relevance iff its incision function satisfies both smoothness and symmetry of removal.*

Theorem 29 states that the only way for kernel contractions to satisfy relevance is via incision functions satisfying both smoothness and symmetry of removal. Although this informs us which principles an incision function must satisfy, it does not relate such incision functions to the agent's epistemic preference relations. We already know from Section 3 that the choices made by smooth incision functions are based on epistemic preference relations founded on spallings. Thus, to capture relevance, we only need to constrain the behaviour of such spallings. The condition we will put upon the spallings must be sufficiently strong to capture *symmetry of removal*, and at the same time it must be general enough so it dispense with no spallings satisfying *symmetry of removal*. The condition we devise for this purpose is called *concordant-mirroring*, and spallings satisfying this condition will be called mirrored spallings:

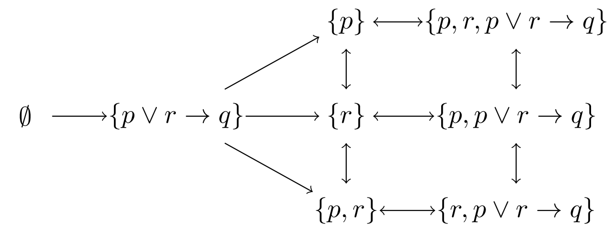

Figure 3: An α-shard satisfying *concordant-mirroring*. The relation is transitive, but to avoid visual pollution we omit edges obtained by transitivity.

Definition 30. *A spalling $\tau : \mathcal{K} \rightarrow \mathcal{P}(\mathcal{P}(\mathcal{K}) \times \mathcal{P}(\mathcal{K}))$ is mirrored iff for each formula α, the relation $\leqslant^\tau_\alpha$ satisfies*

***concordant-mirroring**: if A and B are α-concordant, and $X \leqslant^\tau_\alpha A$ then $X \leqslant^\tau_\alpha B$*

Intuitively, *concordant-mirroring* states that if two sets, say A and B, are α-concordant then they must be as resistant (resp. as vulnerable) as each other. Example IV below illustrates the notion of *concordant-mirroring*:

Example IV. *Let $\mathcal{K} = \{p, r, p \vee r \rightarrow q\}$, and τ be a spalling such that $\leqslant^\tau_\alpha$ is the q-shard depicted in Fig. 3. Except for $\{p \vee r \rightarrow q\}$ and the empty-set, all other sets are equally preferable, and are not resistant. To facilitate visualisation, q-concordant sets are put in the same column. It is easy to see in this way that $\leqslant_q$ satisfies concordant-mirroring. For instance, the sets $\{p\}, \{r\}$ and $\{p, r\}$ are all q-concordant. According to concordant-mirroring all of them must present the same preferences. Thus, as $\{p \vee r \rightarrow q\} \leqslant_q \{p\}$, the sets $\{r\}$ and $\{p, r\}$ must mimic such a preference, that is, $\{p \vee r \rightarrow q\} \leqslant_q \{r\}$ and $\{p \vee r \rightarrow q\} \leqslant_q \{p, r\}$.*

Note also that the only two resistant sets are $\{p \vee r \rightarrow q\}$ and $\emptyset$. Thus, $\delta_\tau(q) = \{p, r\}$, and $\mathcal{K} \dot{-}_\tau q = \{p \vee r \rightarrow q\}$. Note that this contraction satisfies relevance, as putting back either of the two removed formulae will restore q.

An effacing built upon a mirrored spalling will be called a mirrored effacing, and a kernel contraction built upon a mirrored effacing will be called a mirrored kernel contraction.

Definition 31. *A mirrored kernel contraction is a kernel contraction $\dot{-}_\tau$ such that δ_τ is a mirrored effacing.*

Example IV illustrates a mirrored kernel contraction satisfying relevance. Indeed, *concordant-mirroring* is all we need to capture symmetry of removal, and therefore relevance.

Theorem 32. *Mirrored effacings satisfy symmetry of removal, and every mirrored kernel contraction satisfies relevance.*

We have reached the first part of our second representation theorem between relevance and kernel contractions. To complete it, we need to show that *concordant-mirroring* preserves all fully rational kernel contractions. This corresponds to showing that the incision function of every fully rational kernel contraction is actually a mirrored effacing. We achieve this by employing the same strategy we used in Section 3 based on shadowing. Recall that the incision function σ of a fully rational kernel contraction function is smooth. Therefore, we already know from Proposition 22 and Corollary 21 that σ is indeed an effacing whose spalling is its shadowing $\mathcal{T}_\sigma$. The only thing missing is to show that the shadowing $\mathcal{T}_\sigma$ indeed satisfies *concordant-mirroring*:

Proposition 33. *If a smooth kernel contraction function $\dot{-}_\sigma$ satisfies relevance then the shadowing of σ is mirrored.*

We have already illustrated, in Example II, an α-projection from the shadowing of a smooth incision function whose kernel contraction satisfies relevance. Such an α-projection is shown in Fig. 2, and it does satisfy *concordant-mirroring*. We reach our second representation theorem:

Theorem 34. *A smooth kernel contraction satisfies relevance iff its a mirrored kernel contraction.*

5 Related Works

Few efforts have been made to investigate the connection between kernel contraction and relevance. Falappa, Fermé, and Kern-Isberner (2006) investigate the explicit connection between partial meet functions and kernel contraction functions in terms of their choice mechanisms: incision functions, for kernel contraction; and selection function for partial meet functions. The work focuses on translating selection functions and incision functions in terms of each other. For this inter-translation to work, it is assumed that the underlying smooth kernel contraction indeed satisfies relevance. However, it is not investigated which condition an incision function must satisfy in order to capture relevance. In our work, we fill this gap by showing precisely such a condition: *symmetry of removal*. Moreover, their correspondence between smooth incision functions and selection functions are confined to classical proposition logics, whereas our results are presented for more general logics, precisely for Tarskian compact logics. Booth et al. (2014) show that in Horn Logics (Horn, 1951), partial meet is too strong for contraction, and they introduce a more general operation, called infra-contraction, which coincides with kernel contraction. However, infra-contractions do not capture relevance.

For theories, several ways of specifying epistemic preferences are known, such as Epistemic Entrenchment (Gärdenfors, 1988), and Grove's system of spheres (Grove, 1988; Gärdenfors, 1988). However, few efforts have been made to construct epistemic preferences for belief bases. There are two main contraction operators that consider epistemic preferences for belief bases: *safe-contraction* (Alchourrón and Makinson, 1985) and *ensconcements* (Williams, 1994). Kernel contractions are generalisations of *safe-contractions*, in which epistemic preferences are represented as a single hierarchy: an acyclic strict total order on the formulae of a belief base. Safe-contractions, however, are too restrictive and, even for theories, there are fully rational contractions that cannot be defined as safe-contractions. This occurs not only because of the impossibility results listed in Section 3, but also because safe-contraction cannot handle cycles. Williams (1994) has introduced the ensconcements relations, which are in essence generalisations of epistemic entrenchment relations, originally defined for contraction on theories, to deal with belief bases. Fermé, Krevneris, and Reis (2008) have shown that ensconcements are so general that they not only violate *relevance*, but also they violate *core-retainment*.

Concordant-mirorring, introduced in Section 4, is inspired by the *mirroring* condition introduced by Ribeiro, Nayak, and Wassermann (2018). In that work, mirroring is introduced with the purpose to rationalise epistemic preference relations, over interpretations, on non-compact logics. While mirroring enforces that any two incomparable models must mimic each other's preferences, in our setting *concordant-mirroring* is a bit more relaxed and imposes the mimicking behaviour only on α-concordant sets. Mirroring is suitable for contraction on theories and was designed to capture the supplementary postulates of the AGM paradigm (Alchourrón, Gärdenfors, and Makinson, 1985).

6 Discussion and Future Works

In this work, we have restored the connection between kernel contraction and the postulate of relevance by constraining the behaviour of incision functions in two different, but equivalent, ways: (i) via the principle of *symmetry of removal*, and (ii) based on a suitable class of epistemic preference relations. For the former, we have shown that smooth incision functions can only satisfy relevance when they obey *symmetry of removal*. For the latter, we have started by proposing to represent an agent's epistemic preferences via binary relations on the subsets of its belief base. We have then defined the *mirrored-concordance* condition that translates symmetry of removal in terms of epistemic preferences, and shown a representation theorem connecting relevance and kernel contractions based on such preference relations. Our results are presented for logics that are Tarskian and compact. There are further interesting research questions worth to explore:

Smoothness and symmetry of removal: when an agent's epistemic state is represented as a theory, smooth kernel contractions are strong enough to capture relevance. Although this is not the case for belief bases, it is not hard to find belief bases where smoothness is still capable of capturing relevance. For instance, for the belief base $\mathcal{K} = \{p, r, p \vee r \rightarrow q, r \rightarrow q\}$, the contraction $\mathcal{K} \dot{-}_\sigma q$ satisfies relevance, as long as σ is smooth. Thus, another way of establishing the connection between kernel contraction and relevance is to frame precisely the class of belief bases in which smoothness implies *symmetry of removal*.

Inconsistency/Culpability measures: Recently, Ribeiro and Thimm (2021) have proposed to use inconsistency/culpability measures as a means to automatically disclose an agent's epistemic preference relation for contracting inconsistencies, known as *consolidation*. Their results are confined to smooth incision functions, and therefore are not strong enough to capture relevance. It will be worth to investigate how such inconsistency/culpability measures could be strengthened in order to capture relevance, and how such measures could be broadened to perform contraction beyond consolidation and classical propositional logics.

Acknowledgements

This research is supported by the German Research Association (DFG), project number 424710479. I would like to thank Dominik Klumpp for suggesting an intuitive reformulation of α-concordant sets.

References

Alchourrón, C. E., and Makinson, D. 1985. On the logic of theory change: Safe contraction. *Stud Logica* 44(4):405–422.

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.

Booth, R.; Meyer, T.; Varzinczak, I. J.; and Wassermann, R. 2014. On the link between partial meet, kernel, and infra contraction and its application to horn logic. *CoRR* abs/1401.3902.

Cóbe, R., and Wassermann, R. 2015. Ontology repair through partial meet contraction. In *DARe@IJCAI*, volume 1423 of *CEUR Workshop Proceedings*. CEUR-WS.org.

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.

Dixon, S. E. 1994. *Belief revision: A computational approach*. Ph.D. Dissertation, 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, 402–406. IOS Press.

Fermé, E.; Krevneris, M.; and Reis, M. D. L. 2008. An axiomatic characterization of ensconcement-based contraction. *J. Log. Comput.* 18(5):739–753.

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.

Grove, A. 1988. Two modellings for theory change. *J. Philos. Log.* 17(2):157–170.

Guerra, P. T., and Wassermann, R. 2019. Two AGM-style characterizations of model repair. *Annals of Mathematics and Artificial Intelligence* 87(3):233–257.

Hansson, S. O., and Wassermann, R. 2002. Local change. *Stud Logica* 70(1):49–76.

Hansson, S. O. 1991. Belief contraction without recovery. *Stud Logica* 50(2):251–260.

Hansson, S. O. 1994. Kernel contraction. *J. Symb. Log.* 59(3):845–859.

Hansson, S. O. 1999. *A textbook of belief dynamics - theory change and database updating*, volume 11 of *Applied logic series*. Kluwer.

Hansson, S. O. 2017. *Descriptor Revision: Belief Change Through Direct Choice*, volume 46. Springer.

Horn, A. 1951. On sentences which are true of direct unions of algebras. *J. Symb. Log.* 16(1):14–21.

Horridge, M. 2011. *Justification based explanation in ontologies*. Ph.D. Dissertation, University of Manchester, UK.

Hunter, A., and Konieczny, S. 2005. *Approaches to Measuring Inconsistent Information*. Berlin, Heidelberg: Springer Berlin Heidelberg. 191–236.

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.

Kalyanpur, A.; Parsia, B.; Horridge, M.; and Sirin, E. 2007. Finding all justifications of OWL DL entailments. In *ISWC/ASWC*, volume 4825 of *Lecture Notes in Computer Science*, 267–280. Springer.

Katsuno, H., and Mendelzon, A. O. 2003. On the difference between updating a knowledge base and revising it. In Gärdenfors, P., ed., *Belief revision*, 183–203. Cambridge University Press.

Nebel, B. 1990. *Reasoning and Revision in Hybrid Representation Systems*, volume 422 of *Lecture Notes in Computer Science*. Springer.

Ribeiro, J. S., and Thimm, M. 2021. Consolidation via tacit culpability measures: Between explicit and implicit degrees of culpability. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning, KR 2021*, 529–538.

Ribeiro, M. M.; Wassermann, R.; Flouris, G.; and Antoniou, G. 2013. Minimal change: Relevance and recovery revisited. *Artif. Intell.* 201:59–80.

Ribeiro, J. S.; Nayak, A.; and Wassermann, R. 2018. Towards belief contraction without compactness. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, 287–296. AAAI Press.

Simari, G. R., and Loui, R. P. 1992. A mathematical treatment of defeasible reasoning and its implementation. *Artif. Intell.* 53(2-3):125–157.

Williams, M. 1994. On the logic of theory base change. In *Logics in Artificial Intelligence, European Workshop, JELIA '94, Proceedings*, volume 838 of *Lecture Notes in Computer Science*, 86–105. Springer.