Proceedings, Fifteenth International Conference on
Principles of Knowledge Representation and Reasoning (KR 2016)

Consolidating Probabilistic Knowledge Bases via Belief Contraction

Glauber De Bona
University of São Paulo, Brazil
debona@ime.usp.br

Marcelo Finger
University of São Paulo, Brazil
mfinger@ime.usp.br

Márcio M. Ribeiro
University of São Paulo, Brazil
marciomr@usp.br

Yuri D. Santos
University of São Paulo, Brazil
yurids@ime.usp.br

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

Abstract

This paper is set to study the applicability of AGM-like operations to probabilistic bases. We focus on the problem of consistency restoration, also called consolidation or contraction by falsity. We aim to identify the reasons why the set of AGM postulates based on discrete operations of deletions and accretions is too coarse to treat finely adjustable probabilistic formulas. We propose new principles that allow one to deal with the consolidation of inconsistent probabilistic bases, presenting a finer method called liftable contraction. Furthermore, we show that existing methods for probabilistic consolidation via distance minimization are particular cases of the methods proposed.

1 Introduction

This paper studies the problem of restoring the consistency of probabilistic bases. We start by analyzing the inadequacy of the existing AGM paradigm, which was conceived within a discrete framework allowing only for the operations of insertion and deletion. But the continuous nature of probabilities demand for finer control if the goal is to obtain a *minimal change* when restoring the consistency of an inconsistent probabilistic base.

A probabilistic knowledge base may be inconsistent due to several reasons. For instance, the probabilities may have come from different sources or experts, and even a single source may fail to infer its own inconsistency with respect to the axioms of probability. Inconsistency undermines classical inference, calling for the *consolidation* of the probabilistic base, as in the following.

Problem 1.1 (Rain City). An agent has to decide whether to take an umbrella to a one-day trip to "Rain City". Gathering all the information available to her, she comes across two assessments:

- A website states the probability of raining as at least 60%.
- A friend estimates that probability as at most 30%.

Neither taking the umbrella nor getting wet is cost free, so the agent needs a probability bound to take her rational decision. How does she consolidate these statements into consistent probability bounds?

In this work we focus on the consolidation of conditional probabilities assigned to propositions of the form $P(\varphi|\psi) \geq q$, called *probabilistic conditional*, which reads "the probability of φ being true given that ψ is true is at least q". This is a generalization of the precise probabilities approach ($P(\varphi|\psi) = q$) of many logical systems, as the *probabilistic conditional logic* (Rödder 2000; Thimm 2013).

This formal system had its origins with Boole (1854), and was rediscovered several times (e.g. (de Finetti 1930; Hailperin 1965)) until Nilsson (1986) introduced it into the AI community under the name of *probabilistic logic*. Hansen et al. (2000) extended the language to embed conditional probabilities; an alternative semantics based on the maximum entropy principle was investigated by Kern-Isberner and Lucasiewicz (2004). The problem of handling inconsistent probabilities in propositional probabilistic logics has been approached via distance minimization methods by several authors (Jaumard, Hansen, and Poggi de Aragao 1991; Muiño 2011; Potyka and Thimm 2014; De Bona and Finger 2015), which is our focus.

The task of restoring consistency among propositions within a classical setting was formally tackled within the AGM framework (introduced by Alchourrón, Gärdenfors and Mankinson (1985)) as a contraction by the contradiction (Hansson 1997). That is, to consolidate a set of beliefs is to give up — *i.e.*, contract by — the belief in the contradiction. In this setting, consolidation is formalized via formulas discarding; however, in probabilistic logic it is usually the case that probability bounds are changed instead of being ruled out (Muiño 2011). An inconsistent belief of the form of $P(\varphi|\psi) \geq q$ may be weakened to $P(\varphi|\psi) \geq q'$, for some $q' < q$, to achieve consistency without losing all information conveyed by it. One way to minimize this information loss is to minimize the distance between the original, inconsistent probability bounds, and the new, consistent ones, *within a vector space*. This difference in approach, discarding versus changing formulas, makes the AGM paradigm unsuitable for the probabilistic consolidation through distance minimization to be grounded in. The main difficulty is that AGM contractions of a set of beliefs (a *base*) are constrained by rationality postulates that require, among other things, that the contracted base be a subset of the original one; but when a probability bound is changed, this requirement is violated.

Indeed, the AGM paradigm has room for weakening formulas while giving up beliefs, in what Hansson (1999) called *pseudo-contractions*. The idea is that a base being contracted by a proposition may receive new beliefs implied by the old ones. But until recently there was little clue on how to choose the consequences of the original formulas to include in the pseudo-contracted base. When contracting by the contradiction, this issue becomes even more evident, since the classical consequences of an inconsistent set of formulas is the whole logical language. It follows that any base can be a pseudo-contraction by the contradiction of any other base. Recently, Santos, Ribeiro and Wassermann (2015) have put forward a formal method to construct pseudo-contractions that can avoid this problem by using an arbitrary subclassical consequence operation. With such weaker consequence, it is not the case that any classical consequence of the original base may appear in the contraction, but just some of them.

In this work, we propose a method for the probabilistic consolidation via distance minimization founded on the AGM theory, starting with Santos, Ribeiro and Wassermann's (2015) approach. The concept of *liftable contraction* is introduced as a generalization of both contraction and pseudo-contraction; then the probabilistic consolidation via probability bound changing can be characterized as a *liftable contraction by the contradiction*. Thus we obtain a method for probabilistic contraction and consolidation. Even though the probabilistic consolidation methods via distance minimization are well-known in the literature (Thimm 2011; Muiño 2011), they are shown to be particular cases of our general probabilistic contraction methods, which are themselves original to the best of our knowledge.

The rest of the paper is organized as follows: in Section 2, we introduce the classical propositional and the probabilistic logics in order to fix notation; the approach to consolidate probabilistic knowledge bases via distance minimization is presented in Section 3; the postulates and procedures of AGM contraction and consolidation are summarized in Section 4; in Section 5, we relate the probabilistic consolidation via probability bound changing to the AGM consolidation through discarding formulas, using partial meet constructions and a weak consequence operation; the liftable contraction (and consolidation) operator is defined in Section 6 together with the postulates that characterize it, and probabilistic consolidation through probabilistic changing is shown to be a liftable consolidation; in Section 7, we discuss when and how the probabilistic consolidation can be extended to a contraction; finally, Section 8 brings conclusions and future work.

2 Preliminaries

In this work we deal mainly with knowledge bases formed by (imprecise) conditional probabilities on propositions from classical logic. A propositional logical language is a set of formulas formed by atomic propositions combined with logical connectives, possibly with punctuation elements (parentheses). We assume a finite set of symbols $X_n = \{x_1, x_2, x_3, \ldots, x_n\}$ corresponding to *atomic propositions (atoms)*. Formulas are constructed inductively with connectives ($\neg, \wedge, \vee, \rightarrow$) and atomic propositions as usual. The set of all these well-formed formulas is the propositional language over X_n, denoted by $\mathcal{L}_{X_n}$. Additionally, $\top$ denotes $x_i \vee \neg x_i$ for some $x_i \in X_n$, and $\bot$ denotes $\neg\top$.

Given a signature X_n, a *possible world* w is a conjunction of $|X_n| = n$ atoms containing either x_i or $\neg x_i$ for each $x_i \in X_n$. We denote by $W_{X_n} = \{w_1, \ldots, w_{2^n}\}$ the set of all possible worlds over X_n and say a $w \in W_{X_n}$ *entails* a $x_i \in X_n$ ($w \models x_i$) iff x_i is not negated in w. This entailment relation is extended to all $\varphi \in \mathcal{L}_{X_n}$ as usual. A set $B \subseteq \mathcal{L}_{X_n}$, a *belief base*, is said to be *consistent* (or *satisfiable*) iff there is a $w \in W_{X_n}$ such that $w \models \varphi$ for all $\varphi \in B$. Given a base $B \subseteq \mathcal{L}_{X_n}$, the classical consequence operation $Cn_{Cl}(.)$ can be defined as: $\varphi \in Cn_{Cl}(B)$ iff $B \cup \{\neg\varphi\}$ is inconsistent.

A *probabilistic conditional* (or simply *conditional*) is a statement of the form $P(\varphi|\psi) \geq q$, with the underlying meaning "the probability that φ is true given that ψ is true is at least q", where $\varphi, \psi \in \mathcal{L}_{X_n}$ are propositional formulas and q is a real number in $[0, 1]$. If ψ is a tautology, a conditional like $P(\varphi|\psi) \geq q$ is called an *unconditional probabilistic assessment*, usually denoted by $P(\varphi) \geq q$. We denote by $\mathcal{L}^P_{X_n}$ the set of all conditionals over $\mathcal{L}_{X_n}$, and $P(\varphi|\psi) \leq q$ abbreviates $P(\neg\varphi|\psi) \geq 1 - q$.

A *probabilistic interpretation* $\pi : W_{X_n} \rightarrow [0, 1]$, with $\sum_j \pi(w_j) = 1$, is a probability mass over the set of possible worlds, which induces a probability measure $P_\pi : \mathcal{L}_{X_n} \rightarrow [0, 1]$ given by $P_\pi(\varphi) = \sum\{\pi(w_j) | w_j \models \varphi\}$. A conditional $P(\varphi|\psi) \geq q$ is *satisfied* by π iff $P_\pi(\varphi \wedge \psi) \geq q P_\pi(\psi)$. Note that when $P_\pi(\psi) > 0$, a probabilistic conditional $P(\varphi|\psi) \geq q$ is constraining the conditional probability of φ given ψ; but any π with $P_\pi(\psi) = 0$ trivially satisfies that conditional (this semantics is adopted by Halpern (1990), Frisch and Haddawy (1994) and Lukasiewicz (1999), for instance). A *probabilistic knowledge base* (or simply a base) is a set of probabilistic conditionals $\Gamma \subseteq \mathcal{L}^P_{X_n}$. We denote by $\mathbb{K}$ the set of all probabilistic knowledge bases. A base Γ is *consistent* (or *satisfiable*) if there is a probability interpretation satisfying all conditionals $P(\varphi|\psi) \geq q \in \Gamma$. For a probabilistic base Γ, we denote by $Cn_{Pr}(\Gamma)$ the set of all conditionals $\alpha \in \mathcal{L}^P_{X_n}$ such that, if π satisfies Γ, then π satisfies α — this is known as the *standard semantics* (Haenni et al. 2011). [1]

A finite base $\Gamma \in \mathbb{K}$ is said to be *canonical* if it is such that, if $P(\varphi|\psi) \geq q, P(\varphi|\psi) \geq q' \in \Gamma$, then $q = q'$. That is, for each pair φ, ψ, only one probability lower bound can be assigned to $P(\varphi|\psi)$. We denote by $\mathbb{K}_c$ the set of all canonical probabilistic knowledge bases. Any base $\Gamma \in \mathbb{K}$ is equivalent to a canonical base — one can simply pick the highest probability lower bound of each set of conditionals $P(\varphi|\psi) \geq q$ on the same pair φ, ψ; if the resulting base is not finite, one can additionally repeat this procedure modulo equivalence on $\mathcal{L}_{X_n}$. [2]

[1] An alternative semantics for probabilistic entailment defines that a base Γ entails a conditional α if the probabilistic interpretation π that maximizes entropy while satisfying Γ also satisfies α. Nevertheless, the consolidation problem is the same.

[2] In the literature, instead of using canonical bases, probabilistic

3 Consolidation via Probability Bounds Adjustment

When a probabilistic knowledge base is inconsistent, classical inference trivializes, in the sense that everything logically follows from that base. This calls for inconsistency handling in bases in order to restore consistency. The procedure of restoring the consistency of a knowledge base is called *consolidation*. We call a *consolidation operator* any function $! : \mathbb{K} \to \mathbb{K}$ (or $! : \mathbb{K}_c \to \mathbb{K}_c$) that takes (canonical) probabilistic knowledge bases Γ and returns consistent ones $!(\Gamma) = \Gamma!$. Several approaches to implement this operation have been proposed in the literature, mainly focusing on changing the probabilities' numeric values (Muiño 2011; Potyka and Thimm 2014). When a probabilistic base is inconsistent, it is more intuitive to fix it by adjusting the numeric values instead of the propositions to which the probabilities are assigned, due to the notions of order and distance naturally held by numbers. Since changes to consolidate bases are usually required to be minimal in some sense, numbers provide straightforward ways to define such minimality. We present this approach of consolidation via probabilities adjustment focusing on canonical bases, as it is the commonest form, adapting the work from Muiño's (2011).

When probabilistic knowledge bases are defined as sets of precise probability assessments, consolidations are expected (and forced) to be precise as well (Muiño 2011; Potyka and Thimm 2014). We adopt a slightly different approach, allowing probability lower bounds only to be lowered — equivalently, upper bounds can only be increased —, as Hansen et. al (2000) proposed. A precise probability assessment, like $P(\varphi) = q$, could be represented in our language via the pair $P(\varphi) \geq q, P(\neg\varphi) \geq 1 - q$. While the precise probability approach conceives consolidation operators that modify $P(\varphi) = q$ to $P(\varphi) = q'$, for instance for a $q' < q$, the method here described would return the pair $P(\varphi) \geq q', P(\neg\varphi) \geq 1 - q$. That is, the upper bound would not be strengthened, and the consolidated base would have imprecise probabilities. When these probability changes are minimal in some sense, both approaches yield logically equivalent bases, but allowing imprecise probabilities fits better the AGM paradigm, as we will show.

Given a canonical base $\Gamma \in \mathbb{K}_c$, we define its characteristic function $\Lambda_\Gamma : [0,1]^{|\Gamma|} \to \mathbb{K}$ such that, if $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ and $q' = \langle q'_1, \ldots, q'_m \rangle$ is a vector in $[0,1]^m$, then $\Lambda_\Gamma(q') = \{P(\varphi_i|\psi_i) \geq q'_i | 1 \leq i \leq m\}$. In other words, the function Λ_Γ just changes the lower bounds in Γ. Consolidation operators can then be implemented through characteristic functions: $!(\Gamma) = \Lambda_\Gamma(q')$ for some $q' \in [0,1]^{|\Gamma|}$.

Every inconsistent canonical base Γ can be vacuously consolidated by substituting 0 for all its probability lower bounds, through an operator ! such that $!(\Gamma) = \Lambda_\Gamma(\langle 0, \ldots, 0 \rangle)$ is consistent for all Γ. Nevertheless, this drastic procedure misses all the point of minimal information loss, or minimal change in the base. A natural idea to consolidate a canonical base Γ is then to take the consistent probability lower bounds that are in some sense maximally closer to the bounds in Γ; *i.e.*, to make minimal changes in the probability bounds. To make this notion precise, we define an element-wise relation $\leq \subset \mathbb{R}^m \times \mathbb{R}^m$ between vectors $q, q' \in \mathbb{R}^m$: $q \leq q'$ iff $q_i \leq q'_i$ for all $1 \leq i \leq m$. Analogously, we define $q \geq q'$ iff $q' \leq q$; and $q < q'$ iff $q \leq q'$ and $q \neq q'$. Now we can define when a consolidation is maximal in terms of information — introduced as dominant consolidations in (De Bona and Finger 2015):

knowledge bases have been defined as ordered sets, allowing for repeated formulas; but this loosens the link to the belief revision approach, which operates on non-ordered sets.

Definition 3.1. Given a canonical base $\Gamma = \{P(\varphi_i|\psi_i) \geq q_i | 1 \leq i \leq m\}$ in $\mathbb{K}_c$ and a vector $q' \in [0,1]^m$, the base $\Lambda_\Gamma(q')$ is a *maximal consolidation* of Γ if it is consistent and, for every $q'' \in [0,1]^m$ such that $q' < q'' \leq q$, $\Lambda_\Gamma(q'')$ is inconsistent.

Note that maximal consolidations are always canonical, for only the probability bounds are changed. Intuitively, a maximal consolidation is such that, if some probability lower bound were less relaxed, the base would still be inconsistent. For instance, consider the example below:

Example 3.2 (Rain City Revisited)**.** Consider the canonical base $\Gamma = \{P(\varphi) \geq 0.6, P(\varphi) \leq 0.3\}$. An adjustment in the probability bounds that makes the base consistent is $\Gamma' = \{P(\varphi) \geq 0.3, P(\varphi) \leq 0.6\}$. Nonetheless, another adjustment conserves strictly more information, $\Gamma'' = \{P(\varphi) \geq 0.3, P(\varphi) \leq 0.3\}$, so that Γ' is not a maximal consolidation of Γ. The reader can note that Γ'' is a maximal consolidation, for the probability assessment $P(\varphi) \geq 0.3$ could not be consistently less relaxed.

Typically, an inconsistent canonical base has several (often infinite) maximal consolidations, corresponding to a Pareto frontier. In the example above, any $\Psi = \{P(\varphi) \geq q_1, P(\varphi) \leq q_2\}$ such that $q_1 \in [0.3, 0.6]$ and $q_1 = q_2$ would be a maximal consolidation of Γ. However, some maximal consolidations of Γ can be said to be closer to Γ than others, if one considers probability bounds in a vector space. To construct a consolidation operator that returns some closer maximal consolidation, one can employ methods that minimize distances between these vectors of probability bounds, plus some criterion to guarantee uniqueness. To review the commonest set of distances applied to this end in the literature (Muiño 2011), we define the p-norm of a vector $\langle q_1, \ldots, q_m \rangle \in \mathbb{R}^m$, for any positive integer p, as being:

$$\|q\|_p = \sqrt[p]{\sum_{i=1}^{m} |q_i|^p}$$

Taking the limit $p \to \infty$, we also define $\|q\|_\infty = \max_i |q_i|$. Now the p-norm distance between two vectors $q, q' \in \mathbb{R}^m$ can be defined as:

$$d_p(q, q') = \|q - q'\|_p$$

For instance, d_1 is the Manhattan (or absolute) distance, d_2 is the Euclidean distance and d_∞ is the Chebyshev distance.

Given a canonical base $\Gamma = \Lambda_\Gamma(q)$, each p-norm distance defines a set of consistent probability lower bounds vectors that are closest to q — each defining a "closest" consolidation. Let $D_p : \mathbb{K}_c \to 2^{\mathbb{K}_c}$ be a function that returns the closest consolidations according to a p-norm distance:

$$D_p(\Gamma) = \{\Lambda_\Gamma(q')|\Gamma = \Lambda_\Gamma(q) \text{ , } \Lambda_\Gamma(q') \text{ is consistent and } d_p(q,q') \text{is minimum}\}$$

If Γ has only unconditional probability assessments, any finite $p > 1$ yields a unique closest vector of probability bounds and $D_p(\Gamma)$ is a singleton, but in the general case this does not hold, due to non-convexity. Hence, further constraints must be used to choose a single maximal consolidation in order to specify a consolidation operator as a function. For instance, one can consider the closest base with maximum entropy, the closest base that is the most preferred according to some relation, or even the first closest base found by some implementation that computes the distance minimisation. For any positive p in $\mathbb{N} \cup \{\infty\}$, we define the consolidation operator $!_p : \mathbb{K}_c \to \mathbb{K}_c$ as a function that takes an arbitrary base $\Gamma = \Lambda_\Gamma(q)$ and returns a base $\Lambda_\Gamma(q')$ that minimizes $d_p(q, q')$, employing arbitrary criteria to select a single q'. That is, for each $\Gamma \in \mathbb{K}_c$:

$$!_p(\Gamma) \in D_p(\Gamma)$$

We point out that a myriad of different consolidation operators can implement a $!_p$, depending on how a single natural consolidation is chosen when there are several closest ones. Henceforth, we investigate general consolidation operators that return a maximal consolidation for canonical bases, no matter how it is chosen. We define a *maximal consolidation operator* as a consolidation operator $! : \mathbb{K}_c \to \mathbb{K}_c$ such that, for all canonical bases $\Gamma \in \mathbb{K}_c$, $!(\Gamma)$ is a maximal consolidation operator.

4 Consolidation and Contraction in AGM Theory

The AGM paradigm is a theory that models changes in the belief state of a rational agent (see, for instance, (Hansson 1999)). A belief base is a way to model a belief state via a set of formulas from some logical language, in which the agent believes in the logical consequences of this set. Probabilistic knowledge bases can be understood as belief bases, if we consider a hypothetical underlying agent, which might for instance be implemented in an autonomous system. Three main belief change operations in the AGM paradigm are: expansion, contraction and revision. The second one handles the giving up of beliefs, through which consolidation is defined as the contraction by the contradiction; *i.e.*, to consolidate a belief state is to give up the belief on the contradiction. To see how belief contraction relates to the consolidation operators we presented, we quickly review the AGM approach to giving up beliefs on bases.

The whole AGM theory is grounded in a set of rationality postulates for belief change operations, which have a correspondence with procedures that implement these operations. Even though the original constructions and rationality postulates were introduced by Alchourrón, Gärdenfors and Mankinson (1985) for (logically closed) belief sets, Hansson (1992) generalized them to belief bases, which are sets of logical formulas. Consider a general language $\mathcal{L}$ containing $\perp$ and a consequence operation $Cn(.)$ over it. A *belief base* is any set $B \subseteq \mathcal{L}$. A contraction operator for a base B is formalized as a function $- : \mathcal{L} \to 2^{\mathcal{L}}$ that takes a formula $\alpha \in \mathcal{L}$ as argument and returns a base $-(\alpha) = B - \alpha \in 2^{\mathcal{L}}$ — the contraction of B by α. For a function to be called a *contraction* operator for B, it must satisfy the following two rationality postulates, for any $\alpha \in \mathcal{L}_i$:

(Inclusion) $B - \alpha \subseteq B$.

(Success) If $\alpha \notin Cn(\varnothing)$, $\alpha \notin Cn(B - \alpha)$.

In words, the contraction of B by α should introduce no new beliefs in the base, and the resulting base should not imply α, as long as α is not a tautology. A weaker form of inclusion is (Hansson 1999):

(Logical Inclusion) $Cn(B - \alpha) \subseteq Cn(B)$.

An operator satisfying (Success) and (Logical Inclusion) is called a *pseudo-contraction* (Hansson 1999).

The next postulate deals with the notion of minimal changes or the idea of maximal information preservation. While contracting a base by a belief α, it is sensible not to discard beliefs that are not involved in the derivation of α, which is enforced by a postulate called (Relevance):

(Relevance) If $\beta \in B \setminus B - \alpha$, there is a B' such that $B - \alpha \subseteq B' \subseteq B$, $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$.

A final postulate states that contracting B by formulas that are equivalent (modulo B) should yield the same result:

(Uniformity) If it is the case that, for all $B \subseteq B$, we have $\alpha \in Cn(B)$ iff $\beta \in Cn(B)$, then $B - \alpha = B - \beta$.

Hansson (1999) proved, for classical propositional logic, that any contraction on a belief base satisfying (Success), (Inclusion), (Relevance) and (Uniformity) can be expressed in a specific form. To present it, we need some definitions:

Definition 4.1. Given a base $B \subseteq \mathcal{L}$ and a formula $\alpha \in \mathcal{L}$, the *remainder set* $B \perp \alpha$ is such that $X \in B \perp \alpha$ iff:

- $X \subseteq B$;
- $\alpha \notin Cn(X)$;
- for any set Y, if $X \subsetneq Y \subseteq B$, $\alpha \in Cn(Y)$.

Definition 4.2. A function γ is a *selection function* for a base $B \subseteq \mathcal{L}$ iff:

- if $B \perp \alpha \neq \varnothing$, $\varnothing \neq \gamma(B \perp \alpha) \subseteq B \perp \alpha$;
- otherwise $\gamma(B \perp \alpha) = \{B\}$.

Definition 4.3. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ is a *partial meet contraction* for a base $B \subseteq \mathcal{L}$ if, for any formula $\alpha \in \mathcal{L}$, $B - \alpha = \bigcap \gamma(B \perp \alpha)$ for some selection function γ. If $\gamma(B \perp \alpha)$ is a singleton for every $\alpha \in \mathcal{L}$, $-$ is called a *maxichoice contraction*.

For classical propositional logic, Hansson (1999) linked the rationality postulates to the partial meet construction. To present this and other results for a general language $\mathcal{L}$ with an arbitrary consequence operation $Cn(.)$, we have to impose some restrictions on the latter. Below, we define some useful properties a consequence operation may hold:

Definition 4.4. Consider a general language $\mathcal{L}$, arbitrary bases $B, B', B'' \subseteq \mathcal{L}$, an arbitrary formula $\alpha \in \mathcal{L}$ and a consequence operation $Cn'(.)$. A consequence operation $Cn(.)$ satisfies

- *Monotonicity* if $B \subseteq B'$ implies $Cn(B) \subseteq Cn(B')$;
- *Idempotence* if $Cn(Cn(B)) \subseteq Cn(B)$;
- *Inclusion* if $B \subseteq Cn(B)$;
- the *Upper Bound Property* if, for every $B' \subseteq B$ with $\alpha \notin Cn(B')$, there is a $B'' \supseteq B'$ such that $B'' \in B \perp \alpha$;
- *Cn'-Dominance* if $Cn(B) \subseteq Cn'(B)$;
- *Subclassicality* if Cn satisfies Cn_{Cl}-Dominance.

If Cn satisfies monotonicity, inclusion and idempotence, we say it is *Tarskian*.

For instance, Cn_{Cl} for the propositional language and Cn_{Pr} for probabilistic conditionals are Tarskian consequence operations, as they are defined through models, but only the former fully satisfies the upper bound property, as we shall see.

Now, Hansson's representation theorem can be generalized (Wassermann 2000)[3]:

Theorem 4.5. *Let Cn be a consequence operation satisfying monotonicity and the upper bound property. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $-$ is a partial meet contraction.*

To axiomatize maxichoice contraction, (Relevance) can be strengthened (Hansson 1999):

(Fullness) If $\beta \in B$ and $\beta \notin B - \alpha$, then $\alpha \notin Cn(B - \alpha)$ and $\alpha \in Cn((B - \alpha) \cup \{\beta\})$.

Hansson (1999) proved the result below for classical logic, and Wassermann (2000) generalized it:

Theorem 4.6. *Let Cn be a consequence operation satisfying monotonicity and the upper bound property. The operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies (Success), (Inclusion), (Fullness) and (Uniformity) iff $-$ is a maxichoice contraction.*

Hansson (1997) defines a *partial meet consolidation* as a partial meet contraction by the contradiction ($\perp$), which is here adapted:

Definition 4.7. The operation $B!$ is a *partial meet consolidation* for a base $B \subseteq \mathcal{L}$ if $B! = \bigcap \gamma(B \perp^{\perp})$ for some selection function γ. If $\gamma(B \perp^{\perp})$ is a singleton, ! is called a *maxichoice consolidation*.

Except for (Uniformity), the postulates for contraction apply to consolidation ($B!$), as contraction by the contradiction, just taking $\alpha = \perp$; and (Success) for consolidation is also called (Consistency). Hansson (1999) proved the following representation result, generalized by Wassermann (2000):

[3]All proofs are in the appendix of the version available online through the link
www.ime.usp.br/~mfinger/etc/sub/KR2016.pdf.

Theorem 4.8. *Let Cn be a consequence operation satisfying monotonicity, the upper bound property and such that $\perp \notin Cn(\varnothing)$. $B!$ satisfies (Success), (Inclusion), (Relevance) iff $B!$ is a partial meet consolidation.*

Corollary 4.9. *Let Cn be a consequence operation satisfying monotonicity, the upper bound property and such that $\perp \notin Cn(\varnothing)$. $B!$ satisfies (Success), (Inclusion), (Fullness) iff $B!$ is a maxichoice consolidation.*

By taking canonical probabilistic knowledge bases as belief bases over the language $\mathcal{L}^P_{X_n}$ and considering the probabilistic consequence operation $Cn_{Pr}(.)$, we can evaluate the consolidation operators introduced in Section 3 with respect to the AGM rationality postulates. For a given base, a maximal consolidation operator ! is not a contraction by the contradiction, for formulas are not being discarded, but changed, and (Inclusion) is violated. The postulate of (Logical Inclusion) is satisfied, so ! could be viewed as a pseudo-contraction by the contradiction. Nevertheless, this postulate is vacuous when Γ is inconsistent — which is the only meaningful application of consolidation —, for $Cn_{Pr}(\Gamma) = \mathcal{L}^P_{X_n}$. In the following, we look for ways to embed maximal consolidation operators, like $!_p$, into the AGM paradigm in a more suitable way.

5 Changing Probability Bounds versus Discarding Formulas

When probability lower bounds in a canonical probabilistic knowledge base are decreased to consolidate it, this operation satisfies (Logical Inclusion) in a particular way. Not any consequence of a base Γ can appear in its consolidation, but only consequences of each single probabilistic conditional, just with the probability bound modified. Santos, Ribeiro and Wassermann (2015) investigated contractions on propositional bases considering, besides the classical consequence Cn_{Cl}, an arbitrary Tarskian, subclassical consequence operation $Cn^\star$. Based on their ideas, we explore different ways of connecting probabilistic consolidation via changing lower bounds with consolidation via partial meet. First, we define an element-wise consequence operation Cn_{ew} for the probabilistic language $\mathcal{L}^P_{X_n}$.

Definition 5.1. The function $Cn_{ew} : \mathbb{K} \to \mathbb{K}$ is a consequence operation such that, for all $\Gamma \in \mathbb{K}$:

$$Cn_{ew}(\Gamma) = \{P(\varphi|\psi) \geq q' | P(\varphi|\psi) \geq q \in \Gamma \text{ and } 0 \leq q' \leq q\}$$

Proposition 5.2. *Cn_{ew} is Tarskian and satisfies Cn_{Pr}-dominance.*

Given a canonical probabilistic base $\Gamma \in \mathbb{K}_c$ and a maximal consolidation operator !, we have $!(\Gamma) \subseteq Cn_{ew}(\Gamma)$ — equivalently, by monotonicity, idempotency ($\rightarrow$) and inclusion ($\leftarrow$), $Cn_{ew}(!(\Gamma)) \subseteq Cn_{ew}(\Gamma)$ —, a special kind of inclusion. Furthermore, when Γ is inconsistent, it is not necessarily the case that $Cn_{ew}(\Gamma) = \mathcal{L}^P_{X_n}$, thus this inclusion is not trivially satisfied, as it happens with (Logical Inclusion) using Cn_{Pr}. This motivates the following postulate (Santos, Ribeiro, and Wassermann 2015), for a general logical language $\mathcal{L}$ and an arbitrary consequence operation $Cn^\star$:

(Inclusion*) $B - \alpha \subseteq Cn^\star(B)$.

If $Cn^\star = Cn_{ew}$, (Inclusion*) is the kind of inclusion any maximal consolidation operator respects, when viewed as a contraction by the contradiction. In other words, decreasing probability lower bounds in a canonical base Γ is the same as picking a canonical subset of $Cn_{ew}(\Gamma)$.

In a similar way, Santos et al. (2015) propose starred versions for (Relevance) and (Uniformity).

(Relevance*) If $\beta \in Cn^\star(B) \setminus B - \alpha$, there is a B' such that $B-\alpha \subseteq B' \subseteq Cn^\star(B)$, $\alpha \notin Cn(B')$ and $\alpha \in Cn(B' \cup \{\beta\})$.

(Uniformity*) For all $B' \subseteq Cn^\star(B)$, if it is the case that $\alpha \in Cn(B')$ iff $\beta \in Cn(B')$, then $B - \alpha = B - \beta$.

The postulate of (Relevance*), together with (Inclusion*), forces a contraction — and a consolidation — to be closed under $Cn^\star$, what is called (Enforced Closure*) (Santos, Ribeiro, and Wassermann 2015):

(Enforced Closure*) $B - \alpha = Cn^\star(B - \alpha)$.

As maximal consolidation operators return canonical bases, which are not closed under Cn_{ew}, (Relevance*) is violated. Nonetheless, if we consider the closure of the resulting base under Cn_{ew}, this postulate is recovered:

Proposition 5.3. *Let $!_{MC} : \mathbb{K}_c \to \mathbb{K}_c$ be a maximal consolidation operator and consider $Cn^\star = Cn_{ew}$. For any $\Gamma \in \mathbb{K}$, the consolidation operation $!$ for Γ defined as $\Gamma! = Cn_{ew}(\Gamma!_{MC})$ satisfies (Relevance*).*

In the classical propositional case, Santos et al. (2015) proved that, for a compact, Tarskian, subclassical $Cn^\star$, a contraction operator $-$ for a base B satisfies (Success), (Inclusion*), (Relevance*) and (Uniformity*) iff $B - \alpha = \bigcap \gamma(Cn^\star(B) \perp \alpha)$ for some selection function γ. Consequently, if $!(B) = B - \perp$ for a $-$ satisfying these four postulates, $!$ is a partial meet contraction by the contradiction of the starred closure $Cn^\star(B)$. To prove the analogous result for probabilistic logic, where $\perp$ can denote $P(\perp) \geq 1$, we need an intermediate result, for the probabilistic consequence operation Cn_{Pr} is not compact:

Lemma 5.4 (Upper bound property for probabilistic consolidation)**.** *Let $\Gamma \in \mathbb{K}_c$ be a canonical base. For every consistent $\Delta \subseteq Cn_{ew}(\Gamma)$, there is a Δ' such that $\Delta \subseteq \Delta' \subseteq Cn_{ew}(\Gamma)$ and $\Delta' \in Cn_{ew}(\Gamma) \perp^\perp$.*

Lemma 5.5. *Let $\Gamma \in \mathbb{K}_c$ be a canonical base. $\Psi \in \mathbb{K}_c$ is a maximal consolidation of Γ iff $Cn_{ew}(\Psi) \in Cn_{ew}(\Gamma) \perp^\perp$.*

The lemma above states that besides the remainder set $Cn_{ew}(\Gamma) \perp^\perp$ being well-defined, it is characterized by the maximal consolidations. Now a representation Theorem can be proved:

Theorem 5.6. *Consider a base $\Gamma \in \mathbb{K}_c$. An operation $\Gamma!$ satisfies (Success), (Inclusion*) and (Relevance*), for $Cn^\star = Cn_{ew}$, iff $\Gamma! = \bigcap \gamma(Cn_{ew}(\Gamma) \perp^\perp)$, for some selection function γ.*

From this theorem, the first result connecting maximal consolidation operators and the AGM theory follows, in which $\circ$ denotes function composition:

Corollary 5.7. *Consider the consequence operation $Cn_{ew} = Cn^\star$, a maximal consolidation operator $!_{MC}$ and a consolidation operation $!$ for each $\Gamma \in \mathbb{K}_c$ defined as $\Gamma! = Cn_{ew} \circ !_{MC}(\Gamma)$. For each canonical base $\Gamma \in \mathbb{K}_c$, $\Gamma!$ satisfies (Success), (Inclusion*) and (Relevance*).*

This result relates a consolidation operation $!$ that returns the closure of a maximal consolidation under Cn_{ew} to a partial meet construction. Nonetheless, it is not the case that any construction satisfying (Success), (Inclusion*) and (Relevance*) is the Cn_{ew}-closure of a maximal consolidation. From Lemma 5.5, it is easy to see that any consolidation operation that returns the Cn_{ew}-closure of a maximal consolidation can be seen as a partial meet construction that selects a single element out of the remainder set — a maxichoice contraction. Hence, to fully characterize these operations, we define the starred version of (Fullness):

(Fullness*) If $\beta \in Cn^\star(B)$ and $\beta \notin B - \alpha$, then $\alpha \notin Cn(B - \alpha)$ and $\alpha \in Cn((B - \alpha) \cup \{\beta\})$.

Corollary 5.8. *Consider the consequence operation $Cn_{ew} = Cn^\star$. The consolidation operation $\Gamma!$, for any $\Gamma \in \mathbb{K}_c$, satisfies (Success), (Inclusion*) and (Fullness*) iff there is a maximal consolidation operator $!_{MC}$ such that $\Gamma! = Cn_{ew} \circ !_{MC}(\Gamma)$ for all $\Gamma \in \mathbb{K}_c$.*

If the input base is already closed under Cn_{ew} (*i.e.*, $Cn_{ew}(\Gamma) = \Gamma$), we can link any maximal consolidation operator $!_{MC}$ to the original (not starred) postulates, through a canonical base Ψ such that $Cn_{ew}(\Psi) = \Gamma$. To do that, we need a way to get such canonical base Ψ:

Definition 5.9. For any probabilistic knowledge base Γ such that $Cn_{ew}(\Gamma) = Cn_{ew}(\Psi)$ for a canonical $\Psi \in \mathbb{K}_c$[4]:

$$Cn_{ew}^{-1}(\Gamma) = \{P(\varphi|\psi) \geq q \in \Gamma | \text{ there is no } P(\varphi|\psi) \geq q' \in \Gamma \text{ with } q' > q\}$$

Now a second way to relate a maximal consolidation operator to the partial meet construction follows:

Corollary 5.10. *Consider a base $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$. The consolidation operation $!(\Gamma)$, for any $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$, satisfies (Success), (Inclusion) and (Fullness) iff there is a maximal consolidation operator $!_{MC}$ such that $!(\Gamma) = Cn_{ew} \circ !_{MC} \circ Cn_{ew}^{-1}(\Gamma)$ for all $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}$.*

From Theorem 5.6, it follows that any consolidation operation $\Gamma! = Cn_{ew} \circ !_{MC} \circ Cn_{ew}^{-1}(\Gamma)$, for a Cn_{ew}-closed Γ, is a partial meet consolidation $\Gamma! = \bigcap \gamma(\Gamma \perp^\perp)$. Indeed, $\Gamma! = Cn_{ew}(\Psi)$, for some maximal consolidation Ψ of $Cn_{ew}^{-1}(\Gamma)$.

Until now, we have founded on the AGM paradigm two kinds of consolidation operations based on a maximal consolidation operator $!_{MC}$:

- $Cn_{ew} \circ !_{MC}(\Gamma)$ for each canonical base $\Gamma \in \mathbb{K}_c$ returns a base closed under Cn_{ew}, satisfying (Success), (Inclusion*) and (Fullness*);
- $Cn_{ew} \circ !_{MC} \circ Cn_{ew}^{-1}(\Gamma)$ each base $\Gamma \in \mathbb{K}$ closed under Cn_{ew} returns a base closed under Cn_{ew}, satisfying (Success), (Inclusion) and (Fullness).

[4] This avoids ill-defined cases, such as $\Gamma = \{P(\varphi) \geq q | q < 1\}$.

However, we have still not characterized maximal consolidation operators themselves with postulates or partial meet constructions. Corollary 5.10 points a way to do that via the starred closure of the contraction operator, which motivates a generalization of the contraction concept.

6 Liftable Contractions

Given a contraction operator $-$ for a belief base B over a general language $\mathcal{L}$ and a consequence operation $Cn^\star$, we define the *starred closure* of $-$ as the operator $-_\star$ for $Cn^\star(B)$ such that $Cn^\star(B) -_\star \alpha = Cn^\star(B - \alpha)$. That is, if $-$ takes a formula α and returns $B' = B - \alpha$, $-_\star$ takes α and returns $Cn^\star(B')$. The starred closure $!_\star$ of a consolidation operator ! can be defined through $(Cn^\star(B))!_\star = Cn^\star(B!)$.

The (Logical Inclusion) postulate can be seen as a restriction on the Cn-closure of a contraction operator. That is, taking $Cn^\star = Cn$, (Logical Inclusion) requires that the starred closure of the operator satisfy (Inclusion): $Cn^\star(B) -_\star \alpha = Cn^\star(B - \alpha) \subseteq Cn^\star(B)$. Similarly, if $Cn^\star$ were the identity function, (Inclusion) can be viewed as a restriction on the starred operator as well, for $Cn^\star(B) = B$. For a general consequence operation $Cn^\star$, we can "lift" in this way the (Inclusion) postulate to the starred closure of the operator:

(***Lifted Inclusion)** The operator $-$ for a base B is such that $-_\star$ for $Cn^\star(B)$ satisfies (Inclusion).

Any postulate can be similarly lifted, but note that (Success) is equivalent to (*Lifted Success), given some assumptions:

Proposition 6.1. *If Cn satisfies monotonicity and idempotency, and $Cn^\star$, inclusion and Cn-dominance, (Success) is equivalent to (*Lifted Success).*

We call an operator satisfying (*Lifted Inclusion) and (*Lifted Success) a *liftable contraction*. As we pointed out, contractions and pseudo-contractions are particular cases of liftable contractions, where $Cn^\star(B) = B$ and $Cn^\star(B) = Cn(B)$, respectively. [5] Similarly, we can lift (Relevance), (Uniformity) and (Fullness):

(*Lifted Relevance) The operator $-$ for a base B is such that $-_\star$ for $Cn^\star(B)$ satisfies (Relevance).

(*Lifted Uniformity) The operator $-$ for a base B is such that $-_\star$ for $Cn^\star(B)$ satisfies (Uniformity).

(*Lifted Fullness) The operator $-$ for a base B is such that $-_\star$ for $Cn^\star(B)$ satisfies (Fullness).

The following result shows how the lifted postulates relate to the starred ones:

Proposition 6.2. *Consider a consequence operation Cn that satisfies monotonicity and idempotence and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. The postulate of (Inclusion*) is equivalent to (*Lifted Inclusion). The postulate of (Relevance*) implies (*Lifted Relevance). The postulate of (Uniformity*) implies (*Lifted Uniformity). The postulate of (Fullness*) implies (*Lifted Fullness).*

[5] Liftable contractions are in some sense converse to base-generated contractions, defined in (Hansson 1999). While the former is an operation on bases defined through their $Cn^\star$-closure, the latter is an operation on Cn-closed sets defined via their bases.

To guarantee equivalence between the lifted and the starred postulates, it suffices to require the closure of the resulting base under $Cn^\star$:

Proposition 6.3. *Consider a consequence operation Cn that satisfies monotonicity and idempotence and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. If (Enforced Closure*) is satisfied, then (*Lifted Relevance) is equivalent to (Relevance*), (*Lifted Uniformity) is equivalent to (Uniformity*) and (*Lifted Fullness) is equivalent to (Fullness*).*

To link the lifted postulates to partial meet constructions for the starred closure, we define:

Definition 6.4. We say a liftable contraction is a *liftable partial meet contraction (consolidation)* if its starred closure is a partial meet contraction (consolidation); we say it is a *liftable maxichoice contraction (consolidation)* if its starred closure is a maxichoice contraction (consolidation).

The implementation of a liftable partial meet contraction is only partially determined through its starred closure, unless we require the closure of the resulting base under $Cn^\star$:

Proposition 6.5. *Consider a monotonic consequence operation Cn that satisfies the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. A contraction operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base $B \subseteq 2^{\mathcal{L}}$ satisfies the *lifted versions of the postulates of success, inclusion, relevance (fullness) and uniformity iff it is a liftable partial meet contraction (liftable maxichoice contraction). Furthermore, $-$ additionally satisfies (Enforced Closure*) iff $B - \alpha = \bigcap \gamma(Cn^\star(B)\perp\alpha)$ (and $\gamma(Cn^\star(B)\perp\alpha)$ is a singleton) for all $\alpha \in \mathcal{L}$ and for some selection function γ.*

For consolidations, it is always the case that $\alpha = \perp$, and one can drop (*Lifted Uniformity) from the result above:

Proposition 6.6. *Consider a monotonic consequence operation Cn that satisfies the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. A consolidation operation $B!$ satisfies the *lifted versions of the postulates of success, inclusion, and relevance (fullness) iff it is a liftable partial meet consolidation (liftable maxichoice consolidation). Furthermore, if $B!$ satisfies (Enforced Closure*), then $B! = \bigcap \gamma(Cn^\star(B)\perp^\perp)$ for some selection function γ ($B! \in (Cn^\star(B)\perp^\perp)$).*

Any maximal consolidation operator satisfies the *lifted postulates, including (*Lifted Fullness), but not (Enforced Closure*). Hence, from Proposition 6.6, maximal consolidations are liftable maxichoice consolidations. To uniquely characterize maximal consolidation operators, we need to require that output bases be canonical. Given a general consequence operation $Cn^\star$ for a language $\mathcal{L}$, we say a *canonical form* is a function $f_\star : 2^{\mathcal{L}} \to 2^{\mathcal{L}}$ such that $Cn^\star(B) = Cn^\star(f_\star(B))$, and $Cn^\star(B) = Cn^\star(B')$ implies $f_\star(B) = f_\star(B')$. Note that these conditions also imply $f_\star(f_\star(B)) = f_\star(B)$. For any base B, we say $f_\star(B)$ is its canonical form (regarding $f_\star$). Given a canonical form $f_\star$ for a consequence operation $Cn^\star$, we can define:

($f_\star$-Canonicity) $B - \alpha = f_\star(B - \alpha)$.

With this extra postulate, the partial meet construction for the liftable contraction is uniquely determined by the selection function γ, and (*Lifted Fullness) forces γ to return a singleton:

Proposition 6.7. *Consider a consequence operation Cn that satisfies monotonicity, idempotence and the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. An operator $- : \mathcal{L} \to 2^{\mathcal{L}}$ for a base B is a liftable partial meet contraction (liftable maxichoice contraction) and satisfies ($f_\star$-Canonicity) iff, for any formula $\alpha \in \mathcal{L}$, $B - \alpha = f_\star(\bigcap \gamma(Cn^\star(B)\perp\alpha))$ for some selection function γ (and $\gamma(Cn^\star(B)\perp\alpha)$ is a singleton).*

Corollary 6.8. *Consider a consequence operation Cn that satisfies monotonicity, idempotence and the upper bound property and a Tarskian consequence operation $Cn^\star$ satisfying Cn-dominance. An operation $B!$ is a liftable partial meet consolidation (liftable maxichoice consolidation) and satisfies ($f_\star$-Canonicity) iff, $B! = f_\star(\bigcap \gamma(Cn^\star(B)\perp^\perp))$ for some selection function γ (and $\gamma(Cn^\star(B)\perp^\perp)$ is a singleton).*

As the function Cn_{ew}^{-1} satisfies the properties of a canonical form $f_\star$ for the consequence operation Cn_{ew}, we can say that maximal consolidation operators satisfy the corresponding ($f_\star$-Canonicity). Even though the probabilistic logic is not compact, using Lemma 5.4 we can prove some representation results:

Lemma 6.9. *Consider $Cn^\star = Cn_{ew}$ and $f_\star = Cn_{ew}^{-1}$. The operation $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Fullness) and ($f_\star$-Canonicity) for all $\Gamma \in \mathbb{K}_c$ iff $!_{MC} : \mathbb{K}_c \to \mathbb{K}_c$, with $!_{MC}(\Gamma) = \Gamma!$ for all $\Gamma \in \mathcal{L}$, is a maximal consolidation operator.*

If only (*Lifted Relevance) is required, and not (*Lifted Fullness), we obtain constructions that select the infimum of the lower bounds in a set of maximal consolidations.

Definition 6.10. Consider a base $\Gamma \in \mathbb{K}_c$ and a set of vectors $S \subseteq [0,1]^{|\Gamma|}$. We define $\inf\{\Lambda_\Gamma(q) | q \in Q\} = \Lambda_\Gamma(q')$, where $q' \in [0,1]^{|\Gamma|}$ is such that $q'_i = \inf\{q_i | q \in S\}$ for every $1 \leq i \leq |\Gamma|$.

Lemma 6.11. *Consider $Cn^\star = Cn_{ew}$, $f_\star = Cn_{ew}^{-1}$, a $\Gamma \in \mathbb{K}_c$. $\Gamma!$ satisfies (*Lifted Success), (*Lifted Inclusion), (*Lifted Relevance) and ($f_\star$-Canonicity) iff $\Gamma! = \inf M$, for a set M of maximal consolidations of Γ.*

The set M can be defined through the remainder set, as the elements of the former (maximal consolidations) are canonical forms of elements of the latter: $M = \{f_\star(\Psi) | \Psi \in \gamma(Cn^\star(\Gamma)\perp^\perp)\}$. The *full meet* criterion, which requires that a partial meet contraction be the intersection of all elements in the remainder set, is met when M is the set of all maximal consolidations.

By leaving the maxichoice contraction, allowing M to be a set of maximal consolidations, we do not need methods that specify a single one, but only a set of preferred ones. We said in Section 3 that minimizing the p-norm distance does not necessarily yield a unique maximal consolidation; that is, $\mathcal{C}_p$ requires additional selection criteria to be a well-defined function from bases to bases. Nevertheless, minimizing these distances (d_p) leads to a set D_p of preferred, closest maximal consolidations, inducing a selection function γ. To perform a consolidation, we can employ these distances to construct selection functions γ for the partial meet construction, which corresponds to take the minimum lower bounds of the corresponding maximal consolidations.

Definition 6.12. *Consider a canonical base $\Gamma = \Lambda_\Gamma(q)$. For any positive p in $\mathbb{N} \cup \{\infty\}$,*

$$\gamma_p(Cn_{ew}(\Gamma)\perp^\perp) = \{Cn_{ew}(\Lambda_\Gamma(q)) \in Cn_{ew}(\Gamma)\perp^\perp \\ |d_p(q, q') \text{ is minimum}\} .$$

Example 6.13 (Rain City Revisited II)**.** Consider that we want to consolidate the canonical base $\Gamma = \{P(\varphi) \leq 0.3, P(\varphi) \geq 0.6\}$ through a liftable partial meet contraction by the contradiction, returning a canonical base (in practice, we just want to change the probability bounds). Firstly, we compute the closure of Γ under Cn_{ew}:

$$Cn_{ew}(\Gamma) = Cn_{ew}(\{P(\varphi) \leq 0.3\}) \cup Cn_{ew}(\{P(\varphi) \geq 0.6\}) .$$

Then the remainder set $Cn_{ew}(\Gamma)\perp^\perp$ can be defined through the (Cn_{ew})-closure of the maximal consolidations:

$$Cn_{ew}(\Gamma)\perp^\perp = \{Cn_{ew}(\{P(\varphi) \leq q_1, P(\varphi) \geq q_2\}) \\ |q_1 = q_2, q_1 \in [0.3, 0.6]\} .$$

Now we need a selection function γ. Consider for instance the selection function γ_p based on distance minimization, for $p = 1$, $p = 2$ and $p = \infty$:

$$\gamma_1(Cn_{ew}(\Gamma)\perp^\perp) = Cn_{ew}(\Gamma)\perp^\perp ,$$

$$\gamma_2(Cn_{ew}(\Gamma)\perp^\perp) = Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\}) ,$$

$$\gamma_\infty(Cn_{ew}(\Gamma)\perp^\perp) = Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\}) .$$

When $p = 1$, any $q_1, q_2 \in [0.3, 0.6]$ with $q_1 = q_2$ is such that $d_1(\langle q_1, q_2\rangle, \langle 0.3, 0.6\rangle) = |q_1 - 0.3| + |q_2 - 0.6| = 0.3$ [6]. Hence, γ_1 selects all elements in the remainder set. Nonetheless, $d_2(\langle q_1, q_2\rangle, \langle 0.3, 0.6\rangle)$ has a single minimum with $q_1 = q_2$ at $q_1 = q_2 = 0.45$, and γ_2 selects only the (Cn_{ew}) closure of the corresponding maximal consolidation. The same happens to γ_∞.

To consolidate $Cn_{ew}(\Gamma)$, we take the intersection of the selected bases for each selection function:

$$\begin{aligned} Cn_{ew}(\Gamma)!^1_\star &= \bigcap \gamma_1(Cn_{ew}(\Gamma)\perp^\perp) = \bigcap Cn_{ew}(\Gamma)\perp^\perp \\ &= Cn_{ew}(\{P(\varphi) \leq 0.6, P(\varphi) \geq 0.3\}) , \\ Cn_{ew}(\Gamma)!^2_\star &= \bigcap \gamma_2(Cn_{ew}(\Gamma)\perp^\perp) \\ &= \bigcap Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\}) \\ &= Cn_{ew}(\{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\}) , \\ Cn_{ew}(\Gamma)!^\infty_\star &= Cn_{ew}(\Gamma)!^2_\star . \end{aligned}$$

[6]Technically, $P(\varphi) \leq 0.3$ and $P(\varphi) \leq q_1$ abbreviate $P(\neg\varphi) \geq 1 - 0.3$ and $P(\neg\varphi) \geq 1 - q_1$, but note that $|1 - q_1 - (1 - 0.3)| = |0.3 - q_1| = |q_1 - 0.3|$.

Note that, in this particular case, $Cn_{ew}(\Gamma)!^2_\star$ and $Cn_{ew}(\Gamma)!^\infty_\star$ are maxichoice consolidations, while $Cn_{ew}(\Gamma)!^1_\star$ is a full meet consolidation.

As $Cn_{ew}(\Gamma)!^p_\star$, for any $p \in \{1, 2, \infty\}$, is a partial meet consolidation, any $\Gamma!^p$ such that $Cn_{ew}(\Gamma!^p) = Cn_{ew}(\Gamma)!^p_\star$ will be a liftable partial meet consolidation. To satisfy ($f_\star$-canonicity), we take the canonical form:

$$\Gamma!^1 = Cn_{ew}^{-1}(\Gamma!^1_\star) = \{P(\varphi) \leq 0.6, P(\varphi) \geq 0.3\}$$

$$\Gamma!^2 = Cn_{ew}^{-1}(\Gamma!^2_\star) = \{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\}$$

$$\Gamma!^\infty = Cn_{ew}^{-1}(\Gamma!^\infty_\star) = \{P(\varphi) \leq 0.45, P(\varphi) \geq 0.45\}$$

Due to their starred closures, $\Gamma!^2$ and $\Gamma!^\infty$ are liftable maxichoice consolidations, whereas $\Gamma!^1$ is a liftable full meet consolidation.

Recall that $P(\varphi) \leq q_1$ abbreviates $P(\neg\varphi) \geq 1 - q_1$. Thus, $\Gamma!^1$ is taking the minimum lower bounds from the set of maximal consolidations: $1 - q_1 = 1 - 0.6 = 0.4$ and $q_2 = 0.3$. The operations $\Gamma!^2$ and $\Gamma!^\infty$ return a single preferred maximal consolidation. In the end, by consolidating the closure and taking the canonical form, we are just changing the probability bounds.

Once we have fully characterized the consolidation of canonical probabilistic knowledge bases, by employing — and extending — the AGM approach, a compelling question is how to generalize these results to probabilistic contraction, if possible, which we tackle in the next section.

7 Towards Probabilistic Contraction

Using the upper bound property for probabilistic consolidation (Lemma 5.4), we derived results to characterize (liftable) contractions of canonical probabilistic bases by the contradiction ($\Gamma - \bot$) through their starred closure ($Cn_{ew}(\Gamma) -_\star \bot$), but not by arbitrary formulas $\alpha \in \mathcal{L}^P_{X_n}$. In fact, contracting a probabilistic knowledge base $\Gamma = Cn_{ew}(\Gamma)$ by an arbitrary conditional $\alpha \in \mathcal{L}^P_{X_n}$ via partial meet is not possible, due to the following result:

Proposition 7.1. *Cn_{Pr} does not satisfy the upper bound property.*

As a consequence, there are bases $\Gamma \in \mathbb{K}$ and conditionals α such that the partial meet contraction $\Gamma - \alpha$ is ill-defined. To understand how probabilistic consolidation works fine via partial meet, while contraction in general does not, we extend the probabilistic logic. Let $\bar{\mathcal{L}}^P_{X_n}$ be the language containing the *negated conditionals*: $\bar{\mathcal{L}}^P_{X_n} = \{\neg\alpha | \alpha \in \mathcal{L}^P_{X_n}\}$. The semantics for the language $\bar{\mathcal{L}}^P_{X_n} \cup \mathcal{L}^P_{X_n}$ is given by the semantics for $\mathcal{L}^P_{X_n}$ with an extra rule for negation: a probability interpretation π satisfies $\neg\alpha \in \bar{\mathcal{L}}^P_{X_n}$ if π does not satisfy $\alpha \in \mathcal{L}^P_{X_n}$. The consequence operation Cn_{Pr} can be extended accordingly, and Lemma 5.4, generalized:

Theorem 7.2 (Upper bound)**.** *Let $\Gamma = Cn_{ew}(\Gamma)$ be a probabilistic knowledge base in $\mathbb{K}$, and $\alpha \in \bar{\mathcal{L}}^P_{X_n}$, a negated conditional. For every $\Delta \subseteq \Gamma$ such that $\alpha \notin Cn_{Pr}(\Delta)$, there is a $\Psi \in (\Gamma \bot \alpha)$ such that $\Delta \subseteq \Psi$.*

From which follows that the scope of the well-behaved probabilistic contraction is $\bar{\mathcal{L}}^P_{X_n}$ [7]. With this in mind, we relax in this section the definition of a contraction operator $-$ to allow a domain that is strictly smaller than the language, $\bar{\mathcal{L}}^P_{X_n} \subsetneq \bar{\mathcal{L}}^P_{X_n} \cup \mathcal{L}^P_{X_n}$, in order to obtain representation results:

Corollary 7.3. *The operator $- : \bar{\mathcal{L}}^P_{X_n} \to \mathbb{K}$ for a $\Gamma = Cn_{ew}(\Gamma)$ in $\mathbb{K}_c$ satisfies (Success), (Inclusion), (Relevance) and (Uniformity) iff $\Gamma - \alpha = \bigcap \gamma(\Gamma \bot \alpha)$ for all $\alpha \in \bar{\mathcal{L}}^P_{X_n}$ and for some selection function γ.*

To see that this result holds for consolidation, just take $\alpha = P(\bot) > 0 \in \bar{\mathcal{L}}^P_{X_n}$. That is, we can contract by $\bot = P(\bot) \geq 1$ just because it is equivalent to $\bot' = P(\bot) > 0$. However, it is not the case that any $P(\varphi|\psi) \geq q$ has an equivalent $P(\varphi'|\psi') > q'$.

Since the contraction of a probabilistic base $\Gamma = Cn_{ew}(\Gamma)$ by negated conditionals is well-defined, so is the corresponding liftable contraction. Consequently, Propositions 6.5 and 6.7 also hold in this case.

The fact that only contractions by negated conditionals can be well characterized by the postulates and the partial meet construction within the probabilistic logic may be somewhat strange at first, but thinking about belief revision, such oddness disappears. To revise a probabilistic knowledge base by a conditional $\alpha = P(\varphi|\psi) \geq q$, one can apply these two operations, in some order: contract by $\neg\alpha$; expand by α. And as contractions by $\neg\alpha$ are indeed well-founded on the AGM paradigm, one can define revision as well. For instance, using Levi's identity, to revise a canonical base Γ by a conditional $\alpha = P(\varphi|\psi) \geq q$, one could perform a liftable contraction by $\neg\alpha$, which is a well-defined operation, followed by consistently adding α to the base.

8 Conclusion and Future Work

We investigated consolidation procedures for probabilistic knowledge bases under the AGM framework of belief revision. To characterize consolidation via probability changing, we defined the concept of liftable contraction, which is parametrized by a consequence operation. Using this consequence operation, we lifted the postulates for base contraction operations in order that probabilistic consolidations may be fully characterized through a representation theorem. In the end, we devised a probabilistic contraction operation for a specific language, since in the general case such operation is ill-defined.

Future work includes the investigation of liftable revisions via liftable contractions. To achieve that, one might need to define an expansion operation parametrized by a consequence operation as well. Such revision via liftable contractions could then be characterized by a set of lifted postulates. The (liftable) probabilistic contraction we introduced could be explored in detail to ground a liftable revision operation for probabilistic knowledge bases.

[7] Note that this contraction is regarding the standard semantics (Cn_{Pr}), and does not apply to the maximum entropy entailment.

Acknowledgments

GDB was supported by CAPES. MF is partially supported by CNPq grant PQ 306582/2014-7 and Fapesp project 2015/21880-4. RW is partially supported by CNPq grant PQ 309605/2013-0.

References

Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic* 50(02):510–530.

Boole, G. 1854. *An Investigation of the Laws of Thought: on which are Founded the Mathematical Theories of Logic and Probabilities*. Walton and Maberly.

De Bona, G., and Finger, M. 2015. Measuring inconsistency in probabilistic logic: rationality postulates and dutch book interpretation. *Artificial Intelligence* 227:140–164.

de Finetti, B. 1930. Problemi determinati e indeterminati nel calcolo delle probabilità. *Rendiconti Reale Accademia dei Lincei* 6:367–373.

Frisch, A., and Haddawy, P. 1994. Anytime deduction for probabilistic logic. *Artificial Intelligence* 69(1):93–122.

Haenni, R.; Romeijn, J.; Wheeler, G.; and Williamson, J. 2011. Probabilistic logic and probabilistic networks.

Hailperin, T. 1965. Best possible inequalities for the probability of a logical function of events. *American Mathematical Monthly* 72(4):343–359.

Halpern, J. 1990. An analysis of first-order logics of probability. *Artificial Intelligence* 46(3):311–350.

Hansen, P., and Jaumard, B. 2000. Probabilistic satisfiability. *Handbook of Defeasible Reasoning and Uncertainty Management Systems: Algorithms for uncertainty and defeasible reasoning* 321.

Hansen, P.; Jaumard, B.; De Aragao, M. P.; Chauny, F.; and Perron, S. 2000. Probabilistic satisfiability with imprecise probabilities. *International Journal of Approximate Reasoning* 24(2):171–189.

Hansson, S. 1992. A dyadic representation of belief. *Belief revision* 29:89–121.

Hansson, S. 1997. Semi-revision. *Journal of Applied Non-Classical Logics* 7(1-2):151–175.

Hansson, S. 1999. *A textbook of belief dynamics*, volume 1. Springer.

Jaumard, B.; Hansen, P.; and Poggi de Aragao, M. 1991. Column generation methods for probabilistic logic. *INFORMS Journal on Computing* 3(2):135.

Kern-Isberner, G., and Lukasiewicz, T. 2004. Combining probabilistic logic programming with the power of maximum entropy. *Artificial Intelligence* 157(1):139–202.

Lukasiewicz, T. 1999. Probabilistic deduction with conditional constraints over basic events. *Journal of Artificial Intelligence Research* 10:199–241.

Muiño, D. 2011. Measuring and repairing inconsistency in probabilistic knowledge bases. *International Journal of Approximate Reasoning* 52(6):828–840.

Nilsson, N. 1986. Probabilistic logic* 1. *Artificial Intelligence* 28(1):71–87.

Potyka, N., and Thimm, M. 2014. Consolidation of probabilistic knowledge bases by inconsistency minimization. In *ECAI 2014*.

Rödder, W. 2000. Conditional logic and the principle of entropy. *Artificial Intelligence* 117(1):83–106.

Santos, Y.; Ribeiro, M.; and Wassermann, R. 2015. Between Belief Bases and Belief Sets: Partial Meet Contraction. In *Proceedings of the International Workshop on Defeasible and Ampliative Reasoning (DARe-15)*.

Thimm, M. 2011. Probabilistic reasoning with incomplete and inconsistent beliefs, volume 331 of dissertations in artificial intelligence.

Thimm, M. 2013. Inconsistency measures for probabilistic logics. *Artificial Intelligence* 197:1–24.

Wassermann, R. 2000. *Resource-Bounded Belief Revision*. Ph.D. Dissertation, Institute for Logic, Language and Computation — University of Amsterdam.