Towards a Semantic Construction for Belief Base Contraction (A Preliminary Report)

Jandson S. Ribeiro*

University of Hagen, Germany

Abstract

We introduce a novel class of fully rational contraction operators that work by selecting models, which we call tracked contraction operators. These operators are founded on a plausibility relation on models, called a track, that allows distinguishing between suitable and unsuitable models. We show a representation theorem between tracked contraction operators and the basic rationality postulates of contraction. For the supplementary postulates (conjunction and intersection), we strengthen such operators by imposing the mirroring condition on the track relations. We consider logics that are both Tarskian and compact.

Keywords

Belief Change, Belief Contraction, Belief Base, Models

1. Introduction

The field of *Belief Change* [1, 2, 3] studies how an agent should rationally modify its corpus of beliefs in response to incoming pieces of information. The two most important kinds of change are: contraction, which relinquishes undesirable/obsolete information; and revision, which accommodates new information with the caveat of keeping the corpus of beliefs consistent. Each of these kind of changes are governed by sets of rationality postulates, split into basic and supplementary rationality postulates, which prescribe adequate behaviours of change. Such rationality postulates are motivated by the principle of minimal change: in response to a piece of information, say α, an agent should remove only beliefs that either conflict with α (in the case of revision), or that contribute to entail α (in the case of contraction).

Several classes of belief change operators were proposed that abide by such rationality postulates, called rational belief change operators (see [3], for a list). These classes of operators can be split in two main kinds: syntactic operators and semantic operators. Operators belonging to the first kind select sentences from the language, while operators of the second kind select models. Examples of syntactic operators are partial meet operators [1] and smooth kernel operators [4], while Grove's system of spheres [5, 2] and the faifthful pre-orders of Katsuno and Mendelzon [6] are the main frameworks for constructing semantic operators. In the most

8th Workshop on Formal and Cognitive Reasoning, September 19, 2022, Trier, Germany
*Corresponding author.
jandson.ribeiro@fernuni-hagen.de (J. S. Ribeiro)
https://jandsonribeiro.github.io/home/ (J. S. Ribeiro)
0000-0003-1614-6808 (J. S. Ribeiro)

fundamental case, when an agent's corpus of beliefs is represented as a logically closed set of sentences, called a theory, all these classes of operators are characterised by the rationality postulates of contraction/revision.

Theories, however, are very restrictive, as they do not distinguish between explicit and implicit beliefs. One can achieve this distinction by dropping the logical closure requirement, and simply representing an agent's corpus of beliefs as a set of sentences, called a *belief base* [3]. For bases, however, very few belief change operators are capable of satisfying the rationality postulates of belief change. Precisely, only partial meet, a syntactic operator, remains rational for belief bases [3, 4]. As a result, research on belief base change has focused on partial meet operators or other similar syntactic operators [3, 7, 8]. This poses a severe limitation in advancing belief base change, as syntactic operators are highly dependent on the assumptions made about the underlying logic used to represent an agent's knowledge, as for instance, imposing that the language is closed under classical negation [9]. By devising belief change operators via models, such conditions upon the language of the logics can be easily waived.

In this work, we devise two novel classes of semantic operators for belief base contraction. Our approach consists in imposing a pre-order, called a track, upon the models of the logics. A track indicates the most plausible models, which in turn are selected to perform a contraction. We show a representation theorem between the basic rationality postulate of belief base contraction and such novel class of contraction operators. We then impose the mirroring condition [10] upon such tracks, and we show that tracks satisfying mirroring induce belief base contraction operators that capture the supplementary postulates of belief contraction. It is worth highlighting that, except for safe contraction [11], the study of the supplementary postulates on belief bases has been neglected. As contraction is a central operation in belief change, our result can be extended to provide semantic operators for other kinds of belief change such as revision.

Road map: Section 2 introduces some basic notations and definitions that will be used throughout this work. In Section 3, we briefly review belief contraction, including both basic and supplementary rationality postulates of contraction as well as the partial meet contraction operators. For semantic operators, we review the faithful pre-orders of Katsuno and Mendelzon [6] for revision, and we translate them in terms of belief contraction. We show that such operators, though fully rational for theories, are not rational for belief bases. In Section 4, we introduce our two novel classes of contraction operators and the representation theorem connecting tracks and the basic rationality postulates of contraction. Finally, in Section 5 we conclude the work and discuss some future works. Full proofs of the results can be found in the appendix at https://jandsonribeiro.github.io/home/appendix/FCR_22_appendix.pdf

2. Notation and Technical Background

The power set of a set A is denoted by $\mathcal{P}(A)$. 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 Tarskian and satisfy compactness. A theory is a set of formulae $X \subseteq \mathcal{L}$ such that $X = Cn(X)$.

As we are interested to define semantic operators, we exploit the semantic of the logics. Given a logic $\langle \mathcal{L}, Cn \rangle$ and a set of structures $\mathcal{I}$, an interpretation or a model is an element of $\mathcal{I}$ that gives meaning to the formulae of $\mathcal{L}$; $\mathcal{I}$ is called an *interpretation domain* of that logic, whereas each subset of $\mathcal{I}$ is called an *interpretation set*. For instance, an interpretation domain for the Propositional Logic is the power set of the propositional symbols of the language. A satisfaction relation $\models \subseteq \mathcal{I} \times \mathcal{L}$ is used to indicate on which interpretations a formula is satisfied. If $M \models \alpha$, then we say that M is a model of α. If an interpretation M does not satisfy a formula α, denoted by $M \not\models \alpha$, then we say that M is a counter-model of α. The set of all models of α is given by $[\![\alpha]\!]$, while the set of all counter-models of α is given by $\overline{[\![\alpha]\!]}$.

In Tarskian logics, the consequence operator can be semantically defined as: a formula $\varphi \in Cn(X)$ iff every model that satisfies all formulae in X also satisfies φ [12]. Let $\mathcal{I}$ be an interpretation domain of a logic $\langle \mathcal{L}, Cn \rangle$, and M a model in $\mathcal{I}$. The set of all formulae of $\mathcal{L}$ satisfied by M is the theory $Th(M) = \{\varphi \in \mathcal{L} \mid M \models \varphi\}$. Generalising, given a set of models A, $Th(A) = \{\varphi \mid \forall M \in A, M \models \varphi\}$ is the theory of the formulae satisfied by all models in A. Moreover, given a set $X \subseteq \mathcal{L}$, the set of models that satisfy all formulae in X is $[\![X]\!] = \{M \in \mathcal{I} \mid \forall \varphi \in X, M \models \varphi\}$. For simplicity, given a set of formulae X and a model M, we will write $M \models X$ to mean that M satisfies every formula in X.

Throughout this paper we will provide examples to support the intuition of the proposed contraction operators. Due to its simplicity, we will use classical propositional logics to construct such examples. Observe, however, that our results are not confined to classical propositional logics. As usual, the formulae of classical propositional logics are Boolean formulae constructed from a set AP of atomic propositional symbols, via the operators of conjunction ($\wedge$), disjunction ($\vee$) and classical negation ($\neg$). The models are subsets of AP, and the satisfaction relation is defined as usual.

A pre-order on a domain $\mathcal{D}$ is binary relation $\leqslant \subseteq \mathcal{D} \times \mathcal{D}$ that satisfies transitivity and reflexivity. The minimal elements of a set $A \subseteq \mathcal{D}$ w.r.t $\leqslant$ is given by the set $\min_{\leqslant}(A) = \{a \in A \mid \text{ if } b \leqslant a \text{ then } a \leqslant b, \text{ for all } b \in A\}$. We write $a < b$ to denote that $a \leqslant b$ but $b \nleqslant a$.

3. Belief 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 [13, 14, 15], and (ii) as a more general and expressive approach that distinguishes explicit from implicit beliefs [16, 3]. We follow the latter approach, and therefore a belief base can be infinite.

Let $\mathcal{K}$ be a belief base, a contraction function for $\mathcal{K}$ is a function $\dot{-} : \mathcal{L} \to \mathcal{P}(\mathcal{L})$ 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 [17, 4]:

(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 $\mathcal{K} \dot{-} \alpha = \mathcal{K} \dot{-} \beta$;

(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 [3]. We call the set of rationality postulates listed above as the basic rationality postulates of contraction. A contraction function that satisfies all the basic rationality postulates above will be dubbed a *rational contraction function.*

There are other two postulates, called supplementary postulates [1, 3, 18]:

(***intersection***) $\mathcal{K} \dot{-} \varphi \cap \mathcal{K} \dot{-} \psi \subseteq \mathcal{K} \dot{-} \varphi \wedge \psi$

(***conjunction***) If $\varphi \notin Cn(\mathcal{K} \dot{-} \varphi \wedge \psi)$ then $\mathcal{K} \dot{-} (\varphi \wedge \psi) \subseteq \mathcal{K} \dot{-} \varphi$

It is important to stress that the study of the supplementary postulates has been confined to theories, and very little is known about their behaviours on belief bases. Rational contraction operators that satisfy the supplementary postulates will be dubbed *fully rational.*

3.1. Partial Meet Contraction

Several rational contraction operators were proposed in the literature. One of the most influential ones is partial meet (Definition 3, below), which makes use of remainders.

Definition 1. *Given a belief base $\mathcal{K}$ and formula α, an α-remainder of $\mathcal{K}$ is a set $X \subseteq \mathcal{K}$ such that: $\alpha \notin Cn(X)$, and if $X \subset Y \subseteq \mathcal{K}$, then $\alpha \in Cn(Y)$. The set of all α-remainders of $\mathcal{K}$ is denoted by $\mathcal{K} \perp \alpha$.*

Each member of $\mathcal{K} \perp \alpha$ is called a remainder, and it is a maximal subset of $\mathcal{K}$ that does not entail α. A partial meet operator works by selecting remainders and intersecting them. As a remainder set might have many remainders, a choice must be made about which ones are the best to perform the contraction. This choice is done via an extra-logical mechanism called a *selection function*:

Definition 2. *A selection function γ picks some remainder of $\mathcal{K} \perp \alpha$ such that, (i) $\gamma(\mathcal{K} \perp \alpha) \neq \emptyset$, (ii) $\gamma(\mathcal{K} \perp \alpha) \subseteq \mathcal{K} \perp \alpha$, if $\mathcal{K} \perp \alpha \neq \emptyset$; and (iii) $\gamma(\mathcal{K} \perp \alpha) = \{\mathcal{K}\}$, otherwise. A selection function γ is relational iff there is some binary relation $\leqslant$ on all remainders such that $\gamma(\mathcal{K} \perp \alpha) = \min_{\leqslant}(\mathcal{K} \perp \alpha)$, for all $\mathcal{K} \perp \alpha \neq \emptyset$. If $\leqslant$ is transitive then γ is called transitive relational.*

A selection function works as an extra-logical mechanism that realises the agent's epistemic preferences. In the original work of [1], the authors propose to represent an agent's preferences as a binary relation $\leqslant$ on all remainders. Precisely, a pair $A \leqslant B$ means that the remainder A is at least as preferable as B. The agent picks the most preferable α-remainders w.r.t $\leqslant$.

Remainder sets and selection functions are used to define a contraction operator called *partial meet contraction*:

Definition 3. *Given a belief base $\mathcal{K}$, and a selection function γ, the operation $\dot{-}_\gamma$ defined as $K \dot{-}_\gamma \alpha = \bigcap \gamma(K \perp \varphi)$ is a partial meet contraction function.*

Theorem 4. *[19] A contraction operator is rational iff it is a partial meet contraction operator.*

For theories, the transitive relational partial meet operators are characterised by all the rationality postulates of contraction.

Theorem 5. *[1] On theories, a contraction operator is fully rational iff it is a transitive relational partial meet contraction operator.*

As Hansson [18] shows, the transitive relational partial meet operators are not strong enough to satisfy the two supplementary postulates on belief bases. Hansson proposed to strengthen the transitive relations with a property called maximising. However, a representation theorem was not obtained.

3.2. Semantic Contraction Operators

We start by explaining how belief contraction works on models when the agent's corpora of beliefs are represented as theories. After that, we show why such strategies do not work for belief bases.

In terms of models, in order to contract a formula α from a theory $\mathcal{K}$, it suffices to obtain a theory that is a subset of $\mathcal{K}$ (due to the inclusion postulate) and it is satisfied by some counter-models of α. This can be formalised by taking a function $\sigma : \mathcal{L} \to \mathcal{P}(\mathcal{I})$ that picks, for every non-tautological formula α, some counter-models of α. For tautological formulae α, we make $\sigma(\alpha) = \emptyset$, as tautologies have no counter-models. Moreover, if two formulae α and β are logically equivalent, then $\sigma(\alpha) = \sigma(\beta)$. This guarantees that the choice function is not syntax sensitive. We say that σ is a model choice function.

Definition 6. *The contraction function induced by a model choice function σ is the operator $\mathcal{K} \dot{-}_\sigma \alpha = \{\varphi \in \mathcal{K} \mid \sigma(\alpha) \models \varphi\}$.*

Indeed, the basic rationality postulates of contraction characterise such class of semantic contraction operators for theories:

Theorem 7. *[10, 12] In classical propostionl logics, a contraction function $\dot{-}$ on a theory $\mathcal{K}$ is rational iff it is induced by some model choice function σ.*

For full rationality, there are two main classes of belief operators: the revision operators based on faithful pre-orders of Katsuno and Mendelzon (KM, for short) [6] and the revision operators based on Grove's spheres[5]. Although both classes of operators were originally framed for belief revision, they can be easily translated to contraction. In the following, we present a translation of KM operators based on faithful pre-orders in terms of contraction:

Definition 8. *[6][1] Given a belief base $\mathcal{K}$, a pre-order $\leqslant_{\mathcal{K}}$ is faithful w.r.t $\mathcal{K}$ iff it satisfies the two following conditions: (1) if $M, M' \in [\![\mathcal{K}]\!]$ then $M \not<_{\mathcal{K}} M'$; (2) if $M \in [\![\mathcal{K}]\!]$ and $M' \notin [\![\mathcal{K}]\!]$ then $M <_{\mathcal{K}} M'$.*

Definition 9. *Given a faithful pre-order $\leqslant_{\mathcal{K}}$ on a belief base $\mathcal{K}$, the faithful contraction operator founded on $\leqslant_{\mathcal{K}}$ is the operation $\dot{-}_{\leqslant_{\mathcal{K}}}$ such that $[\![\mathcal{K} \dot{-}_{\leqslant_{\mathcal{K}}} \alpha]\!] = [\![\mathcal{K}]\!] \cup \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$. If $\leqslant_{\mathcal{K}}$ is total then $\dot{-}_{\leqslant_{\mathcal{K}}}$ is a total faithful contraction operator.*

A faithful pre-order works as an epistemic preference relation on models. In order to contract a formula α, the agent chooses exactly the most plausible counter-models of α. In the current presentation, KM operators are suitable only for theories, because, for belief bases, there is no guarantee that $\mathcal{K} \dot{-}_{\leqslant_{\mathcal{K}}} \alpha$ outputs a subset of $\mathcal{K}$, as the inclusion postulate demands. Towards this end, in order to satisfy the inclusion postulate we need only to rewrite faithful contraction in the spirit of Definition 6: get the greatest subset of $\mathcal{K}$ satisfied by the minimal counter-models of the formula α to be contracted. Indeed, within classical propositional logics, each KM operation is a special kind of contraction induced by a model choice function as per Definition 6. In classical propositional logics, for theories, the faithful contraction operators on total pre-orders are fully rational:

Theorem 10. *In classical propositional logics, a contraction operator on a theory $\mathcal{K}$ is fully rational iff it is a total faithful contraction operator.*

Caridroit et al. [20] obtain an analogous of Theorem 10 via Levi and Harper identities on the KM faithful revision operators. Observe that the representation theorems above (Theorem 7 and Theorem 10) are established only for theories. Indeed, as Example I below illustrates, both representation theorems breaks down for bases, which is due to violation of the relevance postulate.

Example I. *Consider the belief base $\mathcal{K} = \{p, q, p \vee q, \neg q \vee p\}$, expressed in classical propositional logics, with $AP = \{p, q\}$. We want to contract the formula $p \wedge q$. There are only three rational contraction results: $A_1 = \{p, p \vee q, \neg q \vee p\}$, $A_2 = \{q, p \vee q\}$ and $A_3 = \{p \vee q\}$. Not every model choice function, however, induces a rational contraction operator. To see this, note that we have only four models $M_1 = \{p, q\}$, $M_2 = \{p\}$, $M_3 = \{q\}$ and $M_4 = \emptyset$. Observe that $\overline{[\![p \wedge q]\!]} = \{M_2, M_3, M_4\}$. Let $\leqslant_{\mathcal{K}}$ be the following faithful pre-order on $\mathcal{K}$: $M_1 \leqslant_{\mathcal{K}} M_4 \leqslant_{\mathcal{K}} M_3 \leqslant_{\mathcal{K}} M_2$. Let σ be a model choice function such that $\sigma(p \wedge q) = \min_{\leqslant_{\mathcal{K}}}(\overline{[\![p \wedge q]\!]}) = \{M_4\}$. The only formula of $\mathcal{K}$ that M_4 satisfies is $\neg q \vee p$. Thus, $\mathcal{K} \dot{-}_{\sigma} p \wedge q = \{\neg q \vee p\}$. However, this does not correspond to any of the three possible rational contraction results listed above.*

4. Tracks and Mirrors: Belief Base Contraction on Models

In this section, we provide a novel class of semantic contraction operators for belief bases.

In terms of models, contracting a formula α from a theory $\mathcal{K}$ consists in picking some counter-models of α and maintaining the formulae in $\mathcal{K}$ satisfied by all such picked counter-models.

[1] Originally, KM defines an assignment that maps each formula to a pre-order, and defines such an assignment to be faithful. This assignment has only the purpose to provide general contraction operators. As here we focus on local contraction, we opt to remove this complication and operate directly on the pre-orders.

While this strategy yield rational contractions for theories (Theorem 7), it fails for belief bases as Example I illustrates. This occurs because some counter-models of α might satisfy less formulae than allowed by the relevance postulate. For instance, looking back at Example I, according to relevance the formula $p \vee q$ must be kept. Observe that this formula appears in all the three possible rational contraction results. The counter-model M_4, however, does not satisfy $p \vee q$, which makes it unsuitable for performing a rational contraction, as picking it would remove $p \vee q$. The main hurdle is to properly distinguish between suitable and unsuitable models. To solve this problem, we establish a plausibility relation $\leqslant$ on the models. Intuitively, a pair $M \leqslant M'$ means that the model M is at least as plausible as M'. Towards this end, in order to contract a formula α, only the most plausible counter-models of α w.r.t $\leqslant$ should be chosen, that is, only models within $\min_{\leqslant}(\overline{[\![\alpha]\!]})$. The question at hand is which properties a pre-order on models should satisfy in order to be an adequate plausibility relation that distinguish between suitable and unsuitable models.

Here, we propose such plausibility relations to be defined upon the notion of information preservation. Intuitively, the more information from $\mathcal{K}$ a model preserves the more plausible it is. The set of all formulae from $\mathcal{K}$ satisfied by a model M is given by the set $Pres(M \mid \mathcal{K}) = \{\varphi \in \mathcal{K} \mid M \models \varphi\}$. Definition 11 below formalises a class of pre-orders based on this notion, which we call tracks.

Definition 11. *A track of a belief base $\mathcal{K}$ is a pre-order $\leqslant_{\mathcal{K}} \subseteq \mathcal{I} \times \mathcal{I}$ such that*

1. *If $Pres(M \mid \mathcal{K}) = Pres(M' \mid \mathcal{K})$ then $M' \leqslant_{\mathcal{K}} M$ and $M \leqslant_{\mathcal{K}} M'$; and*
2. *If $Pres(M \mid \mathcal{K}) \subset Pres(M' \mid \mathcal{K})$ then $M' <_{\mathcal{K}} M$*

In short, a track relation imposes models that strictly preserve more information to be strictly more plausible (condition 2), while models that preserve the same set of information are equally plausible (condition 1). Thus, in every track for a belief a base $\mathcal{K}$, the models of $\mathcal{K}$ are the most plausible ones, and they are also all equally plausible.

Proposition 12. *If $\mathcal{K}$ is a consistent belief base and $\leqslant_{\mathcal{K}}$ is a track of $\mathcal{K}$ then $\min_{\leqslant_{\mathcal{K}}}(\mathcal{I}) = [\![\mathcal{K}]\!]$.*

The less pairs a track contains, the more permissive it is. A most permissive track of a belief base $\mathcal{K}$ will be called a least track of $\mathcal{K}$. It is easy to see that each belief base has exactly one least track.

Definition 13. *A least track of a belief base $\mathcal{K}$ is a track $\leqslant_{\mathcal{K}}^{-}$ such that for every track $\leqslant_{\mathcal{K}}$ of $\mathcal{K}$, $\leqslant_{\mathcal{K}}^{-} \subseteq \leqslant_{\mathcal{K}}$*

Observation 14. *Every belief base has a unique least track $\leqslant_{\mathcal{K}}^{-}$.*

Example II (continued from Example I)**.** *The beliefs in $\mathcal{K} = \{p, q, p \vee q, p \vee \neg q\}$ preserved by each of the four models are:*

$$Pres(M_1 \mid \mathcal{K}) = \mathcal{K} \qquad Pres(M_2 \mid \mathcal{K}) = \{p, p \vee q, p \vee \neg q\}$$
$$Pres(M_3 \mid \mathcal{K}) = \{q, p \vee q\} \qquad Pres(M_4 \mid \mathcal{K}) = \{\neg q \vee p\}$$

Fig. 1 (on the right) illustrates the set inclusion relation between the preservation sets of each model, while Fig. 1 (on the left) depicts the least track relation of $\mathcal{K}$. As M_1 is the only model of $\mathcal{K}$, it is strictly more plausible than all other models. Models M_2 and M_3 are incomparable, since they preserve different beliefs in $\mathcal{K}$. For the same reason, M_4 and M_3 are incomparable. However, M_2 is strictly more plausible than M_4, as M_2 preserves strictly more information than M_4.

At this point, we can see that a track can distinguish between suitable and unsuitable models. According to this track, both models M_2 and M_3 are the most plausible counter-models of $p \wedge q$. If we choose either M_2 or M_3 then we get a rational contraction: either $A_1 = \{p, p \vee q, \neg q \vee p\}$, or $A_2 = \{q, p \vee q\}$. By picking both models we get the last rational contraction $A_3 = \{p \vee q\}$. The only not rational contraction are those involving the model M_4 which is not among the most plausible one (the suitable ones). Also observe that other tracks exist: for instance augmenting the illustrated track by making M_2 and M_3 comparable or even M_3 and M_4 comparable. However, for any of the possible tracks, M_4 is never among the suitable ones, as it must be strictly less plausible than M_2, due to condition 2 of track's definition. This suggests that tracks can be used as an adequate class of plausibility relations to distinguish between suitable and unsuitable models.

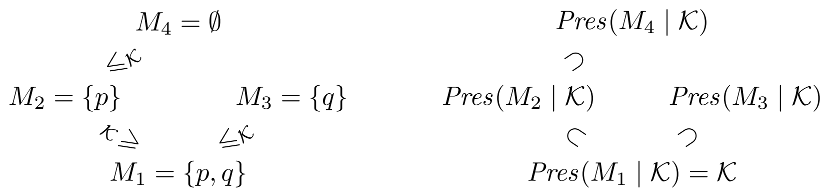

Figure 1: The least track relation $\leqslant_{\mathcal{K}}$(on the left), and the set inclusion relation on the preservation set of the models (on the right).

As tracks establish an adequate notion of plausibility between models, then most plausible ones to contract a formula α are the minimal counter-models of α. In classical propositional logics, such minimal models always exist, as there is only a finite number of models. However, for more expressive logics, such as First Order Logics and several Description Logics [21], there are formulae with an infinite number of (counter-)models. In the presence of an infinite amount of models, some tracks arrange the models through infinite chains. In general, these infinite chains prevents identifying the most plausible counter-models for some formulae. Thus, we need to constrain ourselves to tracks that do not present such bad behaviour, that is, tracks that are *founded*:

Definition 15. *A relation $\leqslant \subseteq \mathcal{I} \times \mathcal{I}$ is founded iff $\min_{\leqslant}(\overline{[\![\alpha]\!]}) \neq \emptyset$ for every non-tautological formula α.*

Relying on founded tracks guarantees that for every non-tautological formula α, there is at least one counter-model to be picked to perform such a contraction. In fact, as long as the underlying Tarskian logic satisfies compactness, every belief base presents at least one founded track: its least track.

Theorem 16. *If a logic $\langle \mathcal{L}, Cn \rangle$ is Tarskian and compact then for every belief base $\mathcal{K} \subseteq \mathcal{L}$, the least track is founded.*

We can then define a function that selects among the most plausible models:

Definition 17. *Let $\leqslant_{\mathcal{K}}$ be a founded track. A tracking selection function on $\leqslant_{\mathcal{K}}$ is a function $\delta_{\leqslant_{\mathcal{K}}} : \mathcal{L} \to \mathcal{P}(\mathcal{I})$ such that*

1. $\delta_{\leqslant_{\mathcal{K}}}(\alpha) \subseteq \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$
2. $\delta_{\leqslant_{\mathcal{K}}}(\alpha) \neq \emptyset$, *if α is not a tautology*
3. *if α and β are logically equivalent then $\delta_{\leqslant_{\mathcal{K}}}(\alpha) = \delta_{\leqslant_{\mathcal{K}}}(\beta)$.*

A tracking selection function works similarly to the model choice function for theories. The main difference is that model choice functions can choose any counter-model of a formula α, while tracking selection functions choose only among the most plausible (w.r.t a track relation) counter-models of α. Condition 3 is related to the postulate of uniformity, and guarantees that a tracking selection function is not syntax sensitive. When it is clear from context, we drop the subscript $\leqslant_{\mathcal{K}}$ and simply write δ.

Following the same strategy as for theories, a contraction on a belief base is performed by keeping the formulae from the current belief base that are satisfied by all the counter-models selected by a tracking selection function.

Definition 18. *Let δ be a tracking selection function. The tracked contraction founded on δ is defined as*

$$\mathcal{K} \dot{-}_{\delta} \alpha = \{\varphi \in \mathcal{K} \mid \delta(\alpha) \models \varphi\}.$$

Example III (continued from Example II)**.** *Let $\leqslant_{\mathcal{K}}^{-}$ be the least track of the belief base $\mathcal{K} = \{p, q, p \vee q, \neg q \vee p\}$. Observe that $\min_{\leqslant_{\mathcal{K}}^{-}}(p \wedge q) = \{M_2, M_3\}$. Then, we can choose any combination of M_2 and M_3 to contract $p \wedge q$. Let δ_1, δ_2 and δ_3 be tracked selection functions founded on $\leqslant_{\mathcal{K}}^{-}$ such that $\delta_1(p \wedge q) = \{M_2\}, \delta_2(p \wedge q) = \{M_3\}$ and $\delta_3(p \wedge q) = \{M_2, M_3\}$. They induce the following tracked contraction operators: $\mathcal{K} \dot{-}_{\delta_1} \neg q \vee p = \{p, p \vee q, \neg q \vee p\}$, $\mathcal{K} \dot{-}_{\delta_2} \neg q \vee p = \{q, p \vee q\}$, and $\mathcal{K} \dot{-}_{\delta_3} \neg q \vee p = \{p \vee q\}$. As one can easily check, each one of them is a rational contraction operator.*

Theorem 19. *Every tracked contraction function is rational.*

Proof sketch. Postulates of *success, inclusion, vacuity* and *uniformity* are easy to prove. We focus on relevance. Let $\beta \in \mathcal{K} \setminus (\mathcal{K} \dot{-}_{\delta_{\leqslant}} \alpha)$. Thus, there is some model $M \in \delta_{\leqslant_{\mathcal{K}}}(\alpha)$ such that $M \not\models \beta$. As $M \in \delta_{\leqslant_{\mathcal{K}}}(\alpha)$, we have that $M \models \mathcal{K} \dot{-}_{\delta_{\leqslant}} \alpha$ and $M \in \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$. Thus, $\mathcal{K} \dot{-}_{\delta_{\leqslant}} \alpha \subseteq Pres(M \mid \mathcal{K}) \subseteq \mathcal{K}$. Let us suppose for contradiction that $\alpha \notin Cn(Pres(M \mid \mathcal{K}) \cup \{\beta\})$. Thus, there is some model $M' \in \overline{[\![\alpha]\!]}$ such that $M' \models Pres(M \mid \mathcal{K}) \cup \{\beta\}$. This means that, $Pres(M \mid \mathcal{K}) \subset Pres(M' \mid \mathcal{K})$ which implies that $M' <_{\mathcal{K}} M$. Therefore, $M \notin \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$, which is a contradiction. □

Theorem 20. *Every rational base contraction function is a tracked contraction function.*

Since a track establishes a plausibility relation between models, it is natural to expect that a track also works as an epistemic preference relation. Therefore, instead of simply picking some of the most plausible models w.r.t a track, it would be rational to pick all such most plausible models. We will call contraction operators that follow this strategy full tracked contraction:

Definition 21. *Let* $\leqslant_{\mathcal{K}}$ *be a founded tracking of a belief base* $\mathcal{K}$*. The full tracked selection of* $\leqslant_{\mathcal{K}}$ *is the function* $\mu_{\leqslant_{\mathcal{K}}}$ *such that* $\mu_{\leqslant_{\mathcal{K}}}(\alpha) = \min_{\leqslant_{\mathcal{K}}}(\overline{[\![\alpha]\!]})$*. Tracked contraction operators founded on full tracking selection functions are full tracked contraction operators.*

Full tracked contraction operators do satisfy *intersection*, due to the transitivity of tracks.

Theorem 22. *Every full tracked contraction satisfies intersection.*

Although tracks capture *intersection*, they are not strong enough to capture *conjunction*. Observe that tracks form a special case of faithful pre-orders (Definition 9). It would be natural then to simply impose totality upon the tracks in the hope of capturing *conjunction*. Totality, however, has been criticised in the literature for being too demanding, as an agent might be indifferent or ignorant on how to grade some of its beliefs [22]. Moreover, works such as [10, 23] have observed that totality is not strong enough to capture *conjunction*, even for theories, in more expressive logics. As a solution, Ribeiro et al. [10] has introduced mirroring:

mirroring: if $A \nleqslant B$ and $B \nleqslant A$ but $C \leqslant A$ then $C \leqslant B$.

According to mirroring, if two models are incomparable then they should agree upon their preferences. We will show here that by employing mirroring upon tracks, *conjunction* is also captured for belief bases.

Theorem 23. *If a founded track satisfies mirroring then its full tracked contraction operator satisfies conjunction.*

5. Conclusion and Future Works

While both syntactic and semantic operators are well known for belief theory contraction (and other forms of belief change), only syntactic operators are known to be rational on belief bases. In this work, we have introduced two new classes of semantic contraction operators for belief bases: *tracked contraction operators* and *full tracked contraction operators* on mirroring. These operators rely on plausibility relations between models, called tracks. In order to contract a formula α, the agent seeks the most plausible counter-models of α w.r.t a track relation, and chooses some of these counter-models (the ones the agent deems most reliable). We have established a representation theorem between tracked contraction operators and the basic rationality postulates of contraction. A track unveils an agent's epistemic preferences: the most plausible models coincides with the most reliable ones, and the agent picks all these models. Tracked contractions following this strategy are called full tracked contraction. We have shown that tracks that satisfy the mirroring condition yield full tracked contraction satisfying the two supplementary postulates.

As future work we shall investigate if mirroring suffices to establish a representation theorem between full tracked contractions and the supplementary postulates. This connection with the supplementary postulates is important, because the study of such postulates has been restricted to belief change operators on theories. Particularly, the connection between contraction operators and the supplementary postulates has been established via epistemic preferences relations such as Epistemic Entrenchment[2] and Hierarchies (for safe contraction)[3]. Although all such epistemic preferences work well for theories, their connection with such rationality postulates easily disappears for bases. The only known exception is safe contraction, which still connects with the supplementary postulates only when a base $\mathcal{K}$ is finite and it is as expressive as a its theory $Cn(\mathcal{K})$: for each formula $\alpha \in Cn(\mathcal{K})$ there is a formula in $\mathcal{K}$ logically equivalent to α.

We shall extend our results for more expressive logics by dispensing with compactness and widening our results to Tarskian logics. Although we have focused on contraction, our results can be easily translated to revision: instead of selecting counter-models, one needs only to select models of the formulae α to be revised.

Acknowledgments

The author is supported by the German Research Association (DFG), project number 424710479.

References

[1] C. E. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: partial meet contraction and revision functions, The Journal of Symbolic Logic 50 (1985) 510–530.

[2] P. Gärdenfors, Knowledge in flux: Modeling the dynamics of epistemic states., The MIT press, 1988.

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

[4] S. O. Hansson, Kernel contraction, J. Symb. Log. 59 (1994) 845–859.

[5] A. Grove, Two modellings for theory change, J. Philos. Log. 17 (1988) 157–170.

[6] H. Katsuno, A. O. Mendelzon, Propositional knowledge base revision and minimal change, Artificial Intelligence 52 (1991) 263–294.

[7] J. S. Ribeiro, M. Thimm, 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, 2021, pp. 529–538.

[8] J. S. Ribeiro, Kernel contraction and the order of relevance, in: G. Kern-Isberner, G. Lakemeyer, T. Meyer (Eds.), Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel. July 31 - August 5, 2022, 2022.

[9] M. M. Ribeiro, R. Wassermann, G. Flouris, G. Antoniou, Minimal change: Relevance and recovery revisited, Artif. Intell. 201 (2013) 59–80.

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

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

[12] J. S. R. Santos, Belief change without compactness, Ph.D. thesis, University of São Paulo, 2020.

[13] B. Nebel, Reasoning and Revision in Hybrid Representation Systems, volume 422 of *Lecture Notes in Computer Science*, Springer, 1990.
[14] S. E. Dixon, Belief revision: A computational approach, Ph.D. thesis, University of Sydney, 1994.
[15] M. Dalal, Investigations into a theory of knowledge base revision, in: Proceedings of the 7th National Conference on Artificial Intelligence, AAAI Press / The MIT Press, 1988, pp. 475–479.
[16] A. Fuhrmann, Theory contraction through base contraction, J. Philos. Log. 20 (1991) 175–203.
[17] S. O. Hansson, Belief contraction without recovery, Stud Logica 50 (1991) 251–260.
[18] S. O. Hansson, Reversing the levi identity, J. Philos. Log. 22 (1993) 637–669.
[19] S. O. Hansson, R. Wassermann, Local change, Stud Logica 70 (2002) 49–76.
[20] T. Caridroit, S. Konieczny, P. Marquis, Contraction in propositional logic, Int. J. Approx. Reason. 80 (2017) 428–442. doi:10.1016/j.ijar.2016.06.010.
[21] F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, Cambridge University Press, 2017.
[22] T. A. Meyer, W. A. Labuschagne, J. Heidema, Refined epistemic entrenchment, J. Log. Lang. Inf. 9 (2000) 237–259.
[23] J. S. Ribeiro, A. Nayak, R. Wassermann, Belief update without compactness in non-finitary languages, in: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019, ijcai.org, 2019, pp. 1858–1864.