Towards Belief Contraction without Compactness

Jandson S. Ribeiro
Macquarie University, Australia
University of São Paulo, Brazil
jandson@ime.usp.br

Abhaya Nayak
Macquarie University, Australia
abhaya.nayak@mq.edu.au

Renata Wassermann
University of São Paulo, Brazil
renata@ime.usp.br

Abstract

In the AGM paradigm of belief change the background logic is taken to be a supra-classical logic satisfying compactness among other properties. Compactness requires that any conclusion drawn from a set of propositions X is implied by a finite subset of X. There are a number of interesting logics such as Computational Tree Logic (CTL, a temporal logic) which do not possess the compactness property, but are important from the belief change point of view. In this paper we explore AGM style belief contraction in non-compact logics as a starting point, with the expectation that the resulting account will facilitate development of corresponding accounts of belief revision. We show that, when the background logic does not satisfy compactness, as long as the language in question is closed under classical negation and disjunction, AGM style belief contraction functions (with appropriate adjustments) can be constructed. We provide such a constructive account of belief contraction that is characterised exactly by the eight AGM postulates of belief contraction. The primary difference between the classical AGM construction of belief contraction functions and the one presented here is that while the former employs remainders of the belief being removed, we use its complements.

Keywords: AGM Theory, Belief Contraction, Compactness.

1 Introduction

Knowledge management involves keeping an agent's body of knowledge up to date in light of new information deemed acceptable by an agent. There are two main approaches to model how a rational agent responds to new information. One, called *belief revision*, is based on the seminal work (Alchourrón, Gärdenfors, and Makinson 1985) further developed in many works such as (Gärdenfors 1988; Hansson 1999; Rott 2001). The other, called *belief update*, is founded on (Katsuno and Mendelzon 1991). The difference between these two approaches is often motivated by the (debatable) claim that belief revision is appropriate when the new information indicates correction (or addition) to the existing knowledge, whereas belief update is more appropriate when the new information indicates the domain in question has changed – in the former the "world" is considered static, and in the latter it is dynamic. In both the approaches, three forms of belief change operations are taken for granted – one for adding the new information without worrying about whether or not the resulted beliefs are jointly consistent; one for removing an existing belief; and one for incorporating the new information with the caveat that the resulting beliefs will be jointly consistent. The assumed behaviour of these operations (captured by corresponding *rationality postulates*) are in some form or other driven by the desideratum that the change involved should be minimised since no belief should be gained or lost without a reason. The latter two of the operations – for belief *removal* and for consistent accommodation of new information – are inter-translatable via the *Levi Identity* and the *Harper Identity*, and both of them involve some extra-logical choice mechanism such as *epistemic entrenchment* as a tie breaker. This extra-logical mechanism is employed to provide an explicit construction of the relevant belief change operator whose behaviour is "described" by the rationality postulates. Our work in this paper is confined to belief removal in a "static world", typically called *belief contraction*, as opposed to *belief erasure* which is the phrase used to name belief removal in a "dynamic world". We assume that the results obtained can shed light on, and can be easily extended to, belief revision.

Much of the research carried out in the field of belief change assumes that the beliefs of an agent are represented as sentences of a non-modal, propositional language, and the background logic used is the classical propositional logic. In particular, the background logic is represented as a Tarskian *consequence operator* satisfying, among other things, *Compactness*, that is any consequence of a set of sentences X is a consequence of a finite subset X' of X. The issue of belief change has also been examined by many using a non-classical logic as the background logic, such as (Restall and Slaney 1995) and (Wassermann 2011).

In these approaches knowledge of an agent is expressed in the object language associated with a background logic, whereas the belief change itself is carried out using some "extra-logical" operations such as contraction and revision. In contrast, efforts have been made to represent belief change dynamics in the object-language itself. A substantive approach toward this end is that of Dynamic Doxastic Logic, DDL in short, (Leitgeb and Segerberg 2007) which is a modal logic in which belief change operators are expressed in the object-language via modal operators. The semantics of DDL is given in terms of hypertheories in(Segerberg 1995). In this tradition, Cantwell (2000) proposed to represent contraction functions through hypertheories in DDL. Cantwell also considers extensions of DDL towards certain non-compact modal logics. A different approach to capture belief change via modal operators is presented in (Bonanno 2009) which uses temporal logics to represent the belief dynamics in a temporal framework and temporal operators to represent the interaction among an agent's beliefs over time. What all these works have in common is that they represent the belief dynamics of an agent in the object-language of a modal logic, while an agent's knowledge is still represented in classical logics. Unlike these works, we represent the knowledge of an agent in non-classical logics, while the belief dynamics operators are kept as extra-logical features in the spirit of the AGM approach.

In this paper we examine belief contraction without assuming any specific non-classical logic as the background logic, but only relaxing a requirement on the background logic – in particular, we no longer assume that the background logic satisfies compactness. For convenience we call such logics *Non-compact Logics*. What lends significance to this endeavour is that a number of interesting logics that play important roles in computing sciences including Artificial Intelligence do not satisfy compactness. These include logics such as temporal logics, e.g., Computation Tree Logic (CTL) and Linear Temporal Logic (LTL) (Clarke, Grumberg, and Peled 2001), that are widely used both in formal specification and verification of systems as well as in Planning. Branched temporal logics (Gabbay, Hodkinson, and Reynolds 1994) such as CTL allow explicit representation of and reasoning about incomplete information about temporal events.

There are a few natural paths one might take when exploring belief contraction with non-compact logics:

1. Suppose we construct a belief contraction function analogous to the way it is constructed in the AGM tradition, such as *Partial Meet Contraction, Transitively Relational Partial Meet Contraction* or *Safe Contraction*, assuming a non-compact logic as the background logic. Will the standard AGM postulates of contraction still be appropriate? If not, what alternative postulates will characterise such belief contraction functions?

2. Suppose we agree upon a set of belief contraction postulates that are arguably rational in the context of a background logic that is non-compact. Can we construct a belief contraction function that satisfy some, or all, of those postulates? If so, how?

3. Since AGM Contraction Postulates have survived the test of time, we may, for a start, assume that they are appropriate even if the background logic is non-compact. What other assumptions must we make about the background logic? Can we construct a belief contraction function that will satisfy them, and if so, how?

Our chosen path is the last one which has been paved to an extent by Flouris (2006) who showed that the existence of a contraction function that satisfies the first six of the AGM contraction postulates (called the *basic contraction postulates*)[1] does not require the background logic to satisfy all the AGM assumptions. He calls such background logics *AGM-Compliant*.

In (Ribeiro et al. 2013) the authors have identified some sufficient conditions for a logic to be AGM-Compliant and examined if the Partial Meet Contraction, which is characterised by the basic AGM postulates of contraction when the background logic satisfies the AGM assumptions such as compactness, is also so characterised in the presence of these sufficiency conditions alone. To the best of our knowledge, the existing efforts to apply belief revision in non-classical logics either presuppose compactness and rely on Partial Meet functions, or drop some of AGM postulates and propose a different class of contraction functions. In this paper, on the other hand, we identify the features of a background non-compact logic (and its associated language) that will guarantee its AGM-Compliance, and construct belief contraction functions that, apart from satisfying the basic AGM postulates of contraction as required by AGM-Compliance, satisfy the two *Supplementary AGM Postulates of contraction* as well. Specifically, we require that the language supporting such non-compact logics be closed under classical negation and disjunction.

We introduce in Section 2 some preliminary concepts and the notation used in this paper. Then in Section 3 we review the AGM theory of belief change, in particular belief contraction, discuss the notion of *AGM-Compliance*, and prove AGM-Compliance for the class of logics closed under classical negation and disjunction. In the subsequent two sections, Section 4 and 5 we provide constructions of two belief contraction functions that are respectively characterised by the basic AGM contraction postulates and the full set of AGM contraction postulates. Finally, in the concluding section, we provide a brief discussion of some interesting issues that will be taken up in our future work. Only occasionally we have sketched the outline of the proofs in the text of this paper; the proofs will be provided as part of a planned extended work.

2 Notation and Technical Background

Given a set A, the power set of A will be denoted as 2^A. We use the terms formula and sentence interchangeably. We will use upper case Roman letters (A, B, ...) to denote sets, and lower case Greek letter ($\alpha, \beta, \dots, \varphi, \psi, \dots$) will be used to denote formulas. We will reserve the upper case letter K for a special kind of sets called belief sets, and the Greek lower case letter γ to denote a kind of function called *selection function*. The upper case letter Γ is reserved to denote a collection of sets. Propositional symbols will be denoted by lower case Roman letters (p, q, r, ...). The letter M will denote a model. Moreover we will use the symbol $\subseteq$ for subset, whereas $\subset$ will denote proper subset. The maximal elements of a set A with respect to a binary relation $<$ is given by $max_<(A) = \{a \in A \mid \neg\exists b \in A, a < b\}$.

We consider a logic as a pair $\langle L, Cn\rangle$, where L is a language and Cn: $2^L \to 2^L$ is a logical consequence operator

[1]Please see Section 3 for the list of AGM contraction postulates.

that maps a set of formulas to the set of all its inferred formulas. For readability, for any formula φ, the set $Cn(\{\varphi\})$ will often simply be written as $Cn(\varphi)$. We will often pretend that the consequence operation Cn itself represents a logic when no confusion is imminent. We limit ourselves to logics that are Tarskian, that is, logics whose consequence operator satisfies the following three properties:

1. **(Monotonicity)**: $A \subseteq B$ iff $Cn(A) \subseteq Cn(B)$;
2. **(Idempotence)**: $Cn(A) = Cn(Cn(A))$; and
3. **(Inclusion)**: $A \subseteq Cn(A)$.

Apart from being Tarskian, the consequence operation is often granted some other properties in the AGM belief change literature, and they are often dubbed *AGM Assumptions*.

- **(deduction)**: $\varphi \in Cn(A \cup \{\psi\})$ iff $\psi \to \varphi \in Cn(A)$;
- **(supraclassicality)**: if φ is a logical consequence of A in classical propositional logic, then $\varphi \in Cn(A)$;
- **(compactness)**: if $\varphi \in Cn(A)$ then there is a finite subset A' of A such that $\varphi \in Cn(A')$.

We will say that a logic $\langle L, Cn\rangle$ is closed under classical negation iff the language L is closed under the negation operator, for each formula $\varphi \in L, Cn(\varphi) \cap Cn(\neg\varphi) = Cn(\emptyset)$, and $Cn(\{\varphi, \neg\varphi\}) = L$. In other words, the negation is interpreted classically. Analogously, the logic is closed under the disjunction if the language is closed under such a connective and it is also classically interpreted, that is, if $\varphi \in Cn(X)$ then $\varphi \vee \psi \in Cn(X)$, for every $\psi \in L$ and $X \subseteq L$. A set of formulas K closed under consequence relation, i.e. $K = Cn(K)$, is called a theory.

One of the inconveniences of dealing with non-compact logics is that it is not possible to define a finite proof system for these logics. We exploit the semantics of the logical consequence operator Cn as a remedial measure. Given a logic $\langle L, Cn\rangle$ and a set of structures $\mathcal{I}$, an interpretation or a model is an element of $\mathcal{I}$ that gives meaning to the formulas of L; $\mathcal{I}$ is called an *interpretation domain* of that logic, whereas each subset of $\mathcal{I}$ is called an *interpretation set*. A satisfaction relation $\models$ binds each pair in $\mathcal{I} \times L$ to a truth value T (true) or F (false), meaning that a formula is respectively true or false in that model. For instance, an interpretation domain for the Propositional Logic is the power set of the propositional symbols of the language.

In Tarskian logics, the consequence operator can be semantically defined as: a formula $\varphi \in Cn(X)$ iff every model that satisfies all formulas in X also satisfies φ. Let $\mathcal{I}$ be an interpretation domain of a logic $\langle L, Cn\rangle$, and M a model in $\mathcal{I}$. The set of all formulas of L satisfied by M is the theory $Th(M) = \{\varphi \in L \mid M \models \varphi\}$. Generalizing, given a set of models A, $Th(A) = \{\varphi \mid \forall M \in A, M \models \varphi\}$ is the theory of the formulas satisfied by all models in A. Moreover, given a theory $X \subseteq L$, the set of models that satisfy all formulas in X is $[\![X]\!] = \{M \in I \mid \forall \varphi \in X, M \models \varphi\}$. For simplicity, given a set of formulas X and a model M, we will write $M \models X$ to mean that M satisfies every formulas in X. Furthermore, we will call a theory T a *complete theory* if and only if, for every sentence φ, either $\varphi \in T$ or $\neg\varphi \in T$, and use T_L to denote all such complete theories.

Unlike in Propositional Logic, in non-classical logics, e.g. modal logics, two different models, say M and M', may satisfy exactly the same set of formulas, that is, $Th(M) = Th(M')$. We say that such models are semantically equivalent, and call a set closed under such equivalence *Completed Interpretation Sets* (CIS):

Definition 1. *A set of interpretations* A *is a Completed Interpretation Set (CIS) if and only if* A *is closed under semantic equivalence.*

This notion will prove useful to demonstrate AGM-Compliance of Tarskian Logics closed under classical negation and disjunction.

3 AGM Contraction and AGM-Compliance

All the beliefs of an agent as a whole is represented as a set of sentences K, called a *belief set* (or *theory*), that is assumed to be closed under logical consequence: $K = Cn(K)$. For notational convenience we take $K + \varphi$ to mean $Cn(K \cup Cn(\varphi))$, for any belief set K and sentence φ.[2] In the AGM paradigm of belief contraction, as well as other forms of AGM belief change (Alchourrón, Gärdenfors, and Makinson 1985), the background logic Cn is assumed to satisfy a set of properties called *AGM assumptions*, namely, it is Tarskian, supra-classical, compact, closed under all boolean connectives (conjunction, disjunction, implication and negation), and satisfies deduction. Let **K** be the collection of all belief sets. Then any function $f : \mathbf{K} \times L \to \mathbf{K}$ is a belief change operation. The full set of AGM rationality conditions for belief contraction are given below. For any theory K, sentences φ and ψ, and belief contraction function $-$:

$(K1)$	$K - \varphi = Cn(K - \varphi)$	(closure)
$(K2)$	$K - \varphi \subseteq K$	(inclusion)
$(K3)$	If $\varphi \notin K$, then $K - \varphi = K$	(vacuity)
$(K4)$	If $\varphi \notin Cn(\emptyset)$, then $\varphi \notin K - \varphi$	(success)
$(K5)$	$K \subseteq (K - \varphi) + \varphi$	(recovery)
$(K6)$	If $Cn(\varphi) = Cn(\psi)$, then $K - \varphi = K - \psi$	(extensionality)
$(K7)$	$(K - \varphi) \cap (K - \psi) \subseteq K - (\varphi \wedge \psi)$	
$(K8)$	If $\varphi \notin K - (\varphi \wedge \psi)$ then $K - (\varphi \wedge \psi) \subseteq K - \varphi$	

Postulates (K1)-(K6) are the basic AGM postulates for contraction, and the last two constitute the supplementary postulates. Discussion of and rationale behind these postulates can be found in (Gärdenfors 1988; Nayak 1994), among others. We will call any belief change operation that satisfies postulates (K1)-(K6) an *AGM rational belief contraction*. Any AGM rational belief contraction that also satisfies the supplementary postulates (K7)-(K8) will be said to be *fully AGM rational*.

The postulates (K1)-(K8) prescribe a good set of behaviours for a contraction function, but do not tell us where to find such a function. There are different available constructions of (fully) AGM-rational belief contraction functions. One of them uses as tool what is called a *remainder set*:

[2] This is actually called the *belief expansion* operator that is used to add beliefs without consideration of whether or not the result is consistent.

Definition 2. *A remainder set $K \perp \varphi$ of sentence φ with respect to a belief set K is the set of maximal subsets of K that do not entail φ. That is, $K \perp \varphi = \{K' \subseteq K \mid \varphi \notin Cn(K')$, and if $K' \subset K'' \subseteq K$, then $\varphi \in Cn(K'')\}$.*

The members of $K \perp \varphi$ are called the *remainders* of φ with respect to a belief set K, each of which may be seen as the result of removing a logically minimal amount of information from K in order to remove φ. Since generally $K \perp \varphi$ will have multiple members, a choice needs to be made as to which remainders to use for constructing a belief contraction function. This is done through the use of an extra-logical selection mechanism called a *selection function*.

Definition 3. *A selection function γ maps a remainder set $K \perp \varphi$ to a set of remainders in it such that*

$$\gamma(K \perp \varphi) = \begin{cases} \emptyset \neq \chi \subseteq K \perp \varphi & \text{if } K \perp \varphi \neq \emptyset \\ \{K\} & \textit{otherwise.} \end{cases}$$

Since no remainder in $K \perp \varphi$ entails φ, it follows that the beliefs common to the remainders chosen by γ will not entail φ. Indeed it can be shown that each remainder is a theory, as is the intersection of any set of remainders. Hence a contraction function $-_\gamma$, called the *partial meet contraction function*, can be constructed thus:

Definition 4. *Given a theory K, a sentence φ and a selection function γ, the operation $-_\gamma$ defined $K -_\gamma \varphi = \bigcap \gamma(K \perp \varphi)$ is called a partial meet contraction function.*

The selection function γ represents the epistemic choices made by an agent. However, it could be quite arbitrary, and hence we might wish to rationalise it by assuming that this choice behaviour reveals some underlying hidden preference of the agent. If we take that preference to be a transitive preference relation $\leq$ over members of the remainder set, and the selection function $\gamma_\leq$ always selects the best (modulo $\leq$) elements in $K \perp \varphi$, then the contraction function $-_{\gamma_\leq}$ is called a *transitively relational partial meet contraction function*. The following representation result is used to justify the AGM postulates for contraction as well as the construction of contraction functions via selection from remainder sets (Alchourrón, Gärdenfors, and Makinson 1985).

Theorem 1. *Let f be a belief change function.*

1. *The function f is an AGM rational contraction function if and only if it is a partial meet contraction function.*
2. *Furthermore, f is a fully AGM rational contraction function if and only if it is a transitively relational partial meet contraction function.*

Belief contraction in the AGM framework, as presented above, makes strong assumptions about the background logic Cn. However, it is not clear if, when we relax such assumptions, we can still construct remainders, let alone construct contraction functions that are (fully) AGM rational. A logic Cn is said to be *AGM-Compliant* if and only if, with Cn as the background logic, it is possible to define a contraction operation that satisfies all the six contraction AGM postulates. Flouris (2006) has shown that the existence of AGM-rational contraction functions does not depend on the background logic Cn satisfying the AGM assumptions. Rather it depends on what he calls its *decomposability* property. Decomposability in turn is defined in terms of *relative complement*.

Definition 5. *Let $\langle L, Cn\rangle$, and $A, K \in 2^L$ be such that $K = Cn(K)$ and $Cn(\emptyset) \subset Cn(A) \subset K$. The complement of A relative to K, denoted $K^-(A)$, is the collection of all sets K' such that $Cn(K') \subset Cn(K)$ and $Cn(K' \cup A) = Cn(K)$.*

To illustrate, let us take the theory $K = Cn(\{p, q\})$ over the set of atomic symbols $\{p, q\}$, and $A = \{p\}$. Clearly, $Cn(\emptyset) \subset Cn(A) \subset K$. If we let $K' = \{p \to q\} \subset Cn(K)$, then $Cn(K' \cup A) = K$. Thus, K' is in a complement of A relative to K. Other sets such as $\{q\}$ and $\{p \leftrightarrow q\}$ are also in $K^-(A)$.

Definition 6. *A logic $\langle L, Cn\rangle$ is decomposable iff $K^-(A) \neq \emptyset$, for every $A, K \in 2^L$ such that $K = Cn(K)$ and $Cn(\emptyset) \subset Cn(A) \subset K$.*

Decomposability has a straightforward relation with AGM-compliance:

Theorem 2. *(Flouris 2006) A logic $\langle L, Cn\rangle$ is decomposable iff it is AGM-Compliant.*

The rest of this section is devoted to show that the class of Tarskian logics closed under classical negation and disjunction are AGM-Compliant (Theorem 4). Towards this end we first need to establish some preliminary results. From now on, unless stated otherwise, any logic we mention is assumed to be a Tarskian Logic closed under classical negation and disjunction. We recall the functions $\dot{[\![]\!]}$ and $\dot{Th()}$ introduced in Section 2. The proof of the following result is straightforward.

Observation 1. *Let A be an interpretation set. Then, $Th(A) = \bigcap_{M \in A} Th(M)$.*

A related immediate result is that any theory obtained from a model is a complete theory due to the presence of the classical negation.

Observation 2. *The theory $Th(M)$ corresponding to a model M is a complete theory.*

As in the classical case, we show that the more models a CIS has, the smaller is its theory.

Proposition 1. *If $K' \subset K$, then $[\![K]\!] \subset [\![K']\!]$, for all theories K and K'.*

Proposition 2 (whose proof is analogous to the classical case) and Corollary 1 are of immediate assistance for the subsequent Proposition 3.

Proposition 2. *Let Γ be a set of distinct complete theories and B a consistent complete theory which is not in Γ. Then, $\left(B \cap \bigcap \Gamma\right) \subset \bigcap \Gamma$.*

Corollary 1. *Let A and B be two CISs. If $M \notin A$, then*

(D1) *$Th(A \cup \{M\}) \subset Th(A)$ and $Th(A \cup \{M\}) \subset Th(M)$; and*

(D2) *if $M \in B$, then $Th(A \cup B) \subset Th(A)$ and $Th(A \cup B) \subset Th(B)$.*

The following proposition assures that no two different theories have the same CIS. An immediate result of this is that there is a bijection between the realm of theories and the class of CISs (see Theorem 3 below).

Proposition 3. *An interpretation set A is a CIS iff $A = [\![K]\!]$ for some theory K.*

We are now ready to define a function that maps each theory to a CIS. An important aspect of such a function is that it maps each theory K exactly to the only completed interpretation set from which K follows. We shall show that this function is a bijection which will be used to demonstrate that logics closed under classical negation are AGM-Compliant.

Theorem 3. *Let $\mathcal{T}$ be the the set of all theories from a logic $\langle L, Cn\rangle$, and $\mathcal{I}_C$ the class of all complete interpretation sets of that logic. Then the function $\tau : \mathcal{T} \to \mathcal{I}_C$, defined $\tau(K) = [\![K]\!]$, is a bijection.*

As τ is a bijection, we have $\tau^{-1}(Y) = Th(Y)$, where Y is a complete interpretation set. Proposition 4 and 5 and Corollary 2, whose proofs are quite straightforward, are useful to prove Theorem 4.

Proposition 4. *Let A and B be two CISs. If $A \subset B$ then $Th(B) \subset Th(A)$.*

Proposition 5. *Let K and K' be two theories. Then, $Th([\![K]\!] \cap [\![K']\!]) = Cn(K \cup K')$.*

Corollary 2. *Let A and B be two sets of formulas, then $Cn(A \cup B) = Cn(A \cup Cn(B))$.*

Now we have all the ingredients to show the last result in this section:

Theorem 4. *Every Tarskian logic closed under classical negation and disjunction is AGM-Compliant.*

4 Construction: AGM Basic Rationality

In the previous section we showed that Tarskian Logics closed under classical negation and disjunction are AGM-Compliant. Question arising: which class of contraction functions is characterized by the AGM postulates in such logics? Unfortunately, we can no longer rely on partial meet contraction functions for this purpose since, as pointed out in (Ribeiro et al. 2013), partial meet functions do not satisfy the AGM postulates in some non-compact logics. We need to devise a new class of functions.

In this section, we construct a new contraction function which satisfies the postulates $(K1)$ to $(K6)$. Recall that T_L denotes the set of all consistent complete theories of a logic $\langle L, Cn\rangle$, and an agent's belief set K is closed under the consequence operation Cn. In the AGM approach the partial meet contraction depends on remainder sets whose existence is guaranteed by the compactness property of the background logic. But since we do not have the compactness property to fall back upon, the contraction function we define will depend on a selection of complete consistent theories which will be intersected with the belief set K. Accordingly we assume a Choice Function $\delta : L \to 2^{T_L}$ that maps each formula φ of L to a set of complete theories $\delta(\varphi)$, subject to the following conditions:

1. $\delta(\varphi) \neq \emptyset$;
2. if $\varphi \notin Cn(\emptyset)$, then $\delta(\varphi) \subseteq \{S \in T_L \mid \varphi \notin S\}$;
3. for any formulas φ and ψ, if $\varphi \equiv \psi$ then $\delta(\varphi) = \delta(\psi)$.

The purpose of the choice function δ is to pick the best complete theories that do not entail a non-tautological formula φ (condition 2). Condition 1 dictates that for any formula, at least one complete theory has to be selected. The last condition assures that the choice mechanism used is not syntax-sensitive. The main difference between this choice function and the selection function used in the construction of the partial meet contraction is that while the latter picks remainders for a given theory (belief set) and formula pair, the former picks complete theories for a given formula, and the role of the theory is postponed to the construction stage of the contracted set. In this sense, our approach better corresponds to the semantic counterpart of the AGM partial meet contraction function: $[\![K - \varphi]\!] = [\![K]\!] \cup \delta'([\![\neg\varphi]\!])$ where δ' is an appropriate choice function that picks models from an input set.

Definition 7. *Let K be a theory, φ a formula, and δ a choice function. An operation $-_\delta$ is an Exhaustive Contraction Function (ECF) iff*

1. *$K -_\delta \varphi = K \cap \bigcap \delta(\varphi)$, if* both:
   (i) $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$ and
   (ii) *either $\neg\varphi \notin Cn(\emptyset)$ or $\bot \in K$;*
2. *$K -_\delta \varphi = K$, otherwise.*

Let us look at the constraints imposed on the ECF by the above definition. A formula φ may only be retracted from a belief set K when: (1) K is not simply a set of tautological formulas, (2) φ itself is not a tautology and (3) φ is in K. These constrains are jointly expressed as $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$. Besides, if K is consistent, φ also needs to be consistent (otherwise φ is not a belief and hence its removal is vacuous).

For illustration purposes, let us contract a formula from a theory expressed in a non-compact logic, namely, the Linear Temporal Logic (Clarke, Grumberg, and Peled 2001). For simplicity, we will consider only two temporal operators of that logic: G and X. The former means *Globally (always) in the future*, and X means *in the "neXt" time instant*. We will keep the disjunction and negation of that logic which are interpreted classically. For our purpose, it will suffice to note two properties regarding these two operators. First, $Cn(\{Gp, p, Xp, X^2p, \ldots, X^np, \ldots\} \subseteq Cn(Gp)$. In other words, Gp implies that p is true in the current time instant and in all next future instants. As the disjunction and negation are interpreted classically, we note that formulas such as $\neg Xp \vee Gp$ and $\neg\neg p$ also belong to $Cn(Gp)$. The second point we need to note is that $Cn(\{p, Xp, \ldots, X^np, \ldots\}) = Cn(Gp)$. For more details of this logic and its semantics, readers may consult (Clarke, Grumberg, and Peled 2001).

Example 1. *Consider the theory $K = Cn(Gp)$ and we wish to contract Xp from it. Let our choice function be δ_1 where:*

1. *$\delta_1(Xp) = Cn(\{p, \neg Xp, Xp \to Gp, \neg X^2p, \ldots\})$;*

2. *Else, if* $\delta_1(\psi) = T_L$ *then* $Cn(\psi) = Cn(\emptyset)$*;*
3. *Else,* $\delta_1(\psi) = \{S \in T_L \mid \psi \notin S\}$*.*

If ψ *is a tautology, we just let* $\delta_1(\psi) = T_L$*. The first constraint above concerns the complete theories chosen to contract the formula* Xp*. As this is the only formula we are interested to retract, for all other non-tautological formulas* ψ *we let* δ_1 *choose all the complete theories that do not imply* ψ *(third constraint). So,* $K -_{\delta_1} Xp = K \cap \delta_1(Xp) = Cn(\{p, Xp \to Gp, X^2p \to Gp, \dots\})$*.*

It is easy to notice that $-_{\delta_1}$ *satisfies* $(K1)$ *to* $(K4)$*. The extensionality postulate* $(K6)$ *follows from condition 3 of the definition of a choice function. For Recovery (K5), it suffices to note that* $Xp \to Gp$ *is in both* $\delta_1(Xp)$ *and* K*, whereby,* $(K -_{\delta_1} Xp) + Xp = K$*.*

Having seen how the contraction function ECF works in practice, let us see how well behaved it is.

Proposition 6. *If a non-tautological formula* φ *belongs to a theory* K *then* $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$*, and either* φ *is consistent or* K *is inconsistent.*

Lemma 1. $[\![Cn(K \cup K')]\!] = [\![Cn(K)]\!] \cap [\![Cn(K')]\!]$*, for all theories* K *and* K'*.*

Lemma 2. $Cn(K \cap K') = Th\big([\![K]\!] \cup [\![K']\!]\big)$*, for all theories* K *and* K'*.*

An immediate consequence of Lemma 2 is that the intersection of a set of theories corresponds to the union of the CIS of the theories being intersected:

Corollary 3. $[\![\bigcap X]\!] = \bigcup_{Y \in X} [\![Y]\!]$*, for all sets of theories* X*.*

This leads us to show that ECF is quite well behaved – it satisfies the six basic AGM contraction postulates:

Theorem 5. *Every ECF satisfies* $(K1)$ *to* $(K6)$*.*

The above result, Theorem 5, establishes half of our first representation result. The following two results, Proposition 7 and Lemma 3, are useful to prove the other half of this representation result (Theorem 6).

Proposition 7. *Let* $-$ *be a function that satisfies* $(K1)$ *to* $(K6)$*. Then,* $K - \varphi \neq K$ *iff* $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$*, and either* $\neg\varphi \notin Cn(\emptyset)$ *or* $\bot \in K$*.*

Lemma 3. *For any theory* K *there is a set of complete theories* X *such that* $K = \bigcap X$*.*

Theorem 6. *If a contraction function* $-$ *satisfies the six basic contraction postulates, then there exists ECF* $\dot{-}$ *such that for any theory* K *and formula* φ*,* $K - \varphi = K \dot{-} \varphi$*.*

Theorems 5 and 6 together show that ECF is indeed very well behaved, and constitute our first representation result.

5 Construction: AGM Full Rationality

In this section we introduce a class of contraction functions that satisfy all eight AGM contraction postulates, the *Blade Contraction*. The main idea is to constrain the class of ECFs in a way that the choice function employed by each such ECF can be represented by a binary relation which respects two conditions, to be named *Maximal Cut* and *Mirroring*.

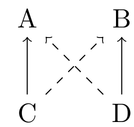

Figure 1: Mirroring. A is preferred to C, hence to D.

Given a theory K and a formula φ, a way to contract K by φ is to intersect K with some consistent complete theories that do not imply φ. In a logic closed under classical negation, these consistent complete theories entail $\neg\varphi$. A consistent complete theory that does not imply a formula φ will be called a complement of φ, and the class of all complements of φ is given by the set $\omega(\varphi) = \{S \in T_L \mid \varphi \notin S\}$. We highlight that, due to the classical negation, $\omega(\varphi \wedge \psi) = \omega(\varphi) \cup \omega(\psi)$. Also note that, since the logic is closed under classical negation, $\omega(\varphi)$ is empty if and only if φ is a tautology,

The basic idea is that the agent's choice function δ does not behave arbitrarily; its behaviour is "rationalised" by the agent's preference: the agent's preference is revealed by its choice behaviour. We represent this preference as a binary relation $<$ over complements of φ. The two constraints we impose on $<$ are:

(Maximal Cut) for every non-tautological formula $\varphi \in L$, $\omega(\varphi)$ has a maximal element w.r.t. $<$.

(Mirroring) if $S_1 \not< S_2$ and $S_2 \not< S_1$; then for any $S' \in T_L$, if $S_1 < S'$ then $S_2 < S'$

The first condition on $<$, *maximal cut*, is similar to the *Limit Assumption* of (Lewis 1976) and the *Finite Stopperedness* of (Gärdenfors and Makinson 1994). It guarantees that for every formula φ, an agent chooses at least one complement theory of φ. The purpose of the Maximal cut is to ensure that every formula to be dropped will be successfully relinquished, that is, the theory being contracted will be intersected by complements of φ.

The second condition, *Mirroring*, is similar to the modular relation defined in (Meyer, Labuschagne, and Heidema 2000) which was based on *modular* partial orders of (Ginsberg 1986) and (Lehmann and Magidor 1992). Though the concept of modular relation is confined to be a partial order, we impose no such restriction. The intuition behind mirroring is that if an agent has no preference between two theories A and B, then those that are preferable to A should also be preferable to B and vice versa. For instance, when dropping a formula φ an agent may choose among the four complements of φ: A, B, C and D. It prefers A to C and B to D, that is, its preference relation is $\{(C, A), (D, B)\}$ which is depicted in Figure 1 by solid arrows. So, it will choose both A and B to contract φ. However, there is no preference between C and D. According to mirroring, all theories that are preferable to C are also preferable to D (and vice versa). Thus, the pairs (C, B) and (D, A) also need to be present in the relation (depicted by dashed arrows in Figure 1).

The rest of this section will proceed as follow. First, we introduce a relational choice function whose elements are picked according to a binary relation that satisfies maximal cut and show that contraction functions defined via such relational choice functions satisfy the postulate $K7$. Then we show that if the binary relation in question also satisfies mirroring, then the contraction functions defined via it satisfies both the postulates $K7$ and $K8$.

From now on, unless otherwise specified, we consider only relations that satisfy maximal cut. We call a binary relation that satisfis maximal cut *contra-headed*. First, we explain the idea of relational choice functions. Let δ be a choice function and φ a formula. Moreover, let $<$ be a binary relation over the class of all complete theories. Recall that, given a set of complete theories A, $max_<(A)$ is the set of all maximal theories in A w.r.t. $<$. The choice function $\delta(\varphi)$ is restricted to pick only from the complement theories of φ, that is, $\delta(\varphi) \subseteq \omega(\varphi)$. Thus, we will say that δ is relational if there is a binary relation $<$ over the set of all complete theories such that the elements picked by $\delta(\varphi)$ coincide with the maximal elements of $\omega(\varphi)$ ordered by $<$, that is, $\delta(\varphi) = max_<(\omega(\varphi))$. We note that $\delta(\varphi)$ has to pick some complements of φ, for every non-tautological formula φ. In the case that δ is relational, it follows that there is a maximal element in $\omega(\varphi)$ w.r.t. the corresponding binary relation $<$; that is, $<$ respects the maximal cut property, and so is contra-headed. A relational choice function whose corresponding binary relation is contra-headed will be called an annulment. We will use μ instead of δ to denote annulment functions.

Definition 8. *Let $<$ be a contra-headed relation over T_L. An annulment is a function $\mu_< : L \to 2^{T_L}$ such that*

$$\mu_<(\varphi) = \begin{cases} max_<\big(\omega(\varphi)\big) & \text{if } Cn(\varphi) \neq Cn(\emptyset); \\ \text{some } \emptyset \neq X \subseteq T_L & \text{otherwise} \end{cases}$$

Note that, for the tautologies, an annulment is free to chose any element, since tautologies do not have any complement. As expected an annulment is indeed a choice function.

Lemma 4. *Every annulment is a choice function.*

Now we proceed to define a new contraction function that we call the *Blade Contraction Function (BCF)*. The idea is similar to an ECF, we simply restrict the choice function to be an annulment function. This will ensure that every BCF satisfies the postulate $(K7)$.

Definition 9. *Let $\mu_<$ be an annulment w.r.t. a contra-headed relation $<$. A Blade Contraction Function (BCF) $\dot{-}_<$ is constructed from $<$ as follows:*

1. $K \dot{-}_< \varphi = K \cap \bigcap \mu_<(\varphi)$, *if* both
   (i) $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$, and
   (ii) *either* $\neg\varphi \notin Cn(\emptyset)$ *or* $\bot \in K$;
2. $K \dot{-}_< \varphi = K$, *otherwise.*

Notice that the definitions of BCF and ECF are quite similar, the only difference is that a BCF involves an annulment, and an ECF involves a choice function. It trivially follows from Definition 9 and Lemma 4 that:

Corollary 4. *Every Blade Contraction Function is an Exhaustive Contraction Function.*

$$\begin{tikzcd}
A = Cn(\neg p, q) & B = Cn(\neg p, \neg q) \arrow[l] \\
C = Cn(p, q) \arrow[u] \arrow[ur] \arrow[r] & D = Cn(p, \neg q) \arrow[ul, dashed]
\end{tikzcd}$$

Figure 2: A depiction of Example2.

Example 2. *Let $K = Cn(p, q)$ be a propositional theory, and $<$ the binary relation depicted by solid arrows in Figure 2. Moreover, let $\mu_<$ be the respective annulment function. The complements of p, q and $p \wedge q$ are respectively $\omega(p) = \{A, B\}$, $\omega(q) = \{B, D\}$ and $\omega(p \wedge q) = \{A, B, D\}$. According to $<$ which is contra-headed: $\mu_<(p) = \{A\}$, $\mu_<(q) = \{B, D\}$ and to $\mu_<(p \wedge q) = \{A, D\}$. Thus we have:*

$$\begin{array}{lll} K \dot{-}_< p & = K \cap \bigcap \mu(p) & = Cn(q) \\ K \dot{-}_< q & = K \cap \bigcap \mu_<(q) & = Cn(\neg q \vee p) \\ K \dot{-}_< p \wedge q & = K \cap \bigcap \mu_<(p \wedge q) & = Cn(p \vee q) \end{array}$$

A blade contraction behaves similarly to an ECF, the difference is that the choice function picks the complements of a formula φ according to a binary relation. The binary relation guarantees that a BCF satisfies postulate K7.

Theorem 7. *Every BCF satisfy $(K7)$.*

Though blade contractions satisfy postulate K7, maximal cut alone is not enough to satisfy postulate $K8$. This can be seen on the blade function of the Example 2. First, from the example we have that $K \dot{-}_< q = Cn(\neg q \vee p)$ and $K \dot{-}_< p \wedge q = Cn(p \vee q)$. Now, notice that $q \notin K \dot{-}_< p \wedge q$ and that $Cn(p \vee q) \not\subseteq Cn(\neg q \vee p)$ which implies that $K \dot{-}_< p \wedge q \not\subseteq K \dot{-}_< q$. This means that $\dot{-}_<$ does not satisfy postulate K8.

To capture $K8$ it will suffice to consider contra-headed relations that also satisfy mirroring. Lemma 8 will help explain how postulates K7 and K8 dictate the way complements of a conjunction $\varphi \wedge \psi$ are chosen in an ECF. To understand why mirroring is related to postulate K8, we will need the condition C2 of that lemma.

According to condition C2 of Lemma 8, if some complement A of a formula φ is chosen to contract a conjunction $\varphi \wedge \psi$, then all theories chosen to contract φ must also be picked to contract the conjunction $\varphi \wedge \psi$. Another way of reading this is: if a complement K' of φ is not chosen in the contraction of φ, the complements of φ that were not chosen for the contraction of this formula are not chosen in the contraction of $\varphi \wedge \psi$. In terms of a relation, this is equivalent to say that if a complement is not maximal in $\omega(\varphi)$ then it is also not maximal in $\omega(\varphi \wedge \psi)$.

To see how mirroring guarantees condition C2, let us first look at the contraction function of Example 2. Let $\mu_<$ be the annulment function of that example. We see that though the theory $Cn(\neg p \wedge q)$ belongs both to $\mu_<(p)$ and $\mu_<(p \wedge q)$, the theory $Cn(p \wedge \neg q)$ is not in $\mu_<(p \wedge q)$. So, let us see what modifications are necessary in order to turn $<$ into a relation that satisfies C2. First, we notice that the only reason why $<$ does not satisfy mirroring is because though $Cn(p, \neg q)$ and $Cn(\neg p \wedge \neg q)$ are not comparable, $Cn(p, \neg q)$ is not less preferable than $Cn(\neg p, q)$. Note that this is the same reason why $\mu_<$ does not satisfy C2: though $Cn(p, \neg q)$ and $Cn(\neg p \wedge \neg q)$ are chosen to contract q, only $Cn(\neg p \wedge \neg q)$ was chosen to contract $p \wedge q$. As $Cn(p \wedge \neg q)$ is already less preferable than $Cn(\neg p \wedge q)$, the theory $Cn(p \wedge \neg q)$ cannot be chosen to contract $p \wedge q$. So, an option to make it satisfy C2 is to make $Cn(p \wedge \neg q)$ also less preferable then $Cn(\neg p \wedge q)$. This new relation $<_1$ is depicted in Figure 2 as arrows. So we have that,

$$\begin{array}{lll} K \dot{-}_{<1} p & = K \cap Cn(\neg p \wedge q) & = Cn(q) \\ K \dot{-}_{<1} q & = K \cap \bigcap \mu_<(q) & = Cn(\neg q \vee p) \\ K \dot{-}_{<1} p \wedge q & = K \cap Cn(\neg p \wedge q) & = Cn(q) \end{array}$$

It is easy to see that the $\dot{-}_{<1}$ satisfies C2 and also K8.

Now we proceed to prove that every Blade Contraction function whose annulment respects mirroring satisfies postulate $(K8)$. This claim is proven as Theorem 8 which brings us to Lemma 5 and Lemma 6.

Lemma 5. *Let A and B be two sets ordered by a contra-headed relation $<$ that satisfies mirroring. If some maximal element of A is also a maximal element of $A \cup B$, then $max_<(A) \subseteq max_<(A \cup B)$.*

Lemma 6. *Let $\mu_<$ be an annulment. Every complete theory S in $\mu_<(\varphi \wedge \psi)$ is such that either $S \in \mu_<(\varphi)$ or $S \in \mu_<(\psi)$.*

Theorem 8. *Let $\dot{-}$ be a BCF w.r.t. an annulment $\mu_<$. If $<$ satisfies mirroring, then $\dot{-}$ satisfies $(K8)$.*

So, from Theorem 7 and Theorem 8, we conclude that every BCF satisfies (K7) and (K8).

The next step is to prove that every contraction function that satisfies all eight AGM contraction postulates is a BCF. First we recall that every function that satisfies the six basic AGM postulates is an ECF (Theorem 6), whereby every function that satisfy the eight AGM postulates is also an ECF. Now, it will be sufficient to show that for every ECF $\dot{-}_\delta$ that satisfies $(K7)$ and $(K8)$, the choice function δ is in fact an annulment, that is, $\delta(\varphi) = max_<(\omega(\varphi))$ for some binary relation $<$. This is a bit tricky. Note that we are considering only contra-headed binary relations. From here onwards, unless stated otherwise, any choice function mentioned is taken to concern an ECF that satisfies all eight AGM postulates. We will need to show how to construct a contra-headed relation for a choice function.

We start from a basic idea. Given a choice function δ, and a formula φ we will need a relation $<_\varphi$ such that the maximal elements of $\omega(\varphi)$ ordered via $<_\varphi$ coincide with $\delta(\varphi)$, that is, $\delta(\varphi) = max_{<\varphi}(\omega(\varphi))$. This is easy, and it suffices to make $<_\varphi = \{(A, B) \in \omega(\varphi) \times \omega(\varphi) \mid A \notin \delta(\varphi) \text{ and } B \in \delta(\varphi)\}$.

The tricky bit is to construct the general relation for the whole class of complete theories, because simply compiling every $<_\varphi$ together as a relation $< = \bigcup_{\varphi \in L} <_\varphi$ will not satisfy mirroring. To notice that, let us take an example. Let us take a choice function δ, such that for a formula φ, it chooses a complete theory A among the set of complete theories A, B and C, whereas for a formula ψ it chooses the theory D over

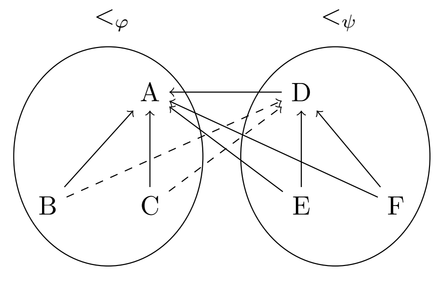

Figure 3: Three relations: $<_\varphi$, $<_\psi$ and $<_{\varphi \wedge \psi}$.

the set of complete theories D, E and F. Moreover, for the conjunction $\varphi \wedge \psi$ it chooses only D. Summarizing, we have

$$\begin{aligned} \delta(\varphi) &= \{A\} \subset \omega(\varphi) = \{A, B, C\} \\ \delta(\psi) &= \{D\} \subset \omega(\psi) = \{D, E, F\} \\ \delta(\varphi \wedge \psi) &= \{D\} \subset \omega(\varphi \wedge \psi) = \{A, B, C, D, E, F\}. \end{aligned}$$

The relations $<_\varphi$, $<_\psi$ and $<_{\varphi \wedge \psi}$ are illustrated in Figure 3. The relation $<_{\varphi \wedge \psi}$ consists of all solid arrows in the figure, whereas the relations $<_\varphi$, $<_\psi$ are depicted by the arrows inside its respective spheres. It is easy to see that $<_{\varphi \wedge \psi}$ is a contra-headed relation for δ. This relation, however, does not satisfy mirroring: observe that $B \not< E$ and $E \not< B$, however $E < D$ and $B \not< D$ violating mirroring. Hence more is needed. This example suggests that if we relate all the non maximal elements of $<_\varphi$ to the maximal elements of $<_\psi$, that is, add the pairs (B, D) and (C, D), the resulting relation does satisfy mirroring. This operation that relates the non-maximal elements of a relation $<_\varphi$ to the maximal ones of another relation $<_\psi$ will be called a triangulation.

Definition 10 (Triangulation)**.** *Let δ be a choice function, φ and ψ be two formulas. The triangulation of φ and ψ w.r.t. δ is the relation $\nabla(\varphi, \psi) = \{(A, B) \in \omega(\varphi) \times \omega(\psi) \mid A \in \omega(\varphi) \setminus (\delta(\varphi) \cup \omega(\psi))$, and $B \in \delta(\psi)\}$.*

For instance, the triangulation for φ and ψ, in the example of Figure 3, is $\delta(\varphi, \psi) = \{(B, D), (C, D)\}$ which is depicted by dashed arrows in that same figure.

Our intention is to show that every ECF that satisfies (K7) and (K8) is a blade contraction. Hence we need to show that there exists a binary relation $<$ satisfying mirroring such that $\delta(\varphi) = max_<(\omega(\varphi))$, that is, δ is an annulment. Employing triangulation, we can construct such a relation that we will call the *shadow* of δ. Theorem 9 shows that every choice function from an ECF for (K7) and (K8) is an annulment, and Theorem 10 shows that the shadow of δ satisfies mirroring.

Definition 11. *Let δ be a choice function. Then the shadow of δ is a relation $< \subseteq T_L \times T_L$ such that (A, B) is in $<$ iff, for some formulas φ and ψ:*

$$\text{either } (A, B) \in \Big(\bigcup_{\varphi \in L} <_\varphi \Big)$$
$$\text{or } (A, B) \in \nabla(\varphi, \psi) \cup \nabla(\psi, \varphi).$$

Lemmas 7, 8, 9 and Corollary 5 below are of assistance to prove Theorem 9. Lemma 8 correlates how a choice function δ behaves in the presence of the two supplementary postulates. Given two formulas φ and ψ, according to $(K7)$, in order to contract the conjunction of the formulas φ and ψ, the function δ can choose only from the complete theories picked to individually contract φ and ψ, that is, from $\delta(\varphi) \cup \delta(\psi)$. For the postulate $(K8)$, if some complete theory picked to contract $\varphi \wedge \psi$ was also picked to contract φ, then all complete theories chosen to contract φ must also be picked to contract the conjunction $\varphi \wedge \psi$. One immediate consequence of these two conditions, as shown in Corollary 5, is that if a complete theory is chosen either to contract φ or to contract ψ, then the set of complete theories picked to contract the conjunction $\varphi \wedge \psi$ is precisely the complete theories chosen to contract both φ together with those picked to contract ψ, that is, the set $\delta(\varphi) \cup \delta(\psi)$. Lemma 9 exhibits a similar property of Corollary 5 but in terms of the shadow relation of δ. It states that the maximal elements of $<_\varphi$ and $<_\psi$ are also maximal elements of $<_{\varphi\wedge\psi}$, when at least one theory from the intersection of $\delta(\varphi) \cap \delta(\psi)$ is chosen. This is because the elements chosen by each $\delta(\varphi)$ are preserved as maximal elements of the shadow of δ.

Lemma 7. *Let K be a theory, and A and B be two sets of complete theories. If $K \cap \bigcap A \subseteq K \cap \bigcap B$, then $B \subseteq A$.*

Lemma 8. *Let $\dot{-}_\delta$ be an ECF, then*

(C1) If $\dot{-}_\delta$ satisfies (K7), then $\delta(\varphi \wedge \psi) \subseteq \delta(\varphi) \cup \delta(\psi)$;

(C2) If $\dot{-}_\delta$ satisfies (K8), then $\delta(\varphi) \subseteq \delta(\varphi \wedge \psi)$, given there is some theory S of $\delta(\varphi)$ in $\delta(\varphi \wedge \psi)$.

Corollary 5. *Let $\dot{-}_\delta$ be an ECF that satisfies $(K7)$ and $(K8)$. If, $\delta(\varphi) \cap \delta(\psi) \neq \emptyset$ then $\delta(\varphi \wedge \psi) = \delta(\varphi) \cup \delta(\psi)$.*

Proposition 8. *Let $\dot{-}_\delta$ be an ECF, and φ and ψ two formulas. If $C \notin \delta(\varphi)$, $C \in \delta(\psi)$ and $C \in \delta(\varphi \wedge \psi)$, then $\dot{-}_\delta$ does not satisfy (K7) or (K8).*

Lemma 9. *Given an ECF $\dot{-}_\delta$ satisfying (K7) and (K8), and formulas φ and ψ, if $\delta(\varphi) \cap \delta(\psi) \neq \emptyset$ then $<_\varphi \cup <_\psi \subseteq <_{\varphi\wedge\psi}$.*

Theorem 9. *If an ECF $\dot{-}_\delta$ satisfies (K7) and (K8), then the choice function δ is an annulment.*

So far, we have shown that every function that satisfies all the eight postulates is an ECF whose choice function is an annulment, that is, a Blade Contraction function. The only other thing we need to finish off the representation result is to show that the shadow of such a choice function respects mirroring.

Theorem 10. *Let $\dot{-}_\delta$ be an ECF that satisfies (K7) and (K8). Then, the shadow $<$ of δ satisfies mirroring.*

From Theorem 7 and Theorem 8, together with Theorem 10, the desired representation theorem between the eight AGM contraction postulates and the Blade Contraction function class easily follows:

Theorem 11. *Every contraction function that satisfies all the eight postulates is equivalent to a BCF.*

6 Conclusion and Future Works

We argued that it is important to develop accounts of belief contraction when the background logic does not guarantee the compactness property, with the expectation that it would facilitate a corresponding account of belief revision. While it was known that satisfaction of the six basic AGM contraction postulates did not necessitate the background logic to satisfy the AGM assumptions such as compactness, not much research was done as to the necessary features of the background logic. We showed in Section 3 that as long as the background logic is Tarskian, and the associated language is closed under classical negation and disjunction, the existence of contraction functions that satisfy the basic AGM contraction postulates is guaranteed. Then we showed in the following section how such a contraction function can be constructed. This construction is analogous to the partial meet contraction in the sense that it uses a selection function over some complete theories that may be taken to be analogues of the remainders used in the AGM tradition. Finally in Section 5 we showed how the selection function in question can be rationalised so that the resulting contraction function will be fully AGM rational.

There are interesting related issues worth exploring:

1. In the classical framework, when the contraction function is fully AGM rational, the corresponding revision function defined via the Levi Identity is also appropriately AGM rational. However we have not examined if such nice behaviour of the revision function is assured in our case when the logic is more general.

2. Belief change has a close connection with non-monotonic logics such as *cumulative logics* developed and explored in (Gärdenfors and Makinson 1994). In such accounts, the non-monotonicity of the cumulative logic is taken to piggy-back on a classical logic satisfying compactness among others. For instance, in the *Right Weakening* property, namely: if $\vdash \beta \rightarrow \gamma$ and $\alpha \mid\sim \beta$ then $\alpha \mid\sim \gamma$, the logic represented by $\vdash$ is taken to be classical. Our work will facilitate exploration of such non-monotonic logics when $\vdash$ is more general and non-compact.

3. In order to rationalise the choice function via a preference relation so that the contraction function based on it will be fully AGM rational, we imposed two properties on that relation that are similar to the Limit assumption of (Lewis 1976) and the modularity imposed by (Meyer, Labuschagne, and Heidema 2000). It will be worth while to explore the connection between such works and ours and discover deeper interconnections.

4. Finally, we believe research into belief update when the background logic is more general could be productive.

These are issues we plan to take up in our future work.

Acknowledgements

The first author is supported by a scholarship from the CAPES Foundation within the Brazilian Government's Ministry of Education, as well as a scholarship from Macquarie University under a Cotutelle agreement. This research has also been partially supported by the Australian Research Council (ARC), Discovery Project: DP150104133.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic* 50(2):510–530.

Bonanno, G. 2009. Belief revision in a temporal framework. *New Perspectives on Games and Interaction* 4:45–80.

Cantwell, J. 2000. Logics of belief change without linearity. *The Journal of Symbolic Logic* 65(4):1556–1575.

Clarke, E. M.; Grumberg, O.; and Peled, D. A. 2001. *Model checking*. MIT Press.

Flouris, G. 2006. *On Belief Change and Ontology Evolution*. Ph.D. Dissertation, University of Crete.

Gabbay, D. M.; Hodkinson, I.; and Reynolds, M. 1994. *Temporal Logic (Vol. 1): Mathematical Foundations and Computational Aspects*. New York, NY, USA: Oxford University Press, Inc.

Gärdenfors, P., and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artif. Intell.* 65(2):197–245.

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

Ginsberg, M. L. 1986. Counterfactuals. *Artif. Intell.* 30(1):35–79.

Hansson, S. O. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*, volume 11. Kluwer Academic.

Katsuno, H., and Mendelzon, A. O. 1991. On the difference between updating a knowledge base and revising it. In *KR*, 387–394.

Lehmann, D. J., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artif. Intell.* 55(1):1–60.

Leitgeb, H., and Segerberg, K. 2007. Dynamic doxastic logic: why, how, and where to? *Synthese* 155(2):167–190.

Lewis, D. 1976. *Counterfactuals*. Oxford.

Meyer, T. A.; Labuschagne, W. A.; and Heidema, J. 2000. Refined epistemic entrenchment. *Journal of Logic, Language and Information* 9(2):237–259.

Nayak, A. C. 1994. Foundational belief change. *J. Philosophical Logic* 23(5):495–533.

Restall, G., and Slaney, J. 1995. Realistic belief revision. In *Proceedings of the Second World Conference on Foundations of Artificial Intelligence (WOCFAI-95)*, 367–378.

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

Rott, H. 2001. *Change, Choice and Inference*. Oxford.

Segerberg, K. 1995. Belief revision from the point of view of doxastic logic. *Logic Journal of the IGPL* 3(4):535–553.

Wassermann, R. 2011. On AGM for non-classical logics. *Journal of Philosophical Logic* 40(2):271–294.