
Belief change based on knowledge measures

GIOVANNI CASINI, *Istituto di Scienza e Tecnologie dell'Informazione, CNR, 56100 Pisa, Italy.*

CAIR, *University of Cape Town, 7700 Cape Town, South Africa.*

E-mail: giovanni.casini@isti.cnr.it

UMBERTO STRACCIA *Istituto di Scienza e Tecnologie dell'Informazione, CNR, 56100 Pisa, Italy.*

E-mail: umberto.straccia@isti.cnr.it

Abstract

Belief change (BC) is the process of changing beliefs (in our case, in terms of contraction, expansion and revision) taking into account a new piece of knowledge, which possibly may be in contradiction with the current belief. In this work, we present a novel quantitative BC framework based on the principle of minimizing the surprise from an information-theoretic perspective. Central to our approach is the *Principle of Minimal Surprise (PMS)*, which asserts that when confronted with the uncertainty about which is the actual world, an agent should tend to favour the most expected, i.e. least surprising, outcomes. To formalize this, we make use of *knowledge measures*, which quantify the amount of information contained in a knowledge base. Guided by the PMS, our framework encourages belief change operations that minimize the informational amount deviation from the original belief, i.e. those that introduce the least surprise.

Keywords: belief change, knowledge measures

1 Introduction

An important topic in knowledge representation and reasoning, more generally in artificial intelligence, epistemology and cognitive science, deals with the formal definition of how rational agents should adapt their knowledge or beliefs while facing new pieces of information, possibly in contradiction with their current beliefs, and aiming at preserving the overall logical consistency of their beliefs. This process and resulting research area is called *Belief Change (BC)* [21, 29].¹ Main ingredients of BC concern how a knowledge-holding agent should incorporate new information, resolve contradictions and update its world-view consistently according to some criteria. Among typical BC operations we may mention (but, see [22] and references therein for others) the following:

- (1) *contraction*: deals with how to remove a belief without replacing it with a new one, often used as a step in revision;
- (2) *belief revision*: deals with how one needs possibly to give up to some current beliefs to accommodate a new one;
- (3) *belief update*: deals with the case in which the world itself may change, and how one has to adapt the beliefs accordingly;
- (4) *belief merging (fusion)*: deals with the combination of multiple, possibly contradictory sources of information, into a coherent whole.

¹ See also <https://plato.stanford.edu/archives/spr2022/entries/logic-belief-revision/>

2 Belief change based on knowledge measures

Various axiomatizations have been defined that are believed such BC operations should satisfy and, to date, a vast multitude of different ways for performing these operations have been proposed (see, e.g. [21]).

In this work, we will focus on the contraction and belief revision operations within the well-known AGM approach [2, 20], and, in line with many approaches connected to the AGM framework [46], we model the changes in the agent's beliefs through operations on a possible-worlds semantics.

In this setting, we are going to present a novel quantitative BC framework based on the principle of minimizing the surprise from an information-theoretic perspective.

The main ingredient of our approach is the *Principle of Minimal Surprise* (PMS), which asserts that when confronted with the uncertainty about which is the actual world, an agent should tend to favour the most expected, i.e. least surprising, outcomes. To formalize this, we make use of *Knowledge Measures* (KMs) [67], which quantify the amount of information contained in a knowledge base. Guided by the PMS, our framework encourages BC operations that minimize the informational amount deviation from the original belief, i.e. those that introduce the least surprise.

Roughly, the basic idea is as follows. A KB, the agent's beliefs, is seen as a formal, possibly incomplete, specification of the actual world. The more models a KB has, the more uncertain the agent is about which is the actual world, which in turn means the less he knows about the actual world. However, the agent relies on a probability distribution over the worlds and the higher is the probability of a world, the more likely it is the actual world.² This will allow us, via *Shannon's Information Theory* (see, e.g. [34]), to quantify the amount of information, i.e. knowledge, an agent has about the actual world. This measure is called KM (see [67]). Now, according to the BC operations we suggest in this work, an agent should conform to the PMS and, thus, should tend to favour the most expected, i.e. or least surprising results, which in turn means the agent should try to minimize the amount of knowledge deviation from the original amount.

Contribution Our contributions can be summarized as follows:

- we generalize KMs introduced in [67], which were originally defined under the assumption of a uniform probability distribution, in a way that prepares the ground for their application to BC;
- we then propose KM-based BC operators that satisfy the AGM postulates;
- we also show that any BC operator that satisfies the AGM postulates can be represented as a KM-based BC operator;
- we introduce quantitative measures that account for the information change of BC operators. This allows us also to quantify precisely, e.g. how much information/knowledge is lost during a contraction and to compare it with other BC operators;
- we also start investigating *iterated belief revision* (IBR) [16] and show that our revision process satisfies three of the four classical postulates of IBR [16]; and
- we illustrate how one may build from our KM-based contraction operator also one not satisfying the (in)famous *recovery postulate*, by focusing on the so-called *severe withdrawal* [61] model as illustrative example.

Although the area of BC has been extensively explored by many researchers (see, e.g. [3, 20, 21, 55] for overviews), several works have also examined potential connections between BC and probability theory (see, e.g. [5, 35–38, 51, 53, 58, 71]). However, these probabilistic approaches generally adopt a reasoning framework that is quite different from ours. Specifically, they focus

²In the worst case, every world is equiprobable, see [67].

TABLE 1 Symbols of Section 2.

Symbol	Formal meaning	Informal meaning
Σ	$\{p, q, \dots\}$	Propositional alphabet
Σ_ϕ		Set of propositional letters occurring in ϕ
\mathcal{L}	$\{\alpha, \beta, \dots\}$	Language of propositional formulae
L	-	Literal
\bar{L}	-	Negation of literal L
\mathcal{K}	$\{\phi_1, \dots, \phi_n\}$	Knowledge Base (KB)
$\bigwedge \mathcal{K}$	$\bigwedge_{\phi \in \mathcal{K}} \phi$	The conjunction of formulae in a KB
w	Set/concatenation of literals s.t. all propositional letters in Σ occur exactly once	Interpretation or world
\mathcal{W}_Σ	-	Set of all worlds w.r.t. Σ
$w \Vdash \phi$	-	w is a model of/satisfies ϕ
\perp	-	Formula without models
\top	-	Formula satisfied by all worlds
$[\phi]$	$\{w \in \mathcal{W} \mid w \Vdash \phi\}$	Set of all models of ϕ
$\phi \vDash \psi$	-	ψ is a logical consequence of ϕ
$\phi \equiv \psi$	$\phi \vDash \psi$ and $\psi \vDash \phi$	ϕ and ψ are logically equivalent
$\phi \leq \psi$	$\phi \vDash \psi$	ψ is a logical consequence of ϕ
$\phi < \psi$	$\phi \vDash \psi$ and $\psi \not\vDash \phi$	ψ is a logical consequence of ϕ , but not vice-versa
Fw	$\bigwedge_{L \in w} L$	Formula whose unique model is w
$F\mathcal{W}$	$\bigvee_{w \in \mathcal{W}} Fw$	Formula whose models are exactly those in $\mathcal{W} \subseteq \mathcal{W}_\Sigma$

typically on how to revise probability distributions over possible worlds, representing an agent's beliefs, when a BC operation occurs. In contrast, following a *frequentist* interpretation of probability [32] (see also Remarks 1 and 2 later on), our approach assumes a fixed probability distribution that remains unchanged during BC, which is used to guide the revision by minimizing informational deviation from the original belief according to the PMS.

Other model-based BC approaches do incorporate a notion of proximity between worlds (see, e.g. [15, 42, 62, 63]), but these typically rely on syntactic measures, such as counting the number of differing literals between worlds, which fundamentally differs from our approach. Typically, their underlying idea is that, when e.g. revising a knowledge base, an agent should tend to minimize the truth changes. In contrast, in our case e.g. revision favours less surprising situations, where surprise is quantified using Shannon's information theory. That is, instead of minimizing syntactic differences, we minimize informational content deviation.

In summary, to best of our knowledge, we believe that our approach is unique and may offer a novel perspective on how BC may be approached.

In the following, we will proceed as follows. In the next section, we introduce the background notions we will rely on. In Section 3, we will recap KMs and propose a more general information-theoretic version for their application to BC. In Section 4, we will show how one may formulate BC using KMs, illustrate succinctly IBR and severe withdrawal within our framework. For each of these sections, to facilitate the reading, we provide a symbol table (see Tables 1, 2 and 3). In Section 5, we

4 Belief change based on knowledge measures

TABLE 2 Symbols of Section 3.

Symbol	Formal meaning	Informal meaning
\mathbb{P}	-	Probability distribution over worlds. The higher is $\mathbb{P}(w)$, the more likely w is the actual world
\mathbb{P}^u	-	Uniform probability distribution
w^p	$\mathbb{P}(w) = p$	The world w has probability p
$\mathbb{P}(\phi)$	$\sum_{w \in [\phi]} \mathbb{P}(w)$	The probability of ϕ
ϕ^p	$\mathbb{P}(\phi) = p$	ϕ has probability p
$[\phi]^+$	$\{w \mid w \models \phi \text{ and } \mathbb{P}(w) > 0\}$	The set of possible/non-zero probability models of ϕ
$\mathbb{P}0$	$F \{w \mid \mathbb{P}(w) = 0\}$	The formula whose models are impossible/have zero probability
$\phi \leq_{\mathbb{P}} \psi$	$\phi \wedge \neg \mathbb{P}0 \models \psi$	ϕ \mathbb{P} -entails ψ , i.e. all non-zero \mathbb{P} -probability models of ϕ are ψ models, i.e. $[\phi]^+ \subseteq [\psi]$
$\phi <_{\mathbb{P}} \psi$	$\phi \leq_{\mathbb{P}} \psi$ and $\psi \not\leq_{\mathbb{P}} \phi$	ϕ \mathbb{P} -entails ψ , but not vice-versa
$\phi <_{\mathbb{P}} \psi$	$<_{\mathbb{P}} \in \{\leq_{\mathbb{P}}, <_{\mathbb{P}}\}$	-
$\phi \equiv_{\mathbb{P}} \psi$	$\phi \leq_{\mathbb{P}} \psi$ and $\psi \leq_{\mathbb{P}} \phi$	ϕ and ψ are \mathbb{P} -equivalent
$\kappa_S(\phi)$	$-\log_2 \mathbb{P}(\phi)$	Shannon KM of ϕ , i.e. amount of knowledge/information encoded by ϕ .
$\kappa(\phi)$	$-\log_b \mathbb{P}(\phi) = \frac{1}{\log_2 b} \cdot \kappa_S(\phi)$	Amount of knowledge/information encoded by ϕ , according to any KM

address related work and, eventually, Section 6 summarizes our contribution and proposes topics for future work. Some proofs can be found in Appendix A.

2 Background

Let $\Sigma = \{p, q, \dots\}$ be a finite non-empty set of propositional letters (all symbols may have an optional sup- or sub-script or apex). $\mathcal{L} = \{\alpha, \beta, \dots\}$ is the set of propositional formulae based on the set of Boolean operators $\{\wedge, \vee, \neg, \rightarrow, \leftrightarrow\}$ and Σ . A *literal* (denoted L) is either a propositional letter or its negation and with \bar{L} we denote $\neg p$ (resp. p) if $L = p$ (resp. if $L = \neg p$). A *knowledge base* (KB) $\mathcal{K} = \{\phi_1, \dots, \phi_n\}$ is a finite set of formulae ϕ_i and with $\bigwedge \mathcal{K}$ we denote the conjunction of the formulae in \mathcal{K} , i.e. $\bigwedge_{\phi \in \mathcal{K}} \phi$. In the following, whenever we write \mathcal{K} , we consider $\bigwedge \mathcal{K}$ instead, unless stated otherwise. Given a formula ϕ , with $\Sigma_\phi \subseteq \Sigma$ we denote the set of propositional letters occurring in ϕ (we may omit the reference to Σ if no ambiguity arises).

An *interpretation*, or *world*, w w.r.t. Σ is a set of literals such that all propositional letters in Σ occur exactly once in w . \mathcal{W}_Σ is the set of all worlds w.r.t. Σ . We may also denote a world w as the concatenation of the literals occurring in w by replacing a negative literal $\neg p$ with \bar{p} (e.g. the world $\{\neg p, q\}$ may be denoted also as $\bar{p}q$). With $w \models \phi$ we indicate that the world w *satisfies* the formula ϕ , i.e. w is a *model* of ϕ , which is inductively defined as usual: $w \models L$ iff $L \in w$, $w \models \neg \phi$ iff $w \not\models \phi$, $w \models \phi \wedge \psi$ iff $w \models \phi$ and $w \models \psi$, $w \models \phi \vee \psi$ iff $w \models \phi$ or $w \models \psi$, $w \models \phi \rightarrow \psi$ iff $w \models \neg \phi \vee \psi$, and eventually $w \models \phi \leftrightarrow \psi$ iff $w \models (\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$. Furthermore, ϕ is *satisfiable* (resp. *unsatisfiable*) if it has (resp. has no) model. We will also use two special symbols \perp and \top : \perp is the formula without models, while every world satisfies \top . We impose that \top (resp. \perp) cannot occur

TABLE 3 Symbols of Section 4.

Symbol	Formal meaning	Informal meaning
$\phi \perp\!\!\!\perp \alpha$	$\{F([\phi]^+)\}$ if $\top \leq_{\mathbb{P}} \alpha$, else $\{F([\psi]^+) \mid \phi \leq_{\mathbb{P}} \psi, \psi \not\leq_{\mathbb{P}} \alpha\}$	The possible remainders of ϕ w.r.t. α
$\phi \perp \alpha$	$\{\psi \in \phi \perp\!\!\!\perp \alpha \mid \text{there is no } \psi' \in \phi \perp\!\!\!\perp \alpha \text{ s.t. } \psi' <_{\mathbb{P}} \psi\}$	The set of remainders of ϕ w.r.t. α .
$\phi_{\alpha}^{\dot{\leftarrow} \kappa}$	$\bigvee \min_{\kappa}(\phi \perp \alpha)$	KM-contraction: ϕ is km-contracted by α $(\min_{\kappa}(\Gamma) = \{\phi \in \Gamma \mid \forall \psi \in \Gamma, \kappa(\phi) \leq \kappa(\psi)\})$
$\phi_{\alpha}^{\dot{\leftarrow}}$	$\bigvee \gamma(\phi \perp \alpha)$	Partial meet contraction: ϕ is contracted by α (γ is a choice function returning a set of formulae), where the uniform distribution \mathbb{P}^u is assumed pre-order on set of interpretations \mathcal{W}_{Σ}
\leq_{ϕ}	-	
$r(w)$	$\begin{cases} 0, & \text{if } w \in \min_{\leq_{\phi}}(\mathcal{W}_{\Sigma}) \\ i, & \text{if } w \in \min_{\leq_{\phi}}(\mathcal{W}_{\Sigma} \setminus \mathcal{W}_{\Sigma}^i) \end{cases}$	ϕ -faithful ranking function, where $\mathcal{W}_{\Sigma}^i = \{w' \mid r(w') < i\}$
σ	-	A sphere, i.e. set of possible worlds
\mathbb{S}_{ϕ}	-	System of spheres centred on ϕ
$\sigma(\neg\alpha)$	$[\phi]^+ \cup \{w \in [\top]^+ \setminus [\phi]^+ \mid \mathbb{P}(w) \geq p_{-\alpha}^{\max}\}$	Smallest sphere intersecting $[\neg\alpha]^+$ $(p_{\beta}^{\max} = \max\{\mathbb{P}(w) \mid w \in [\beta]^+\})$
$\phi_{\alpha}^{\ddot{\leftarrow}}$	$\begin{cases} F\sigma(\neg\alpha), & \text{if } \top \not\leq_{\mathbb{P}} \alpha \\ F[\phi]^+ & \text{otherwise} \end{cases}$	KM-severe withdrawal: a contraction without recovery
$\phi + \alpha$	$\phi \wedge \alpha$	Expansion: α is added to ϕ
$\phi_{\alpha}^{\star \kappa}$	$(\bigvee \min_{\kappa}(\phi \perp \neg\alpha)) \wedge \alpha$	KM-revision: ϕ is km-revised by α (using Levi identity $\phi_{\alpha}^{\star} = (\phi \div \neg\alpha) + \alpha$)
$IC^{\dot{\leftarrow}}(\phi, \alpha)$	$\kappa(\phi_{\alpha}^{\dot{\leftarrow}}) - \kappa(\phi)$	Information change of a contraction \div
$IC^+(\phi, \alpha)$	$\kappa(\phi + \alpha) - \kappa(\phi)$	Information change of expansion
$IC^{\star}(\phi, \alpha)$	$\kappa(\phi \star \alpha) - \kappa(\phi)$	Information change of revision

in any other formula different than \top (resp. \perp) itself. Given formulae ϕ and ψ , let $[\phi]$ be the set of models of ϕ w.r.t. Σ ; define $\phi \models \psi$ iff $[\phi] \subseteq [\psi]$, and ϕ and ψ *equivalent*, denoted $\phi \equiv \psi$, iff $[\phi] = [\psi]$.

In the following, w.l.o.g., for any set of formulae $\mathcal{S} = \{\phi_1, \phi_2, \dots\}$, we always assume that $\phi_i \not\equiv \phi_j$ for each $i \neq j$, and, thus, $|\mathcal{S}| \leq 2^{|\Sigma|}$.

Moreover, let $\phi \leq \psi$ iff $\phi \models \psi$; and $\phi < \psi$ iff $\phi \models \psi$ and $\psi \not\models \phi$. For $w \in \mathcal{W}_{\Sigma}$, let

$$Fw = \bigwedge_{L \in w} L, \quad (1)$$

i.e. Fw is the formula whose unique model is w . For $\mathcal{W} \subseteq \mathcal{W}_{\Sigma}$, let

$$F\mathcal{W} = \bigvee_{w \in \mathcal{W}} Fw, \quad (2)$$

6 Belief change based on knowledge measures

i.e. $F\mathcal{W}$ is the formula whose models are exactly those in \mathcal{W} . We also impose that $F\emptyset = \perp$ and $F\mathcal{W}_\Sigma = \top$. Obviously,

$$\phi \equiv F[\phi] \quad (3)$$

holds.³

EXAMPLE 1. (Running example)

Consider the following KB about ‘birds’ over $\Sigma = \{b, p, o, f, w\}$, where b, p, o, f and w stand for ‘bird’, ‘penguin’, ‘ostrich’, ‘flies’ and ‘wings’, respectively. The KB

$$\mathcal{K} = \{b, b \rightarrow f, p \rightarrow b, o \rightarrow b, \neg(p \wedge o)\}$$

states that we believe that there are birds and that they fly. Also, a penguin is a bird, and an ostrich is a bird, but a penguin cannot be an ostrich and vice versa. The set of models $[\mathcal{K}]$ of \mathcal{K} over Σ is

$$[\mathcal{K}] = \{b\bar{p}\bar{o}f\bar{w}, b\bar{p}of\bar{w}, b\bar{p}\bar{o}f\bar{w}, bp\bar{o}f\bar{w}, b\bar{p}of\bar{w}, b\bar{p}\bar{o}f\bar{w}\} \quad (4)$$

and, in absence of further information, any model w of \mathcal{K} may indeed be the actual world. Furthermore, we believe that it is reasonable to assume that, for any strict non-empty subset $\mathcal{S} \subset [\mathcal{K}]$ (resp. strict superset $\mathcal{S} \supset [\mathcal{K}]$), the formula $F\mathcal{S}$ is less (resp. more) precise in describing who the actual world is w.r.t. \mathcal{K} , i.e. encodes less (resp. more) ‘information’ about who the actual world is than \mathcal{K} . Please note that in the former case $\mathcal{K} \models F\mathcal{S}$ holds, while in the latter case we have $F\mathcal{S} \models \mathcal{K}$.

3 An information-theoretic approach to KMs

In this section we generalize KMs introduced in [67] in a way to be applied to BC. The generalization is inspired on Shannon’s *information theory* (see, e.g. [34]). Roughly, the simple principles we will rely on are the following:⁴

- (1) a formula ϕ is considered as a formal, possibly incomplete, specification of the actual world, which however has to be one of the models of ϕ , without necessarily knowing which it is;
- (2) the more models a formula has, the more uncertain we are about which is the actual world, which in turn means the less we know about the actual world;
- (3) additionally, we assume that a fixed probability distribution \mathbb{P} over the worlds is given. The higher is the probability $\mathbb{P}(w)$ of a world w , the more likely w is the actual world.

Let us recall that the core idea in Shannon’s *information theory* is that the ‘informational value’ of a communicated message (in our case a formula) depends on the degree to which the content of the message (in our case, the models and their probability of a formula) is *surprising*. If a highly likely event occurs, i.e. the probability of a formula is high the surprise is low and, thus, the message carries little information. On the other hand, if a highly unlikely event occurs, the surprise is high and the message is much more informative.

So, assume that a probability distribution $\mathbb{P}: \mathcal{W}_\Sigma \rightarrow [0, 1]$ over the worlds in \mathcal{W}_Σ is given. The intended meaning of $\mathbb{P}(w)$ is to indicate how probable it is that w is indeed the actual world.

³Note that $F[\phi]_\Sigma$ is the formula whose models are exactly those of ϕ .

⁴We refer the reader to Appendix C for a succinct recap on KMs as introduced in [67]. Also note that in [67] each world is equally plausible, i.e. a uniform probability distribution across all possible worlds is implicitly assumed.

REMARK 1.

As mentioned in Section 1, we consider a *frequentist* interpretation of probability and not an *epistemic/Bayesian/subjective* one (see, e.g. [32]). That is, probability values are based on observed frequencies in data or experiments (e.g. 90% of birds fly) and are interpreted as describing objective randomness inherent in the world. So, \mathbb{P} is a fixed probability distribution that remains unchanged by a BC operation.

We say that a world w is *possible*, if it has non-zero probability, i.e. $\mathbb{P}(w) > 0$. For ease of presentation, if no ambiguity arises, we may also denote $\mathbb{P}(w) = n$ as w^n . So, we may also write \mathbb{P} as a set $\{w_1^{p_1}, \dots, w_k^{p_k}\}$, with $\mathbb{P}(w_i) = p_i$, and the convention that worlds not occurring in that set have 0 probability. The *uniform probability distribution* is, of course, defined as $\mathbb{P}^u(w) = 1/|\mathcal{W}_\Sigma|$. The probability that one of the models of ϕ is indeed the actual world, called the *probability of a formula* ϕ , denoted $\mathbb{P}(\phi)$, is defined as

$$\mathbb{P}(\phi) = \sum_{w \in [\phi]} \mathbb{P}(w). \quad (5)$$

Note that $\mathbb{P}(\top) = 1$, $\mathbb{P}(\perp) = 0$ and, if \mathbb{P} is the uniform probability distribution, then $\mathbb{P}(\phi) = |[\phi]|/|\mathcal{W}_\Sigma|$. We may also simply write ϕ^p in place of $\mathbb{P}(\phi) = p$ if clear from the context. With $[\phi]^+$ we denote the *set of possible models* of ϕ , i.e. the set of models of ϕ that have non-zero probability: i.e.

$$[\phi]^+ = \{w \in \mathcal{W}_\Sigma \mid w \models \phi \text{ and } \mathbb{P}(w) > 0\}. \quad (6)$$

REMARK 2.

Please note that, in accordance with Remark 1, two formulae ϕ and ψ , aimed at describing information about the actual world (i.e. one of the worlds in \mathcal{W}_Σ), share the *same* probability distribution \mathbb{P} over worlds. The latter will induce a probability $\mathbb{P}(\phi)$ (resp. $\mathbb{P}(\psi)$) indicating how likely one of the models of ϕ (resp. ψ) is indeed the actual world. This will then allow us to compare the amount of information encoded by ϕ and ψ on equal terms, i.e. based on the same probability distribution.

Note also that a formula ϕ does not impose restrictions on \mathbb{P} , but we only expect that ϕ is ‘compatible’ with the scenario depicted by \mathbb{P} , i.e. one of the models of ϕ has a non-zero probability. Moreover, what is considered an impossible world according to \mathbb{P} has to be false (see \mathbb{P} -entailment later on).

This is different from numerous probabilistic logics proposed in the literature (see e.g. [33]) in which typically a probabilistic KB is a formal specification about probabilities and/or subjective probabilities of worlds. We do not consider this scenario here and leave it for future work.

EXAMPLE 2. (Running example cont.)

Consider Example 1. Let us assume that we have the following probability distribution \mathbb{P} w.r.t. $\Sigma = \{b, p, o, f, w\}$:⁵

$$\mathbb{P} =_{\text{def}} \{b\bar{p}\bar{o}f w^{0.1}, b\bar{p}o\bar{f} w^{0.1}, b\bar{p}\bar{o}\bar{f} w^{0.15}, b\bar{p}\bar{o}f\bar{w}^{0.15}, \\ b\bar{p}o\bar{f}\bar{w}^{0.2}, b\bar{p}\bar{o}f w^{0.2}, \bar{b}\bar{p}\bar{o}f w^{0.07}, \bar{b}\bar{p}\bar{o}f\bar{w}^{0.03}\}.$$

⁵If a world is not listed, then its probability is 0.

8 Belief change based on knowledge measures

Therefore, we have that the following probabilities of formulae:

$$b^{0.9}, p^{0.35}, o^{0.3}, f^{0.45}, w^{1.0}, (p \wedge o)^{0.0}, \\ (b \rightarrow f)^{0.45}, (p \rightarrow b)^{0.9}, (o \rightarrow b)^{1.0}$$

and, in particular, $\mathbb{P}(\mathcal{K}) = 0.35$. That is, the probability that one of the models of \mathcal{K} is indeed the actual world is 0.35 w.r.t. \mathbb{P} .

As next, we extend the notion of entailment to the case in which we have a probability distribution over worlds. To do so, we define the following supraclassical monotonic entailment relation: consider formulae ϕ , ψ , and a probability distribution \mathbb{P} over \mathcal{W}_Σ . Let us define the formula

$$\mathbb{P}0 = F \{w \mid \mathbb{P}(w) = 0, w \in \mathcal{W}_\Sigma\} \quad (7)$$

as the disjunction of all worlds having 0 probability, i.e. the formula whose models are impossible/have zero probability. So, $\mathbb{P}(\mathbb{P}0) = 0$ and, thus, $\mathbb{P}(\neg\mathbb{P}0) = 1$.

Then the notion of ϕ \mathbb{P} -entails ψ , denoted $\phi \leq_{\mathbb{P}} \psi$ or also $\phi \models_{\mathbb{P}} \psi$, is defined as the case in which all non-zero \mathbb{P} -probability models of ϕ are ψ models, i.e.

$$\phi \leq_{\mathbb{P}} \psi \text{ iff } \phi \wedge \neg\mathbb{P}0 \models \psi. \quad (8)$$

That is, $\phi \leq_{\mathbb{P}} \psi$ iff $[\phi]^+ \subseteq [\psi]$. We also define (i) $\phi <_{\mathbb{P}} \psi$ iff $\phi \leq_{\mathbb{P}} \psi$ and $\psi \not\leq_{\mathbb{P}} \phi$; and (ii) $\phi \equiv_{\mathbb{P}} \psi$ (ϕ and ψ are \mathbb{P} -equivalent) iff $\phi \leq_{\mathbb{P}} \psi$ and $\psi \leq_{\mathbb{P}} \phi$. We say that ϕ is \mathbb{P} -consistent, or also \mathbb{P} -satisfiable iff $\phi \not\leq_{\mathbb{P}} \perp$.

REMARK 3.

Note that $\phi \wedge \neg\mathbb{P}0$ has a model only iff ϕ has a non-zero probability model, i.e. ϕ is \mathbb{P} -satisfiable. Therefore, $\phi <_{\mathbb{P}} \psi$ tell us that all non-zero \mathbb{P} -probability models of ϕ are models of ψ and that there is a non-zero \mathbb{P} -probability model of ψ that is not a model of ϕ .

REMARK 4. (Classical case)

Classical propositional logic can be obtained as a special case if, e.g. the uniform probability distribution \mathbb{P}^u is considered. In that case, $\phi \leq_{\mathbb{P}^u} \psi$ iff $\phi \models \psi$ holds.

However, this result does not generalize to any probability distribution \mathbb{P} : while $\phi \models \alpha$ implies $\phi \leq_{\mathbb{P}} \alpha$, the opposite is not true. In fact, e.g. for $\phi =_{\text{def}} a$ and $\mathbb{P}(a\bar{b}) = 0$, we have $\phi \wedge \neg\mathbb{P}0 =_{\text{def}} a \wedge \neg(a \wedge \neg b) \models b$ and thus, $\phi \leq_{\mathbb{P}} b$, but $\phi \not\models b$.

The following proposition will be used later and can easily be shown.

PROPOSITION 1.

Consider any formulae ϕ and ψ w.r.t. Σ , and any probability distribution \mathbb{P} over Σ . If $\phi <_{\mathbb{P}} \psi$ then $\mathbb{P}(\phi) < \mathbb{P}(\psi)$, for $< \in \{\leq, <\}$.

Moreover, the inverse of Proposition 1 does not hold.

EXAMPLE 3. (Running example cont.)

Consider Example 2. It is easily verified that $\mathcal{K} <_{\mathbb{P}} f \wedge (\neg b \rightarrow (p \vee o \vee f))$. Moreover, note that $\mathbb{P}(\neg b \rightarrow (p \vee o \vee f)) = 1$ as $\bar{b}\bar{p}\bar{o}\bar{f}^{0.0}$ holds. It follows that $\mathbb{P}(f \wedge (\neg b \rightarrow (p \vee o \vee f))) = 1 > 0.35 = \mathbb{P}(\mathcal{K})$ in accordance with Proposition 1.

Moreover, $\mathbb{P}(b \wedge p \wedge \neg o \wedge f) = 0.1 < 0.15 = \mathbb{P}(b \wedge \neg p \wedge \neg o \wedge f)$, but $b \wedge p \wedge \neg o \wedge f \not\prec_{\mathbb{P}} b \wedge \neg p \wedge \neg o \wedge f$ and, thus, the converse of Proposition 1 does not hold.

We conclude this part by defining the notion of independent formulae: namely, consider two formulae ϕ and ψ , and a probability distribution \mathbb{P} over Σ , then we say that ϕ and ψ are \mathbb{P} -independent iff

$$\mathbb{P}(\phi \wedge \psi) = \mathbb{P}(\phi) \cdot \mathbb{P}(\psi).$$

EXAMPLE 4. (Running example cont.)

Consider Example 2. Then b and w are \mathbb{P} -independent, while b and f are not.

We are ready now to present our generalization for KMs: given a probability distribution \mathbb{P} over Σ , a KM κ is a total function mapping formulae over Σ into $\mathbb{R}^+ \cup \{0, \infty\}$ satisfying the following three axioms:

- (KM1) $\kappa(\top) = 0$ and $\kappa(\perp) = \infty$;
- (KM2) κ is monotone non-increasing, i.e. if $\mathbb{P}(\phi) \triangleleft \mathbb{P}(\psi)$ then $\kappa(\psi) \triangleleft \kappa(\phi)$, for $\triangleleft \in \{\leq, <\}$;
- (KM3) κ is additive, i.e. if ϕ and ψ are \mathbb{P} -independent, then $\kappa(\phi \wedge \psi) = \kappa(\phi) + \kappa(\psi)$.

Let us shortly explain the above axioms.

Concerning axiom (KM1), note that if $\models \phi$ then for every formula ψ , we have that $\psi \models \phi$ and, thus, $\mathbb{P}(\phi) \geq \mathbb{P}(\psi)$. By axiom (KM2) $\kappa(\psi) \geq \kappa(\phi)$ has to hold, i.e. $\kappa(\phi)$ has to be as small as possible, which motivates $\kappa(\phi) = 0$. Analogously, if ϕ is unsatisfiable then for every formula ψ , we have that $\phi \models \psi$ and, thus, $\mathbb{P}(\phi) \leq \mathbb{P}(\psi)$. By axiom (KM2) $\kappa(\phi) \geq \kappa(\psi)$ has to hold. That is, $\kappa(\phi)$ has to be as large as possible, which motivates $\kappa(\phi) = \infty$.

Concerning axiom (KM3), it encodes a sort of model independence: in this case, we may join ϕ and ψ without any loss of information w.r.t. the sum of both.

Eventually, concerning (KM2), it states that if ψ is more probable than ϕ then ψ is less surprising than ϕ and, thus, according to information theory, the former carries less information/knowledge than the latter.

Also note the following: for $\triangleleft \in \{\leq, <\}$, if $\phi \triangleleft_{\mathbb{P}} \psi$, then by Proposition 1, $\mathbb{P}(\phi) \triangleleft \mathbb{P}(\psi)$ and, thus, by (KM2), $\kappa(\psi) \triangleleft \kappa(\phi)$. So, this case resembles axiom (E) in [67] (see also Appendix C). However, as $\mathbb{P}(\phi) \triangleleft \mathbb{P}(\psi)$ does not imply $\phi \triangleleft_{\mathbb{P}} \psi$, axiom (KM2) is weaker than (E), and it is more appropriate to the present framework since now we deal also with non-uniform probability distributions.

From what we have said above, the following is immediate.

PROPOSITION 2.

Consider any formulas ϕ and ψ w.r.t. Σ , and any probability distribution \mathbb{P} over Σ . If $\phi \triangleleft_{\mathbb{P}} \psi$ then $\kappa(\psi) \triangleleft \kappa(\phi)$, for $\triangleleft \in \{\leq, <\}$.

EXAMPLE 5.

(No surprise) Consider \mathbb{P} over $\Sigma = \{a, b\}$ with $\mathbb{P}(ab) = 1$ (all other probabilities are 0). Then $\top \models_{\mathbb{P}} a \wedge b$ and, thus, by (KM1) and (KM2) we get $0 = \kappa(\top) \geq \kappa(a \wedge b) \geq 0$, i.e. $\kappa(a \wedge b) = 0$. That is, under \mathbb{P} , $a \wedge b$ is no surprise and, thus, does not provide us with any information. On the other hand, according to [67], under uniform probability distributions, we get $\kappa_h(a \wedge b) = 2$.

As known from information theory [34], a function satisfying (KM1)–(KM3) can be defined in the following way: given a probability distribution \mathbb{P} over Σ , the *Shannon KM*, or simply the *S-KM* of a formula ϕ , denoted $\kappa_S(\phi)$, is defined as

$$\boxed{\kappa_S(\phi) = -\log_2 \mathbb{P}(\phi)} \quad (9)$$

where we postulate $\log_2 0 = -\infty$. Note that the probability distribution \mathbb{P} is a parameter of κ_S , which we omit for ease of presentation. In case of ambiguity, we will write $\kappa_S(\phi, \mathbb{P})$ instead. Noteworthy, the function κ_S is the unique solution satisfying (KM1)–(KM3), up to a multiplying constant, which is a direct consequence of the uniqueness result of the *entropy* function proved by Shannon in [64] (see also [9]). That is,

PROPOSITION 3. ([9, 64])

Any KM κ satisfying (KM1)–(KM3) is of the form $\frac{1}{\log_2 b} \cdot \kappa_S$, for a constant $b > 1$, i.e. for any formula ϕ w.r.t. Σ and any probability distribution \mathbb{P} over Σ , $\kappa(\phi) = -\log_b \mathbb{P}(\phi) = \frac{1}{\log_2 b} \cdot \kappa_S(\phi)$.

Now, it can easily be shown that Straccia's KM κ_h [67].

$$\kappa_h(\phi) = |\Sigma| - \log_2 |[\phi]| \quad (10)$$

is a special case of κ_S . That is,

PROPOSITION 4.

For any formula ϕ , given the uniform probability distribution \mathbb{P}^u over Σ , $\kappa_S(\phi) = \kappa_h(\phi)$.

An immediate consequence of Proposition 4 is also:

PROPOSITION 5.

Under uniform probability distribution, and, thus, classical logic, κ_S satisfies the (TELM) axioms illustrated in [67].

EXAMPLE 6. (Running example cont.)

Consider Example 2. Then

$$\kappa_S(\mathcal{K}) = -\log_2 0.35 = 1.515.$$

By referring to Example 3, we also have that $\mathcal{K} <_{\mathbb{P}} f$ and, thus, $1.515 \approx \kappa_S(\mathcal{K}) > \kappa_S(f) = -\log_2 0.45 \approx 1.152$, i.e. \mathcal{K} has more knowledge about the actual world than f , as expected by Proposition 2.

4 KMs-based BC

We now show how we may apply KMs to BC operations within the AGM approach, where the agent's beliefs is represented as a KB [2, 20, 21, 29]. Given a belief ϕ (of an agent), the operations we focus on are as follows:

- (1) *contraction*, a sentence α is removed, i.e. a belief ϕ is replaced by another belief $\phi \div \alpha$ that is a consequence of ϕ not entailing α ;
- (2) *expansion*, a sentence α is added to ϕ and nothing is removed, i.e. ϕ is replaced with $\phi + \alpha = \phi \wedge \alpha$;

- (3) *revision*, a sentence α is added to ϕ , and at the same time other sentences are removed if this is needed to ensure that the resulting belief $\phi \star \alpha$ is consistent.

It is well known that the so-called *Levi identity* allows to define revision in terms of contraction. Alternatively, the *Harper identity* allows to define contraction in terms of revision, i.e. (see also [7]),

$$\begin{aligned} \text{Levi identity} \quad \phi \star \alpha &= (\phi \div \neg\alpha) \wedge \alpha; \\ \text{Harper identity} \quad \phi \div \neg\alpha &= \phi \vee (\phi \star \alpha). \end{aligned}$$

We will follow the former path by defining a KM-based contraction operator and then use the Levi identity to obtain a KM-based revision operator.

To start with, let us fix a probability distribution \mathbb{P} over Σ . Additionally, for the rest of the paper, we will also consider \mathbb{P} -satisfiable formulae only, i.e. our belief must be compatible w.r.t. \mathbb{P} . Please note that \mathbb{P} is the only parameter of our setting and that different distributions may give rise to different BC operators. However, as we will show, all of them will satisfy the AGM postulates. Moreover, any AGM BC operator can be obtained from a specific probability distribution (see Propositions 11 and 12).

4.1 Contraction

A (partial)⁶ function $\div : \mathcal{L}_\Sigma \mapsto \mathcal{L}_\Sigma$ is a *contraction operation* if it satisfies the following postulates (for ease, we write ϕ_α^\div in place of $\phi \div \alpha$):

$$\begin{aligned} (\div 1) \quad \phi &\leq_{\mathbb{P}} \phi_\alpha^\div && \text{(inclusion)} \\ (\div 2) \quad \text{if } \phi &\not\leq_{\mathbb{P}} \alpha, \text{ then } \phi_\alpha^\div &\equiv_{\mathbb{P}} \phi && \text{(vacuity)} \\ (\div 3) \quad \text{if } \top &\not\leq_{\mathbb{P}} \alpha, \text{ then } \phi_\alpha^\div &\not\leq_{\mathbb{P}} \alpha && \text{(success)} \\ (\div 4) \quad \text{if } \alpha &\equiv_{\mathbb{P}} \beta, \text{ then } \phi_\alpha^\div &\equiv_{\mathbb{P}} \phi_\beta^\div && \text{(extensionality)} \\ (\div 5) \quad \phi_\alpha^\div &\wedge \alpha \leq_{\mathbb{P}} \phi && \text{(recovery)} \\ (\div 6) \quad \phi_{\alpha \wedge \beta}^\div &\leq_{\mathbb{P}} \phi_\alpha^\div \vee \phi_\beta^\div && \text{(conjunctive overlap)} \\ (\div 7) \quad \text{if } \phi_{\alpha \wedge \beta}^\div &\not\leq_{\mathbb{P}} \alpha, \text{ then } \phi_\alpha^\div &\leq_{\mathbb{P}} \phi_{\alpha \wedge \beta}^\div && \text{(conjunctive inclusion)} \end{aligned}$$

The reader who is familiar with the AGM literature will notice that we have reformulated the postulates considering formulae, which are the \mathbb{P} -entailment analogues of those defined for classical propositional logic as in [7], instead of logically closed theories.

The intuitive meaning of these postulates is as follows:

- ($\div 1$) ensures that after contraction no new information is obtained and, thus, $\kappa(\phi_\alpha^\div) \leq \kappa(\phi)$;
- ($\div 2$) indicates that if α is not \mathbb{P} -entailed by ϕ , then no change is made during the contraction and, thus, $\kappa(\phi_\alpha^\div) = \kappa(\phi)$;
- ($\div 3$) ensures that the only case for the contraction of ϕ by α to fail is when α is a always true, i.e. $\mathbb{P}(\alpha) = 1$;
- ($\div 4$) reflects the principle of syntactic independence, i.e. if two formulae are \mathbb{P} -equivalent then so are the contractions;
- ($\div 5$) states that is enough to add (by expansion) the removed formula to recover the original belief;
- ($\div 6$) states that any formula not suppressed either in the contraction of ϕ by α or in the contraction of ϕ by β , must not be suppressed by the contraction of ϕ by $\alpha \wedge \beta$;

⁶The function is partial because of the restriction of formulae to be compatible with the probability distribution \mathbb{P} .

12 Belief change based on knowledge measures

(\div 7) states that if α has been suppressed upon contracting ϕ by $\alpha \wedge \beta$, then we expect that any formula suppressed in order to remove α will also be suppressed when $\alpha \wedge \beta$ is removed.

For some more discussion and explanations about the non-probabilistic version of those axioms, we redirect the reader to e.g. [2, 7, 20].

Now, we define the contraction operation associated with a KM κ (w.r.t. \mathbb{P}) as follows. Consider formulae ϕ and α . The set of *possible remainders* of ϕ w.r.t. α , denoted $\phi \perp\!\!\!\perp \alpha$, is defined as⁷

$$\phi \perp\!\!\!\perp \alpha = \begin{cases} \{F([\phi]^+)\} & \text{if } \top \leq_{\mathbb{P}} \alpha \\ \{F([\psi]^+) \mid \phi \leq_{\mathbb{P}} \psi \text{ and } \psi \not\leq_{\mathbb{P}} \alpha\} & \text{otherwise.} \end{cases} \quad (11)$$

Essentially, as we would like to contract α from ϕ and, thus, would like to avoid that the contraction \mathbb{P} -entails α , we weaken ϕ so that we do not \mathbb{P} -entail α anymore. Please note that the set $\phi \perp\!\!\!\perp \alpha$ is *finite*, as for sets of formulae, we assumed that they cannot contain any pair of equivalent formulae. Now, from the set of possible remainders, we select those that are most specific according to \mathbb{P} -entailment, i.e.

$$\phi \perp \alpha = \{\psi \in \phi \perp\!\!\!\perp \alpha \mid \text{there is no } \psi' \in \phi \perp\!\!\!\perp \alpha \text{ s.t. } \psi' <_{\mathbb{P}} \psi\}. \quad (12)$$

This set is called the set of *remainders* of ϕ w.r.t. α .

Note that if $\phi \not\leq_{\mathbb{P}} \alpha$ or $\top \leq_{\mathbb{P}} \alpha$ then the unique remainder is \mathbb{P} -equivalent to ϕ . On the other hand, it is not difficult to see that, if $\phi \leq_{\mathbb{P}} \alpha$, any remainder ψ a formula whose models are those obtained by adding to the non-0 probability models of ϕ a single model of $\neg\alpha$. That is,

PROPOSITION 6.

Consider any probability distribution \mathbb{P} over \mathcal{W}_{Σ} and any formulae ϕ and α such that $\phi \leq_{\mathbb{P}} \alpha$ and $\top \not\leq_{\mathbb{P}} \alpha$. Then $\psi \in \phi \perp \alpha$ iff

$$\psi \equiv F([\phi]_{\Sigma}^+ \cup \{w\})$$

for some $w \in [\neg\alpha]_{\Sigma}^+$. Furthermore, $\mathbb{P}(\psi) = \mathbb{P}(F([\phi]_{\Sigma}^+)) + \mathbb{P}(w)$. On the other hand, if $\phi \not\leq_{\mathbb{P}} \alpha$ or $\top \leq_{\mathbb{P}} \alpha$ holds then $\phi \perp \alpha = \{F([\phi]^+)\}$.

We are ready to define our contraction operation. But before that we illustrate the principles we rely on.

The definition of appropriate procedures for modelling the belief dynamics of a *rational* agent cannot be attained by relying only on purely logical constraints, since they are often insufficient to determine unequivocally which beliefs have to be dropped, as illustrated by the following example (taken from [21, p.1]). A KB contains the statements ‘Juan was born in Puerto Carreño’ (α), ‘José was born in Puerto Ayacucho’ (β) and ‘two people are compatriots if they were born in the same country’ (γ). Assume that we are informed that ‘Juan and José are compatriots’ (δ); if we simply add δ to the KB, we will obtain a contradiction,⁸ hence we have to modify our KB to accommodate the new information still preserving consistency. To this end, we have various options, e.g. we could eliminate either α , β or even γ ; or, among other options, we could weaken α and β into, respectively, ‘Juan was born in Puerto Carreño or Puerto Ayacucho’ and ‘José was born in Puerto Carreño or in Puerto Ayacucho’.

⁷Recall that $F\mathcal{W}$ is the formula whose models are exactly those in the set of worlds \mathcal{W} , see Equation 2.

⁸Despite close to each other, Puerto Carreño is in Colombia, while Puerto Ayacucho is in Venezuela.

To define procedures that model such changes, we may impose some extra-logical *rationality principles* that allow us to choose only one option among the different possible ones. Different principles have been suggested in the literature (see [61] for a more detailed presentation).

One principle that is commonly accepted is the *Principle of Indifference*,⁹ which is a basic principle connected to the formal management of preferences:

$$\frac{\textit{The Principle of Indifference}}{\textit{Objects held in equal regard should be treated equally}}. \quad (13)$$

Another popular one is the *Principle of Informational Economy* [27], which simply states that we should not give up our beliefs beyond necessity: information is ‘precious’, and if we need to change our beliefs, we should drop as little information as possible in the process:¹⁰

$$\frac{\textit{The Principle of Informational Economy}}{\textit{Keep loss of information to a minimum}}. \quad (14)$$

This is generally considered the guiding principle of the AGM approach (but it has also been argued that that is not the case [60]). The Principle of Informational Economy is quite general and formally it can be interpreted in different ways. For example, it can be implemented in the form of the *Principle of Conservatism* [40], where the relative informational difference between two KBs is defined in terms of set-theoretic inclusion: if a logical theory A contains the logical theory B , then A is considered more informative than B and it is preferred or, equivalently, if a formula ϕ implies a formula ψ , ϕ is considered more informative than ψ (compare also with Proposition 2). Adhering to such a principle together with the Principle of Indifference results in the so-called *full-meet contraction*, i.e. to contract ϕ by α we take under consideration *only* the formulas in the set of remainders (according to the Principle of Conservatism), and *all* of them (according to the Principle of Indifference). That is, the full-meet contraction is the following operation:

$$\phi_{\alpha}^{\dot{-}} = \bigvee (\phi \perp \alpha). \quad (15)$$

It is well known [21, section 3.8] that full-meet contraction is not a satisfying approach to contraction, since it tends to drop too much information. A better behaviour is obtained using a *partial-meet contraction*, i.e. considering only *some* of the elements of the set $\phi \perp \alpha$, i.e.

$$\phi_{\alpha}^{\dot{-}} = \bigvee \gamma(\phi \perp \alpha), \quad (16)$$

where γ is a so-called *choice function*.¹¹ Different implementations of the partial-meet contraction depend on the rationale that we consider in the definition of the choice function.

Here we propose to use KMs as a way of modelling the expectations of a rational agent. To this end, we introduce the PMS, where ‘surprise’ is expressed in information-theoretic terms, i.e. the more probable an event is the less it is surprising.

$$\frac{\textit{The Principle of Minimal Surprise}}{\textit{The less surprising option should be preferred}}. \quad (17)$$

⁹Here we refer to the *Principle of Indifference* as stated in the area of BC [61], and not to the homonym principle in the theory of epistemic probability. The latter, also known as *Principle of Insufficient Reason*, states that if we are given a finite set of options and there is no reason to consider one more likely than the others, then we should assign the same probability (interpreted as degree of belief) to each possible option. The two principles, despite being clearly very close, belong to different research areas and have a different meaning.

¹⁰See also the information loss of a contraction in Equation 19.

¹¹ γ returns a set of formulae.

This principle simply states that facing uncertainty about which is the actual world, an agent should opt for the most expected (i.e. the least surprising) option. Such a principle is intuitive, and it has been also assumed in major approaches to conditional non-monotonic reasoning, such as [28, 65]. Our approach follows the same line, but it uses KMs in order to create a bridge between our probabilistic framework and the ranking of the information in terms of level of surprise. Note that in the KM theory, the informative value associated with a piece of information is directly proportional to its exceptionality: the less probable it is, the more informative it conveys (it is an immediate consequence of the axiom **(KM2)**). Hence, a formula ϕ is strictly less surprising than a formula ψ iff $\kappa_S(\phi) < \kappa_S(\psi)$.

Formally, in order to contract ϕ by α we select some formulae in the remainder set $\phi \perp \alpha$, which, relying on KMs and the PMS, results in choosing all the most expected ones. Moreover, to be compliant also with the Principle of Indifference, we select all the remainders of least surprise. The result is Equation 18 in the definition below.

DEFINITION 1. (KM-contraction)

Given any formulae ϕ, α and any KM κ , then the *KM-contraction operator* $\dot{\div}_\kappa$ is defined as

$$\phi_\alpha^{\dot{\div}_\kappa} = \bigvee_{\kappa} \min(\phi \perp \alpha), \quad (18)$$

where, given any set of formulae, Γ

$$\min_{\kappa}(\Gamma) = \{\phi \in \Gamma \mid \forall \psi \in \Gamma, \kappa(\phi) \leq \kappa(\psi)\}.$$

So, first, information is maximized in order to save as much knowledge as possible; second, surprise is minimized among the selected possibilities.

Note that $\phi_\alpha^{\dot{\div}_\kappa}$ is deterministic in the sense that it is uniquely determined by the given probability distribution \mathbb{P} (and different probability distributions may give rise to different contractions).

Given any contraction operator $\dot{\div}$, we may also associate with it a quantitative function that numerically specifies how much information has been lost by contracting ϕ with α . That is, the *information change of a contraction* $\dot{\div}$ (w.r.t. a given probability distribution \mathbb{P}) is defined as

$$IC^{\dot{\div}}(\phi, \alpha) = \kappa(\phi_\alpha^{\dot{\div}}) - \kappa(\phi). \quad (19)$$

Note that $IC^{\dot{\div}}(\phi, \alpha) \leq 0$, i.e. we always loose information via contraction.

Proposition 6 provides us with a simple constructive way to compute a contraction. That is, given \mathbb{P} , ϕ and α , $\phi_\alpha^{\dot{\div}_\kappa}$ can be computed in the following way:

- (1) if $[\phi]^+ \not\subseteq [\alpha]$ or $[\top]^+ \subseteq [\alpha]$ then return $F([\phi]^+)$ and we are done; otherwise
- (2) determine the remainder set $\phi \perp \alpha$ as the set of formulae ψ_i , where for each $w_i \in [\neg\alpha]^+$, $\psi_i = F([\phi]^+ \cup \{w_i\})$; and
- (3) eventually, compute $\kappa(\psi_i)$ from $\mathbb{P}(\psi_i)$, determine so $\min_{\kappa}(\phi \perp \alpha)$ and, conclude by applying Equation 18.

EXAMPLE 7. (Running example cont.)

Assume that we would like to contract \mathcal{K} with f , i.e. let us determine $\mathcal{K} \dot{\div}_{\kappa_S} f$. We already know from Example 3 that $\mathcal{K} \leq_{\mathbb{P}_S} f$ and, thus, \mathcal{K} cannot be itself a possible reminder. Therefore, we

need to weaken it. To do so, at first, we build $\mathbb{P}0$, according to Example 2, and then determine the remainder set by using Proposition 6 and the models of \mathcal{K} identified in Example 1 together with their probabilities as from Example 2.

Now, the non-zero probability models of $\neg f$ are those in \mathbb{P} for which f is false. These can be identified by

$$[\neg f]^+ = \{bp\bar{o}\bar{f}w^{0.15}, b\bar{p}o\bar{f}w^{0.2}, b\bar{p}\bar{o}\bar{f}w^{0.2}\}.$$

As a consequence, we get three remainders, one for each element in $[\neg f]^+$. That is,

$$\begin{aligned}\psi_1 &= F([\mathcal{K}] \cup \{bp\bar{o}\bar{f}w\}) \\ \psi_2 &= F([\mathcal{K}] \cup \{b\bar{p}o\bar{f}w\}) \\ \psi_3 &= F([\mathcal{K}] \cup \{b\bar{p}\bar{o}\bar{f}w\}).\end{aligned}$$

Moreover, it is easily verified that

$$\begin{aligned}\mathbb{P}(\psi_1) &= 0.35 + 0.15 = 0.5 \\ \mathbb{P}(\psi_2) &= 0.35 + 0.2 = 0.55 \\ \mathbb{P}(\psi_3) &= 0.35 + 0.2 = 0.55\end{aligned}$$

and consequently, $\kappa_S(\psi_1) = 1.0$, $\kappa_S(\psi_2) \approx 0.862$ and $\kappa_S(\psi_3) \approx 0.862$. Eventually, the knowledge minimal (least surprising/most probable) ones are ψ_2 and ψ_3 and, thus, we conclude with

$$\begin{aligned}\mathcal{K} \dot{\div}_{\kappa_S} f &\equiv \psi_1 \vee \psi_2 \\ &\equiv (\mathcal{K} \vee F b\bar{p}\bar{o}\bar{f}w \vee F b\bar{p}o\bar{f}w) \\ &\equiv (\mathcal{K} \vee F b\bar{p}\bar{f}w).\end{aligned}\tag{20}$$

Therefore, the contraction weakens the initial belief in Example 1 by making the belief ‘birds not being penguins do not fly and have wings’ possible as well (w.r.t. the given probability distribution). The latter is also the ‘least surprising’ (most probable) candidate among the remainders. Moreover, note that $\mathbb{P}(\mathcal{K} \dot{\div}_{\kappa_S} f) = 0.75$, so $\kappa_S(\mathcal{K} \dot{\div}_{\kappa_S} f) \approx 0.415$, and, thus, the information change of this contraction can be quantified as

$$IC^{\dot{\div}_{\kappa_S}}(\mathcal{K}, f) = 0.415 - 1.515 = -1.1.\tag{21}$$

From Proposition 6 and from what we have seen in Example 7 above, it is then not surprising that the following proposition holds.

PROPOSITION 7.

Consider any probability distribution \mathbb{P} over \mathcal{W} , any KM κ and formulae ϕ and α . Then

$$\phi_{\alpha}^{\dot{\div}_{\kappa}} \equiv \begin{cases} F([\phi]^+ \cup \min_{\kappa}([\neg\alpha]^+)) & \text{if } \phi \leq_{\mathbb{P}} \alpha \\ F[\phi]^+ & \text{otherwise,} \end{cases}\tag{22}$$

where¹²

$$\min_{\kappa}([\neg\alpha]^+) = \{w \in [\neg\alpha]^+ \mid \forall w' \in [\neg\alpha]^+, \kappa(Fw) \leq \kappa(Fw')\}.$$

¹²Please note that \min_{κ} is a function that when applied to a set of formulae returns a set of formulae, while when applied to a set of worlds returns a set of worlds.

A consequence of Proposition 7 is the following proposition establishing exactly the information loss of a contraction of ϕ w.r.t. α in terms of the probability of ϕ and the probability of having a knowledge minimal (least surprising) $\neg\alpha$ world.

COROLLARY 1.

Consider any probability distribution \mathbb{P} over \mathcal{W} , any KM κ and formulae ϕ and α . Then

$$IC^{\div}(\phi, \alpha) = \begin{cases} -\log_2\left(1 + \frac{\mathbb{P}(F \min_{\kappa}([\neg\alpha]^+))}{\mathbb{P}(\phi)}\right) & \text{if } \phi \leq_{\mathbb{P}} \alpha \\ 0 & \text{otherwise.} \end{cases} \quad (23)$$

So, for instance, by referring to Example 7 and Equation 20, we have $\mathbb{P}(\mathcal{K}) = 0.35$ and $\mathbb{P}(F b\bar{p}\bar{f}w) = \mathbb{P}(F b\bar{p}o\bar{f}w) + \mathbb{P}(F b\bar{p}\bar{o}\bar{f}w) = 0.2 + 0.2 = 0.4$ and, thus,

$$\begin{aligned} IC^{\div}(\mathcal{K}, f) &= -\log_2\left(1 + \frac{0.4}{0.35}\right) \\ &= -1.1 \end{aligned}$$

which confirms Equation 21.

The following proposition tells us that a KM-contraction operator¹³ is an AGM contraction operation.

PROPOSITION 8.

For all probability distributions \mathbb{P} and KM-contraction operator \div_{κ} based on \mathbb{P} , \div_{κ} is an AGM contraction operator, i.e. it satisfies postulates $(\div 1)$ – $(\div 7)$.

Next, we prove that every AGM contraction operation can be represented as a KM-contraction operation by appropriately defining a probability distribution underlying the KM-contraction operation. That is, the various AGM contraction operations differentiate each other on a probability distribution over worlds they rely on.¹⁴

To start with, it is well known that AGM operations can be characterized by a total pre-order over the set of interpretations \mathcal{W}_{Σ} [46]. Formally, given a formula ϕ , a total pre-order \leq_{ϕ} on \mathcal{W}_{Σ} , (we denote the related strict relation as $<_{\phi}$, while denote related symmetric relation as \simeq_{ϕ}), is a *faithful assignment* for ϕ iff the following conditions hold:

- (1) if $\{w, w'\} \subseteq [\phi]$, then $w \simeq_{\phi} w'$; and
- (2) if $w \in [\phi]$ and $w' \notin [\phi]_{\Sigma}$, then $w <_{\phi} w'$.

That is, in a faithful assignment for ϕ , ϕ models are ranked lower than non- ϕ models. Eventually, given a total pre-order \leq over worlds, the minimal ψ -models w.r.t. \leq are defined in the usual way:

$$\min_{\leq}([\psi]) = \{w \in [\psi] \mid \forall w' \in [\psi], w \leq w'\}.$$

The notion of a faithful assignment allows us to characterize contraction and revision operations satisfying $(\div 1)$ – $(\div 7)$. The following is a representation result that can be easily derived and is the KM-based analogue to [7, Theorem 14] and [46, Theorem 3.3].¹⁵

¹³We recall that different probability distributions may give rise to different KM-contraction operators.

¹⁴Of course, such a characterization would not be possible using Straccia's κ_h [67].), where the uniform distribution is assumed.

¹⁵The reader may compare it also with Proposition 7.

TABLE 4 ϕ -faithful ranking r with possible r -faithful probability distribution \mathbb{P} , and the relative KM.

r	worlds	\mathbb{P}	κ_{Σ}
2	$\bar{a}\bar{b}c \quad \bar{a}\bar{b}\bar{c}$	1/16	4
1	$\bar{a}bc \quad \bar{a}\bar{b}\bar{c} \quad \bar{a}bc \quad \bar{a}\bar{b}\bar{c}$	2/16	3
0	$abc \quad ab\bar{c}$	3/16	2.415

PROPOSITION 9. ([7, 46])

Let ϕ, α be formulae. An operation \div on ϕ satisfies $(\div 1) - (\div 7)$ if and only if there is a faithful assignment \leq_{ϕ} for ϕ such that

$$\phi_{\alpha}^{\div} \equiv F([\phi] \cup \min_{\leq_{\phi}}([\neg\alpha])).$$

So, this proposition tells us that we may express the set of models of the contraction of ϕ by α as the union of the models of ϕ and of the minimal counter-models of α w.r.t. \leq_{ϕ} .

Given that Σ is finite, the total pre-order \leq_{ϕ} can be easily translated into a ranking function r :¹⁶ namely

$$r(w) = \begin{cases} 0, & \text{if } w \in \min_{\leq_{\phi}}(\mathcal{W}_{\Sigma}) \\ i, & \text{if } w \in \min_{\leq_{\phi}}(\mathcal{W}_{\Sigma} \setminus \{w' \mid r(w') < i\}) \end{cases} \quad (24)$$

which we call a ϕ -faithful ranking. Note that models of ϕ have rank 0. The contraction \div associated with r is then defined as

$$\phi_{\alpha}^{\div r} \equiv F([\phi] \cup \min_r([\neg\alpha])). \quad (25)$$

EXAMPLE 8.

Let $\Sigma = \{a, b, c\}$ and $\phi =_{\text{def}} a \wedge b$. A ϕ -faithful ranking r is shown in Table 4. The other columns will be addressed in Example 9 later on. To determine $(a \wedge b)_b^{\div r}$, we have to consider $\min_r([\neg b]) = \{\bar{a}\bar{b}c, \bar{a}\bar{b}\bar{c}\}$. Hence $[(a \wedge b)_b^{\div r}] = \{abc, ab\bar{c}, \bar{a}\bar{b}c, \bar{a}\bar{b}\bar{c}\}$ and, thus, $(a \wedge b)_b^{\div r} \equiv a$.

Now, given any set of worlds \mathcal{W} and any ϕ -faithful ranking r over \mathcal{W} , we say that \mathbb{P} is a r -faithful probability distribution over \mathcal{W} if it respects the following constraints: (i) for any $w, w' \in \mathcal{W}$, $\mathbb{P}(w) \geq \mathbb{P}(w')$ iff $r(w) \leq r(w')$; and (ii) for any $w \in \mathcal{W}$, $\mathbb{P}(w) > 0$. That is, the lower is the rank of a world the higher is its probability and vice versa.

PROPOSITION 10.

For any formula ϕ , set of worlds \mathcal{W} and ϕ -faithful ranking r over \mathcal{W} , there exists an r -faithful probability distribution \mathbb{P} over \mathcal{W} .

Now, using \mathbb{P} constructed as in the proof of Proposition 10, we can generate a KM-contraction based on \mathbb{P} that corresponds to the AGM-contraction we started with.

¹⁶For ease, we omit the reference to ϕ if obvious from the context.

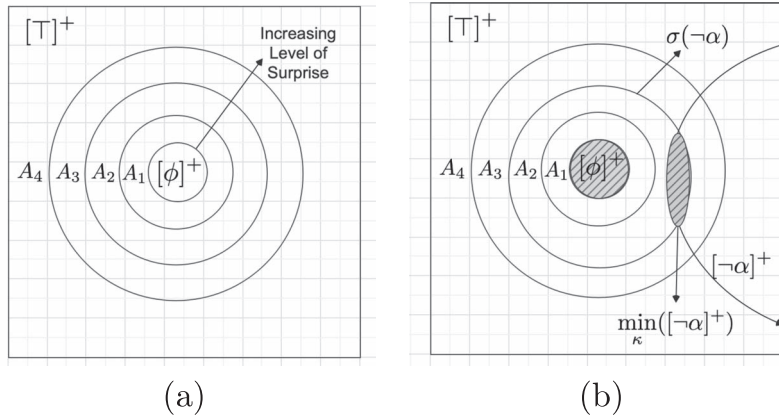


FIGURE 1 (a) Probabilistic system of spheres; (b) probabilistic system of spheres with contraction $\phi_{\alpha}^{\dot{\div}}$ shaded.

PROPOSITION 11.

Let ϕ be any formula and let $\dot{\div}$ be an AGM-contraction. Then there is a KM contraction $\dot{\div}_{\kappa}$ s.t., for all α , $\phi_{\alpha}^{\dot{\div}} \equiv \phi_{\alpha}^{\dot{\div}_{\kappa}}$ holds.

EXAMPLE 9.

Consider Table 4. The \mathbb{P} (resp. κ_S) column indicates the probability (resp. KM) value associated with each world within the same rank according to an r -faithful probability distribution \mathbb{P} , as by Proposition 10. For instance, if $\mathbb{P}(abc) = 3/16$ (resp. $\kappa_S(abc) = 2.415$). To determine $(a \wedge b)_b^{\dot{\div}_{\kappa_S}}$ w.r.t. that \mathbb{P} , we have to consider $\min_{\kappa_S}([-b]) = \{\bar{a}\bar{b}c, a\bar{b}\bar{c}\}$. Hence, $[(a \wedge b)_b^{\dot{\div}_{\kappa_S}}] = \{abc, a\bar{b}\bar{c}, \bar{a}\bar{b}c, a\bar{b}\bar{c}\}$ and $(a \wedge b)_b^{\dot{\div}_{\kappa_S}} \equiv a \equiv (a \wedge b)_b^{\dot{\div}_r}$, in agreement with Proposition 11.

A consequence of Propositions 8 and 11 is the following.

PROPOSITION 12.

Given any formula ϕ , for every formula α , $\phi_{\alpha}^{\dot{\div}}$ is an AGM contraction operation iff it is a KM-contraction operation.

4.2 Sphere-based view of contraction

An interesting way of viewing a contraction operator may be given pictorially in terms of a probabilistic variant of a system of spheres (see, e.g. [31, 61]), as illustrated in Figure 1 (a).

Specifically, consider a formula ϕ and a formula α to be contracted from ϕ . A sphere σ is a set of possible worlds.¹⁷ A system of spheres \mathbb{S}_{ϕ} centred on ϕ ¹⁸ is an order $\sigma_0 < \sigma_1 \dots < \sigma_n$ of nested spheres (in the sense of set inclusion $\sigma_i \subset \sigma_{i+1}$, $0 \leq i < n$) in which the smallest or innermost

¹⁷That is, worlds having non-zero probability.

¹⁸We will omit the reference to \mathbb{S}_{ϕ} if clear from context.

sphere is $\sigma_0 = [\phi]^+$ (the possible models of ϕ), the outermost is $\sigma_n = [\top]^+$ (all possible models) and where

- (1) all possible worlds in the *annulus* $A_{i+1} = \sigma_{i+1} \setminus \sigma_i$ ($0 \leq i < n$) have the same probability, i.e. for $w, w' \in A_{i+1}$, $\mathbb{P}(w) = \mathbb{P}(w')$; and
- (2) the probability of a world decreases from the innermost annulus A_1 to the outermost annulus A_n , i.e. for $w \in A_i$ and $w' \in A_{i+1}$ ($i > 0$), we have $\mathbb{P}(w) > \mathbb{P}(w')$.

Intuitively, the agent believes the actual world to be one of the ϕ -worlds but does not have sufficient information to establish which one. However, the agent may be mistaken, in which case it believes that the actual world is most likely to be one of those in the next greater sphere and so on. That is, with 2 we want to say that the nesting of the spheres defines an order in terms of *increasing surprise* of the worlds an agent considers possible. Moving from the innermost annulus to the outermost one increases the surprise, from an information-theoretic point of view. On the other hand, point 1 encodes the fact that each possible world within an annulus is equally plausible and there is no reason to prefer one world to the other if we accept the Principle of Indifference. Now, in order to contract α from ϕ ¹⁹ we need to add at least one $\neg\alpha$ model to the models of ϕ . As we are inspired by the PMS, to do so we consider the smallest sphere $\sigma(\neg\alpha)$ intersecting $[\neg\alpha]^+$, which is indeed $\min_\kappa([\neg\alpha]^+) = \sigma(\neg\alpha) \cap [\neg\alpha]^+$ (see Figure 1 (b)). By the Principle of Indifference, each world in this intersection is equally plausible and, thus, the possible models of $\phi \dot{-} \alpha$ (see also by Proposition 7) are those in $[\phi]^+ \cup \min_\kappa([\neg\alpha]^+)$, i.e. the shaded ones in Figure 1 (b).

REMARK 5.

We may express $\sigma(\neg\alpha)$ in the following way. For a formula β , let p_β^{\max} be the maximal probability of the possible β -models, i.e.

$$p_\beta^{\max} = \max\{\mathbb{P}(w) \mid w \in [\beta]^+\}.$$

Note that $\mathbb{P}(w) = p_\beta^{\max}$ for all $w \in \min_\kappa[\beta]^+$. Then

$$\sigma(\neg\alpha) = [\phi]^+ \cup \{w \in [\top]^+ \setminus [\phi]^+ \mid \mathbb{P}(w) \geq p_{\neg\alpha}^{\max}\}. \quad (26)$$

4.3 A contraction without recovery

As we have seen, our contraction operator is an AGM contraction and, thus, satisfies the (in)famous *recovery* postulate ($\div 5$) around which there is a longstanding debate. Various contraction proposals have been developed not satisfying ($\div 5$), such as *saturatable contraction* [50], *semi contraction* [23], *severe withdrawal* [61], *recuperative/ring withdrawal* [19] and many more (see e.g. [21]).

The sole purpose of this section is to illustrate how one may build from our KM-based contraction operator one not satisfying the recovery postulate. To do so, we focus on *severe withdrawal* [61], denoted $\dot{-}$, since one may give for it a simple sphere-based pictorial explanation and formalization.²⁰

¹⁹For ease, assume $[\phi]^+ \subseteq [\alpha]^+$ and $[\neg\alpha]^+ \neq \emptyset$.

²⁰For many other approaches and their sphere-based view, see e.g. [61, appendix B].




FIGURE 2 System of spheres for severe withdrawal showing $\phi_{\alpha}^{\#}$ shaded.

From a sphere-based point of view, given a formula ϕ and a formula α to be contracted,²¹ the severe withdrawal of α from ϕ , denoted $\phi_{\alpha}^{\#}$, is the shaded area in Figure 2 (compare with Figure 1 (b)).

Essentially, severe withdrawal relies also on the *Principle of Weak Preference*, according to which ‘if one world is considered at least as plausible as another then the former should be considered if the latter is. That is, an agent has determined a preference over worlds (the spheres) and does not prefer the (closest) $\neg\alpha$ -worlds over the (closer) α -worlds just because it is giving up belief in α . Its preferences are established prior to the change and we assume that there is no reason to alter them in light of the new information’ [61]. Consequently, by following [61], we may define our km-based *severe withdrawal* operator as the formula corresponding to the set of possible worlds in the smallest sphere intersecting $[\neg\alpha]^+$, i.e. $\mathcal{K}_{\alpha}^{\#} = F\sigma(\neg\alpha)$.

Formally, as we did for contraction, we provide next the \mathbb{P} -analogue postulates of severe withdrawal given in [61] as follows: a function $\#$ is a *severe withdrawal* function if it satisfies the following postulates:

- ($\#$ 1) $\phi \leq_{\mathbb{P}} \phi_{\alpha}^{\#}$
- ($\#$ 2) if $\phi \not\leq_{\mathbb{P}} \alpha$ or $\top \leq_{\mathbb{P}} \alpha$, then $\phi_{\alpha}^{\#} \equiv_{\mathbb{P}} \phi$
- ($\#$ 3) if $\top \not\leq_{\mathbb{P}} \alpha$, then $\phi_{\alpha}^{\#} \not\leq_{\mathbb{P}} \alpha$
- ($\#$ 4) if $\alpha \equiv_{\mathbb{P}} \beta$, then $\phi_{\alpha}^{\#} \equiv_{\mathbb{P}} \phi_{\beta}^{\#}$
- ($\#$ 6a) if $\top \not\leq_{\mathbb{P}} \alpha$, then $\phi_{\alpha \wedge \beta}^{\#} \leq_{\mathbb{P}} \phi_{\alpha}^{\#}$
- ($\#$ 7) if $\phi_{\alpha \wedge \beta}^{\#} \not\leq_{\mathbb{P}} \alpha$, then $\phi_{\alpha}^{\#} \leq_{\mathbb{P}} \phi_{\alpha \wedge \beta}^{\#}$

Postulates ($\#$ 1), ($\#$ 3), ($\#$ 4) and ($\#$ 7) are as for AGM contraction. Postulate ($\#$ 2) is as the vacuity postulate (\div 3), but contains an additional antecedent $\top \not\leq_{\mathbb{P}} \alpha$ to take care of the limiting case, which was previously handled with the aid of the recovery postulate, which now is missing. Postulate ($\#$ 6) has been replaced with the stronger *antitony condition* ($\#$ 6a), which states that anything that is

²¹Again, for ease we assume $[\neg\alpha]^+ \neq \emptyset$.




FIGURE 3 System of spheres where $\phi_{\alpha \wedge \beta}^{\diverge_{\kappa}} \leq_{\mathbb{P}} \beta$ holds.

given up to remove a strong sentence (the conjunction of α and β) should also be given up when removing a weaker sentence (α), provided the latter is not logically true [61].

We now show, by relying on \mathbb{P} -analogue construction as proposed by [61], how to define a severe withdrawal function \diverge_{κ} from our KM-based contraction function \diverge_{κ} .

DEFINITION 2. (KM-Severe withdrawal from \diverge)

Given formulae ϕ , α and a KM-contraction \diverge_{κ} based on a probability distribution \mathbb{P} . Then the *KM-severe withdrawal function defined from \diverge_{κ}* is defined as follows:

$$\phi_{\alpha}^{\diverge_{\kappa}} = \begin{cases} \bigwedge \{\beta \mid \phi_{\alpha \wedge \beta}^{\diverge_{\kappa}} \leq_{\mathbb{P}} \beta\} & \text{if } \top \not\leq_{\mathbb{P}} \alpha \\ F[\phi]^+ & \text{otherwise.} \end{cases} \quad (27)$$

Please note that again the set $\{\beta \mid \phi_{\alpha \wedge \beta}^{\diverge_{\kappa}} \leq_{\mathbb{P}} \beta\}$ is finite, as for sets of formulae, we assumed that they cannot contain any pair of equivalent formulae.

Now, observe that in order $\phi_{\alpha \wedge \beta}^{\diverge_{\kappa}} \leq_{\mathbb{P}} \beta$ to hold, it should be the case that the smallest sphere $\sigma(\neg\alpha)$ intersecting $[\neg\alpha]^+$ has to be interior to the smallest sphere $\sigma(\neg\beta)$ intersecting $[\neg\beta]^+$, i.e. $\sigma(\neg\alpha) < \sigma(\neg\beta)$, as otherwise we would add $\neg\beta$ -worlds to $\phi_{\alpha \wedge \beta}^{\diverge_{\kappa}}$ and, thus, $\phi_{\alpha \wedge \beta}^{\diverge_{\kappa}} \leq_{\mathbb{P}} \beta$ would not hold. Consequently, $\sigma(\neg\alpha)$ contains no $\neg\beta$ -worlds, i.e. β is satisfied by all worlds in $\sigma(\neg\alpha)$ (see Figure 3).

Therefore, $\{\beta \mid \phi_{\alpha \wedge \beta}^{\diverge_{\kappa}} \leq_{\mathbb{P}} \beta\}$ is the set of formulae that hold in $\sigma(\neg\alpha)$, which in turns mean that $\phi_{\alpha}^{\diverge_{\kappa}} = F\sigma(\neg\alpha)$. From what we have said, the following proposition is immediate.

PROPOSITION 13.

Given any formulae ϕ , α , any KM-contraction \diverge_{κ} and the KM-severe withdrawal function \diverge_{κ} defined from \diverge_{κ} . Then

$$\phi_{\alpha}^{\diverge_{\kappa}} \equiv \begin{cases} F\sigma(\neg\alpha) & \text{if } \top \not\leq_{\mathbb{P}} \alpha \\ F[\phi]^+ & \text{otherwise,} \end{cases} \quad (28)$$

where $\sigma(\neg\alpha)$ is defined as in Equation 26.

22 Belief change based on knowledge measures

In other words, ‘by contracting α from ϕ according to Definition 2, we keep those formulae β that would be retained when given the choice to remove either α or β (or both)’ [61].

REMARK 6.

In [19] it is shown that one may define a ring withdrawal operator using a severe withdrawal operator or even using an AGM contraction operator. Along a similar line, following what we have done in this section, it not difficult to see that one may define the analogue KM-based ring withdrawal operator using a KM-based severe withdrawal operator or even using a KM-based AGM contraction operator.

By Proposition 13, obviously

$$\phi_{\alpha}^{\dot{\div}\kappa} \leq_{\mathbb{P}} \phi_{\alpha}^{\ddot{\div}\kappa} \quad (29)$$

holds (see also Figure 2). Consequently, both

$$\kappa(\phi_{\alpha}^{\dot{\div}\kappa}) \geq \kappa(\phi_{\alpha}^{\ddot{\div}\kappa})$$

and

$$0 \geq IC^{\dot{\div}\kappa}(\phi, \alpha) \geq IC^{\ddot{\div}\kappa}(\phi, \alpha)$$

hold, where $IC^{\ddot{\div}\kappa}(\phi, \alpha)$ is the obvious adaption of Equation 19 to $\ddot{\div}\kappa$, i.e.

$$IC^{\ddot{\div}\kappa}(\phi, \alpha) = \kappa(\phi_{\alpha}^{\ddot{\div}\kappa}) - \kappa(\phi). \quad (30)$$

This confirms that severe withdrawal loses more knowledge/information than an AGM contraction. Specifically, the information change of the KM-based severe withdrawal function and analogue of Corollary 1 is

COROLLARY 2.

Consider any probability distribution \mathbb{P} over \mathcal{W} and any formulae ϕ and α . Then

$$IC^{\ddot{\div}\kappa}(\phi, \alpha) = \begin{cases} -\log_2 \frac{\mathbb{P}(F\sigma(\neg\alpha))}{\mathbb{P}(\phi)} & \text{if } \top \not\leq_{\mathbb{P}} \alpha \\ 0 & \text{otherwise.} \end{cases} \quad (31)$$

Note that, if $\top \not\leq_{\mathbb{P}} \alpha$, as $[\phi]^+ \subseteq \sigma(\neg\alpha)$, we have $\frac{\mathbb{P}(F\sigma(\neg\alpha))}{\mathbb{P}(\phi)} \geq 1$. So, we may also write $IC^{\ddot{\div}\kappa}(\phi, \alpha)$ as

$$IC^{\ddot{\div}\kappa}(\phi, \alpha) = \begin{cases} -\log_2(1 + \frac{\mathbb{P}(F(\sigma(\neg\alpha) \setminus [\phi]^+))}{\mathbb{P}(\phi)}) & \text{if } \top \not\leq_{\mathbb{P}} \alpha \\ 0 & \text{otherwise.} \end{cases} \quad (32)$$

We conclude this section by showing that $\ddot{\div}\kappa$ is indeed a severe withdrawal function in the sense that it satisfies the postulates for severe withdrawal.

PROPOSITION 14.

Consider any probability distribution \mathbb{P} and any KM-contraction operator $\dot{\div}\kappa$ based on \mathbb{P} . Then the function $\ddot{\div}\kappa$ is a severe withdrawal operator, i.e. it satisfies the postulates ($\ddot{\div}$ 1)-($\ddot{\div}$ 4), ($\ddot{\div}$ 6a) and ($\ddot{\div}$ 7).

We can also prove the opposite direction, i.e. for each severe withdrawal operator $\ddot{\div}$ there is a KM κ s.t. $\ddot{\div}\kappa$ corresponds to $\ddot{\div}$.

PROPOSITION 15.

Given any formula ϕ , for every formula α , ϕ_{α}^{\div} is a severe withdrawal operation iff it is a KM-severe withdrawal operation based on some probability distribution \mathbb{P} .

4.4 Belief revision

Now, we address the other main BC operation, namely *revision*. A function \star is an AGM *revision* in our setting iff it satisfies the following postulates, which are the \mathbb{P} -entailment analogues of those in [7]:²²

- | | | |
|------|--|------------------|
| (★1) | $\phi_{\alpha}^{+} \leq_{\mathbb{P}} \phi_{\alpha}^{\star}$ | (inclusion) |
| (★2) | if $\phi \not\leq_{\mathbb{P}} \neg\alpha$, then $\phi_{\alpha}^{\star} \equiv_{\mathbb{P}} \phi_{\alpha}^{+}$ | (vacuity) |
| (★3) | $\phi_{\alpha}^{\star} \leq_{\mathbb{P}} \alpha$ | (success) |
| (★4) | if $\alpha \equiv_{\mathbb{P}} \beta$, then $\phi_{\alpha}^{\star} \equiv_{\mathbb{P}} \phi_{\beta}^{\star}$ | (extensionality) |
| (★5) | if $\alpha \not\leq_{\mathbb{P}} \perp$, then $\phi_{\alpha}^{\star} \not\leq_{\mathbb{P}} \perp$ | (consistency) |
| (★6) | $(\phi_{\alpha}^{\star})_{\beta}^{+} \leq_{\mathbb{P}} \phi_{\alpha \wedge \beta}^{\star}$ | (superexpansion) |
| (★7) | if $\phi_{\alpha}^{\star} \not\leq_{\mathbb{P}} \neg\beta$, then $\phi_{\alpha \wedge \beta}^{\star} \leq_{\mathbb{P}} (\phi_{\alpha}^{\star})_{\beta}^{+}$ | (subexpansion) |

The intuitive meaning of these postulates is as follows:

- (★1) states that the revision consists of α and the formulae of ϕ that do not contradict α ;
- (★2) states that in the case that if α does not contradict any formula in ϕ there is no reason to remove any of them;
- (★3) states that the new formula α must be incorporated in the revision;
- (★4) states that the revision operation must be syntax independent: \mathbb{P} -equivalent formulae must yield the same result;
- (★5) states that unless α is not \mathbb{P} -consistent the result of the revision must be \mathbb{P} -consistent;
- (★6), (★7) state that if ϕ is to be changed minimally so as to include two formulae α and β such a change should be possible by first revising ϕ w.r.t. α and then expanding ϕ_{α}^{\star} by β , provided that β does not contradict ϕ_{α}^{\star} .

As mentioned at the beginning of this section, every revision operation can be modelled by composing a contraction operation and an expansion one through the Levi Identity [49]: in fact, it is well known that AGM contractions and AGM revisions are interconnected through the Levi Identity.

PROPOSITION 16. ([2])

A revision \star is an AGM revision operator iff it can be defined via the Levi Identity using an AGM contraction operator \div .

Applying now the Levi identity we can immediately define the class of KM-revision operations \star_{κ} by contracting $\neg\alpha$ from ϕ and then adding α to the result obtaining

$$\phi_{\alpha}^{\star_{\kappa}} = \left(\bigvee_{\kappa} \min(\phi \perp \neg\alpha) \right) \wedge \alpha. \quad (33)$$

From Propositions 12 and 16, we can conclude with

²²Note that (★2) implies (★1).

PROPOSITION 17.

Given any formula ϕ , for every formula α , ϕ_α^\star is an AGM-revision operation iff it is a KM-revision operation.

EXAMPLE 10.

Consider Table 4. Assume we believe $(a \wedge b)$ and we are informed that $\neg b$ holds. Then the revision $(a \wedge b)_{\neg b}^{\star_{\kappa_S}}$ corresponds to $(a \wedge b)_{\neg b}^{\dagger_{\kappa_S}} \wedge \neg b$. Since $(a \wedge b)_{\neg b}^{\dagger_{\kappa_S}} \equiv a$, we can conclude that $(a \wedge b)_{\neg b}^{\star_{\kappa_S}} \equiv (a \wedge \neg b)$.

The analogue of Proposition 7 for the revision case is as follows.

PROPOSITION 18.

Consider any probability distribution \mathbb{P} over \mathcal{W} , and any KM κ and formulae ϕ and α Then

$$\phi_\alpha^\star = \begin{cases} F(\min_\kappa([\alpha]^+)) & \text{if } \phi \leq_{\mathbb{P}} \neg\alpha \\ F([\phi]^+ \cap [\alpha]^+) & \text{otherwise.} \end{cases} \quad (34)$$

Finally, likewise contraction, for expansion and revision operators $+$ and \star , respectively, we may also introduce related quantitative measures such as *information change of expansion*

$$IC^+(\phi, \alpha) = \kappa(\phi + \alpha) - \kappa(\phi) \quad (35)$$

and *information change of revision*

$$IC^\star(\phi, \alpha) = \kappa(\phi \star \alpha) - \kappa(\phi) \quad (36)$$

respectively. Clearly, while the former is non-negative, the latter may not.

We conclude this section by providing the analogue of Corollary 1, both for information change of expansion and revision, whereas for the latter we use Proposition 18.

COROLLARY 3.

Consider any probability distribution \mathbb{P} over \mathcal{W} , any KM κ and any formulae ϕ and α . Then

$$IC^+(\phi, \alpha) = -\log_2 \mathbb{P}(\alpha \mid \phi) \quad (37)$$

$$IC^\star(\phi, \alpha) = \begin{cases} \log_2 \frac{\mathbb{P}(\phi)}{\mathbb{P}(F(\min_\kappa([\alpha]^+))} & \text{if } \phi \leq_{\mathbb{P}} \neg\alpha \\ IC^+(\phi, \alpha) & \text{otherwise.} \end{cases} \quad (38)$$

4.5 IBR in brief

In this section, we also give a succinct look into the problem of *iterated revision* [16]. That is, we are going to look at possible ways in which we can apply a sequence of revision operations in our framework. In [16] four well-known extra postulates for iterated revision have been proposed, which we adapt here to our probabilistic setting as follows:

- (C1) if $\psi \leq_{\mathbb{P}} \alpha$, then $(\phi_\alpha^\star)_\psi^\star \equiv_{\mathbb{P}} \phi_\psi^\star$
- (C2) if $\psi \leq_{\mathbb{P}} \neg\alpha$, then $(\phi_\alpha^\star)_\psi^\star \equiv_{\mathbb{P}} \phi_\psi^\star$
- (C3) if $\phi_\psi^\star \leq_{\mathbb{P}} \alpha$, then $(\phi_\alpha^\star)_\psi^\star \leq_{\mathbb{P}} \alpha$
- (C4) if $\phi_\psi^\star \not\leq_{\mathbb{P}} \neg\alpha$, then $(\phi_\alpha^\star)_\psi^\star \not\leq_{\mathbb{P}} \neg\alpha$

TABLE 5 Probability distribution \mathbb{P} , and the relative KM, in Example 11.

r	worlds	\mathbb{P}	κ_S
2	$\bar{p}q \quad \bar{p}\bar{q}$	0.1	3.32
1	pq	0.2	2.32
0	$p\bar{q}$	0.6	0.74

The intuitive meaning of these postulates is as for [16]:

- (C1) states that if ϕ needs to be revised by two formulae α and β , the second being more specific than the former, the first is redundant;
- (C2) states that when two contradictory pieces of evidence arrive ($\psi \leq_{\mathbb{P}} \neg\alpha$), the last one prevails;
- (C3) states that an evidence α should be retained after accommodating a more recent evidence ψ that implies α ;
- (C4) states that if α is not contradicted after seeing ψ , then it should remain uncontradicted when ψ is preceded by α itself.

For further discussions we refer the reader to [16]. Here we just aim to show which of the iterated revisions postulates hold for KM-based belief revision.

The following can be shown.

PROPOSITION 19.

Let \star be a KM-based revision operator. Then \star satisfies (C1), (C3) and (C4).

The next example shows that (C2) is not always satisfied.

EXAMPLE 11.

Let $\Sigma = \{p, q\}$, and let $\mathbb{P} = \{p\bar{q}^{0.6}, pq^{0.2}, \bar{p}q^{0.1}, \bar{p}\bar{q}^{0.1}\}$. Clearly $\kappa(\bar{p}q) = \kappa(\bar{p}\bar{q}) > \kappa(pq) > \kappa(p\bar{q})$. Now, let $\phi \equiv_{\mathbb{P}} (p \wedge q) \vee (\neg p \wedge \neg q)$, $\psi \equiv_{\mathbb{P}} p$, and $\alpha \equiv_{\mathbb{P}} (\neg p \wedge q)$. It can be verified (see Table 5) that

$$\begin{aligned} \psi &\leq_{\mathbb{P}} \neg\alpha \\ \phi &\not\leq_{\mathbb{P}} \neg\psi \\ \phi &\leq_{\mathbb{P}} \neg\alpha \end{aligned}$$

By Proposition 18, the following results follow:

- $[\phi_{\psi}^{\star}] = F([\phi]^+ \cap [\psi]^+) = \{pq\}$;
- $[\phi_{\alpha}^{\star}] = F(\min_{\kappa}([\alpha]^+)) = \{\bar{p}q\}$;
- $[(\phi_{\alpha}^{\star})_{\psi}^{\star}] = F(\min_{\kappa}([\psi]^+)) = \{p\bar{q}\}$.

Therefore, $(\phi_{\alpha}^{\star})_{\psi}^{\star} \neq_{\mathbb{P}} \phi_{\psi}^{\star}$.

From Example 11 it follows that

PROPOSITION 20.

There is a KM-based revision operator \star that does not satisfy (C2).

Not satisfying **(C2)** may appear as an undesired property, but the rationality of Postulate **(C2)** has actually been questioned in various works, as illustrated below. **(C2)** formalizes the idea that if ψ contradicts the previous information α , then ψ completely erases the effect of the ‘temporary’ presence of α in the belief set. Such an imposition, however, appears too strong in many situations, and the desirability of such a postulate has been debated in several works such as [25, 43–45, 47, 48, 66]. Let us consider an extreme example proposed by Stalnaker [66]:

[...] suppose I start out knowing almost nothing, and then receive, at once, a fat encyclopaedia of information. After I have digested it, I am told that there is a typo on p. 576: the Battle of Hastings was in 1066, rather than 1606. Should I then throw out the encyclopaedia and go back to my prior state of almost total ignorance? That does not seem reasonable. [66, p.208]

In our framework, let us consider the following, less extreme, example.

EXAMPLE 12.

Let $\Sigma = \{p, d, c\}$, where p, d and c read, respectively, ‘Bob owns a pet’, ‘Bob owns a dog’ and ‘Bob owns a cat’. Suppose now that, according to the statistics about pet ownership in Bob’s country, we have the following probability distribution:

$$\mathbb{P} = \{\bar{p}\bar{d}\bar{c}^{0.3}, p\bar{d}\bar{c}^{0.3}, p\bar{d}c^{0.25}, pdc^{0.15}\}.$$

Assume that our starting KB (our ϕ) corresponds to the proposition $\neg p$, i.e. ‘Bob does not own a pet’. Also, our probability distribution \mathbb{P} , assigning probability 0 to some worlds, implies some background information: in particular, $d \rightarrow p, c \rightarrow p$ and $p \rightarrow (d \vee c)$ (for the sake of simplicity we assume that the only kinds of pets are dogs and cats).

In the following, let $\alpha := d$ and $\psi := \neg d$.

- If we are informed that ‘Bob does not own a dog’, i.e. we ask to revise ϕ by ψ , by Equation 34, we ‘remain’ in $\bar{p}\bar{d}\bar{c}$. That is, $\phi_\psi^* \equiv_{\mathbb{P}} \phi$.
- Instead, if we are informed that ‘Bob owns a dog’, i.e. we ask to revise ϕ by α , by Equation 34, we ‘move’ to the world $p\bar{d}\bar{c}$ (the most expected/least surprising world in which d is satisfied) and, thus, $\phi_\alpha^* \equiv_{\mathbb{P}} Fp\bar{d}\bar{c}$.

If successively we are informed that ‘Bob does not own a dog’, i.e. we ask to revise ϕ_α^* by ψ , by Equation 34, we have to consider both the worlds $\bar{p}\bar{d}\bar{c}$ and $p\bar{d}\bar{c}$, which are the most expected/least surprising worlds in which Bob does not own a dog. Therefore, $(\phi_\alpha^*)_\psi^* \equiv_{\mathbb{P}} F\bar{p}\bar{d}\bar{c} \vee Fp\bar{d}\bar{c} \equiv_{\mathbb{P}} p \rightarrow (\neg d \wedge c)$. That is, we leave open the possibility that Bob owns a cat.

Therefore, ϕ_ψ^* and $(\phi_\alpha^*)_\psi^*$ are not \mathbb{P} -equivalent and we believe that the inequality is reasonable to hold in this case.

We are aware of another example of revision that does not satisfy **(C2)**, which is *amnesic revision* [60], as it has been proved by Jin and Thielscher [45, proposition 1]. However, amnesic revision corresponds to the well-known full-meet revision [2], which is generally considered a non-satisfying form of revision since it tends to drop too much information at every revision step (we have briefly discussed the associated full-meet contraction operation in Section 4).

5 Related work

Concerning KMs, to the best of our knowledge, we are not aware of other works, except [67], that address the idea of KMs. However, there is an emerging community that deals with measuring

inconsistency degrees, such as [30, 68, 69] (and references therein), which in fact, as pointed out in [67], is an area where KMs may apply.

BC has been the subject of extensive investigation (see, e.g. [3, 20, 21, 55], and references therein). Providing a comprehensive overview of the vast literature is beyond the scope of this paper; instead, we focus only on work that is closely related to our approach.

Many studies have explored the interplay between BC and probability theory (see, e.g. [5, 10–12, 35–38, 41, 51–53, 57–59, 65, 71]).²³ These works generally focus on revising probability distributions over possible worlds, which represent an agent’s beliefs, in response to BC operations. However, these approaches operate within a fundamentally different reasoning framework than ours. Specifically, in our setting, the probability distribution is fixed a priori and remains unchanged after a BC. Instead, it is used to guide the change by minimizing the informational deviation from the original knowledge base, following the *Principle of Minimum Surprise* (PMS). Also note that a KB may not impose any restriction on the probabilistic distribution.

Other model-based approaches to BC make use of a proximity relation between possible worlds (see, e.g. [4, 15, 24, 26, 42, 62, 63, 70]). In most cases, this proximity is defined syntactically, typically based on the number of propositional literals on which two worlds differ. For illustrative purposes, Appendix B provides a brief overview of the two proposals by Satoh [62] and Dalal [15]. We focus in particular on Dalal’s approach, which defines AGM belief revision operators grounded in the Hamming distance between possible worlds. As shown in Proposition 11, since Dalal’s operator satisfies the AGM postulates, corresponding KM-based revision operators can be derived by constructing suitable, ad hoc probability distributions (see Proposition A12). Nonetheless, it is crucial to emphasize that our framework differs in a fundamental way from these approaches in terms of how it interprets the *Principle of Informational Economy*. In Dalal’s and similar models, the preferred worlds are those that differ minimally in the assignment of truth values to propositional variables—essentially minimizing changes to what can be seen as atomic facts in our semantics. The underlying idea is that, when revising a knowledge base, an agent should tend to minimize the truth changes. In contrast, our approach reinterprets the Principle of Informational Economy through the lens of the PMS. Here, revision favours less surprising situations, where surprise is quantified using Shannon’s information theory. That is, instead of minimizing syntactic differences, we minimize informational content deviation.

Worth mentioning are also [39, 50], whose major objective was to illustrate a contraction that does not satisfy the recovery postulate. Interestingly, besides the recovery postulate, Levi’s contraction does neither satisfy conjunctive overlap nor conjunctive inclusion. But, the interesting point is that one may regain these two properties back by relying on the notion of *Information Value* (IV) [39, 50] that should be used to guide the selection of contraction candidates. Essentially, two variants of IVs have been identified:²⁴ (i) *strong monotonicity* of IV \mathcal{V} : if $\alpha \models \beta$ then $\mathcal{V}(\alpha) > \mathcal{V}(\beta)$; and (ii) *weak monotonicity* of IV \mathcal{V} : if $\alpha \models \beta$ then $\mathcal{V}(\alpha) \geq \mathcal{V}(\beta)$. From that, the *value-based Levi contraction* is then defined, roughly, as a disjunction of the \mathcal{V} -maximal elements in $S(\phi, \alpha)$, where this latter is the set of so-called *saturatable contractions* that consist of formulae ψ weaker than ϕ such that $\psi \wedge \neg\alpha$ is maximally consistent.²⁵ Of course, the notion of IV resembles axiom (E) of KMs in [67]²⁶ in the sense that the latter are weak monotonic IVs, though the opposite is not true. Also, the overall

²³There is also related work on BC and possibility theory, such as [14, 17, 56].

²⁴ $\mathcal{V}(\alpha)$ is real-valued.

²⁵In our setting, this set may be defined as the set $S(\phi, \alpha) = \{\psi \mid \phi \leq_{\mathbb{P}} \psi \text{ and } \psi \wedge \neg\alpha \text{ is maximal } \mathbb{P}\text{-consistent}\}$, where β is maximal if there is no \mathbb{P} -consistent γ with $\gamma <_{\mathbb{P}} \beta$.

²⁶If $\phi \models \psi$ then $\kappa(\psi) \leq \kappa(\phi)$.

definition of value-based Levi contraction is based on an IV maximization selection process among saturatable contractions, while we rely on a KM minimization selection process over remainders (w.r.t. a priori fixed probabilistic distribution). Nevertheless, at the time of writing, the relationship between Levi's value-based Levi contraction and our KM-based one is still unclear and we will address it in future work, together with other alternative constructions of BC operators. Let us also mention that we also tried to consider the \max_{κ} operator (with obvious definition) in place of the \min_{κ} operator in Definition 1, though it turned out to be questionable: as shown in Equation 9, KMs reward the less probable options: the less expected is a piece of information, the 'more information it contains' from an information-theoretic point of view. But then, regarding Example 7, if we select the remainders with highest KM, we would have given the priority to a world in which the non-flying birds would have been penguins (i.e. $bp\bar{o}\bar{f}w^{0.15}$) rather than to the more probable world $(b\bar{p}\bar{f}w)^{0.4}$.

6 Conclusions and future work

We have introduced novel quantitative BC operators rooted in KMs, aimed at minimizing the (information theoretic) surprise of the information conveyed by the modified belief.

In particular, our contributions encompass several key aspects: (i) we have generalized KMs for their application to BC; (ii) we have proposed KM-based BC operators that adhere to the AGM postulates and (iii) we have demonstrated that any BC operator meeting the AGM postulates can be represented as a KM-based BC operator, regulated by a specific probability distribution over the worlds. We have additionally introduced quantitative metrics that capture the information change of BC operators. We also have given a succinct look into the problem of iterated revision [16], and have also illustrated how one can construct a non-recovery-postulate-compliant contraction operator from our KM-based framework, by focusing on the so-called severe withdrawal [61] model as illustrative example.

Topics for future research may include the following (besides those already mentioned in the paper here and there): (i) to investigate and/or apply the notion of KMs in other contexts such as First-Order Language, or more in general languages that do not have the finite-model property, to various forms of paraconsistent logics [1]), e.g. measuring the degree of inconsistency [30]), to non-monotonic logics [6, 69], to conditional logics [8] and to uncertain, many-valued and mathematical fuzzy logics [18]; (ii) to investigate on the application of KMs to other approaches related to BC, e.g. belief update, inclusive to non-recovery postulate-compliant variants and (iii) to investigate forms of BC obtained by combining the use of KMs with some kind of measure of similarity/semantic closeness among worlds, such as Dalal distance [15].

Acknowledgements

This work has been partially supported by the H2020 STARWARS Project (GA No. 101086252), type of action HORIZON TMA MSCA Staff Exchanges. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union, which cannot be held responsible for them.

References

- [1] J. M. Abe, S. Akama and K. Nakamatsu. *Introduction to Annotated Logics - Foundations for Paracomplete and Paraconsistent Reasoning, Volume 88 of Intelligent Systems Reference Library*. Springer, 2015.

- [2] C. Alchourrón, P. Gärdenfors and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, **50**, 510–530, 1985.
- [3] T. I. Aravanis and P. Peppas. Theory-relational belief revision. *Annals of Mathematics and Artificial Intelligence*, **90**, 573–594, 2022.
- [4] A. Borgida.. Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, **10**, 565–603, 1985.
- [5] C. Boutilier.. On the revision of probabilistic belief states. *Notre Dame Journal of Formal Logic*, **36**, 158–183, 1995.
- [6] G. Brewka, I. Niemelä and M. Truszczyński.. Nonmonotonic reasoning. In *Handbook of Knowledge Representation, Volume 3 of Foundations of Artificial Intelligence*, F. van Harmelen, V. Lifschitz and B. W. Porter, eds, pp. 239–284. Elsevier, 2008.
- [7] T. Caridroit, S. Konieczny and P. Marquis.. Contraction in propositional logic. *International Journal of Approximate Reasoning*, **80**, 428–442, 2017.
- [8] G. Casini and U. Straccia.. A general framework for modelling conditional reasoning—preliminary report. In *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel. July 31–August 5, 2022*, G. Kern-Isberner, G. Lakemeyer and T. Meyer, eds, IJCAI Organization, 2022.
- [9] C. G. Chakrabarti and I. Chakrabarty.. Shannon entropy: axiomatic characterization and application. *International Journal of Mathematics and Mathematical Sciences*, **2005**, 2847–2854, 2005.
- [10] H. Chan and A. Darwiche.. A distance measure for bounding probabilistic belief change. *International Journal of Approximate Reasoning*, **38**, 149–174, 2005.
- [11] H. Chan and A. Darwiche.. On the revision of probabilistic beliefs using uncertain evidence. *Artificial Intelligence*, **163**, 67–90, 2005.
- [12] K. Chhogyal, A. C. Nayak, Z. Zhuang and A. Sattar.. Probabilistic belief contraction using argumentation. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25–31, 2015*, Q. Yang and M. J. Wooldridge, eds, pp. 2854–2860. AAAI Press, **2015**.
- [13] N. Creignou, O. Papini, R. Pichler and S. Woltran.. Belief revision within fragments of propositional logic. *Journal of Computer and System Sciences*, **80**, 427–449, 2014.
- [14] C. da Costa Pereira and A. G. B. Tettamanzi.. A syntactic possibilistic belief change operator: theory and empirical study. *Web Intell. Agent Syst.*, **12**, 155–169, 2014.
- [15] M. Dalal.. Investigations into a theory of knowledge base revision: preliminary report. In *Proceedings of the Seventh AAAI National Conference on Artificial Intelligence, AAAI’88*, Howard E. Shrobe, Tom M. Mitchell, and Reid G. Smith, eds, pp 475–479. AAAI Press, 1988.
- [16] A. Darwiche and J. Pearl.. On the logic of iterated belief revision. *Artificial Intelligence*, **89**, 1–29, 1997.
- [17] D. Dubois and H. Prade.. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *International Journal of Approximate Reasoning*, **17**, 295–324, 1997.
- [18] D. Dubois and H. Prade.. Possibility theory, probability theory and multiple-valued logics: a clarification. *Annals of Mathematics and Artificial Intelligence*, **32**, 35–66, 2001.
- [19] E. Fermé, M. Garapa, A. Nayak and M. D. Luís Reis.. Relevance, recovery and recuperation: a prelude to ring withdrawal. *International Journal of Approximate Reasoning*, **166**, 109108, 2024.
- [20] E. L. Fermé and S. O. Hansson, eds. AGM 25 years: twenty-five years of research in belief change. *Journal of Philosophical Logic*, **40**, 295–331, 2011.

- [21] E. L. Fermé and S. O. Hansson, eds. *Belief Change: Introduction and Overview*. Springer International Publishing, 2018.
- [22] E. L. Fermé and A. Herzig., and M. V. Martinez. On the logic of theory base change: reformulation of belief bases. In *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25–March 4, 2025, Philadelphia, PA, USA*, T. Walsh, J. Shah and Z. Kolter, eds. pp. 14922–14929. AAAI Press, **39**, 14922, 14929.
- [23] E. L. Fermé and R. O. Rodríguez.. Semi-contraction: axioms and construction. *Notre Dame Journal of Formal Logic*, **39**, 332–345, 1998.
- [24] K. D. Forbus.. Introducing actions into qualitative simulation. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence. Detroit, MI, USA, August 1989*, N. S. Sridharan, eds, pp. 1273–1278. Morgan Kaufmann, 1989.
- [25] M. Freund and D. Lehmann.. *Belief Revision and Rational Inference*, arxiv, Cornell University, 2002.
- [26] A. Fusaoka.. On a linear representation theory for quantitative belief change. In *AI 2002: Advances in Artificial Intelligence, 15th Australian Joint Conference on Artificial Intelligence, Canberra, Australia, December 2–6, 2002, Proceedings, volume 2557 of Lecture Notes in Computer Science*, B. McKay and J. K. Slaney, eds, pp. 48–59. Springer, 2002.
- [27] P. Gärdenfors.. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.
- [28] P. Gärdenfors and D. Makinson.. Nonmonotonic inference based on expectations. *Artificial Intelligence*, **65**, 197–245, 1994.
- [29] P. Gärdenfors.. *Belief Revision*. Cambridge University Press, 2003.
- [30] J. Grant and F. Parisi.. General information spaces: measuring inconsistency, rationality postulates, and complexity. *Annals of Mathematics and Artificial Intelligence*, **90**, 235–269, 2022.
- [31] A. Grove.. Two modellings for theory change. *Journal of Philosophical Logic*, **17**, 157–170, 1988.
- [32] I. Hacking.. *An Introduction to Propability and Inductive Logic*. Cambridge University Press, 2001.
- [33] Y. Joseph.. *Halpern. Reasoning about Uncertainty*. The MIT Press, 2017.
- [34] D. Hankerson, P. D. Johnson and G. A. Harris.. *Introduction to Information Theory and Data Compression*, 1st edn. CRC Press, Inc., 1998.
- [35] S. O. Hansson.. Revising probabilities and full beliefs. *Journal of Philosophical Logic*, **49**, 1005–1039, 2020.
- [36] S. O. Hansson.. A characterization of probability-based dichotomous belief revision. *Studia Logica*, **110**, 511–543, 2022.
- [37] S. O. Hansson.. A basis for AGM revision in Bayesian probability revision. *Journal of Philosophical Logic*, **52**, 1535–1559, 2023.
- [38] S. O. Hansson.. Iterated AGM revision based on probability revision. *Journal of Logic, Language and Information*, **32**, 657–675, 2023.
- [39] S. O. Hansson and E. J. Olsson.. Levi contractions and AGM contractions: a comparison. *Notre Dame Journal of Formal Logic*, **36**, 103–119, 1995.
- [40] G. Harman.. *Change in View: Principles of Reasoning*. MIT Press, Cambridge, MA, USA, 1986.
- [41] J. Hawthorne and D. Makinson.. The quantitative/qualitative watershed for rules of uncertain inference. *Studia Logica*, **86**, 247–297, 2007.
- [42] S. Hegner.. Specification and implementation of programs for updating incomplete informa-

- tion databases. In *Proceedings of the Sixth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, PODS '87*, Moshe Y. Vardi, ed., pp. 146–158, New York, NY, USA. Association for Computing Machinery, 1987.
- [43] A. Herzig, S. Konieczny and L. Perrussel.. On iterated revision in the AGM framework. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, T. D. Nielsen and N. L. Zhang, eds, pp. 477–488, 2003. Springer Berlin Heidelberg.
- [44] Y. Jin and M. Thielscher.. Iterated belief revision, revised. In *Proceedings of the 19th international joint conference on artificial intelligence, IJCAI'05*, Leslie Pack Kaelbling and Alessandro Saffiotti, eds., page 478–483, Morgan Kaufmann Publishers Inc, 2005.
- [45] Y. Jin and M. Thielscher.. Iterated belief revision, revised. *Artificial Intelligence*, **171**, 1–18, 2007.
- [46] H. Katsuno and A. Mendelzon.. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, **52**, 263–294, 1991.
- [47] S. Konieczny and R. P. Pérez.. A framework for iterated revision. *Journal of Applied Non-Classical Logics*, **10**, 339–367, 2000.
- [48] D. Lehmann.. Belief revision, revised. In *Proceedings of the 14th international joint conference on artificial intelligence - volume 2, IJCAI'95*, Chris S. Mellish, ed., pp. 1534–1540, Morgan Kaufmann Publishers Inc, 1995.
- [49] I. Levi.. Subjunctives, dispositions and chances. *Synthese*, **34**, 423–455, 1977.
- [50] I. Levi.. *The Fixation of Belief and its Undoing*. Cambridge University Press, 1991.
- [51] S. Lindström and W. Rabinowicz.. On probabilistic representation of non-probabilistic belief revision. *Journal of Philosophical Logic*, **18**, 69–101, 1989.
- [52] J. Ma, W. Liu and S. Benferhat.. A belief revision framework for revising epistemic states with partial epistemic states. *International Journal of Approximate Reasoning*, **59**, 20–40, 2015.
- [53] D. C. Makinson.. Conditional probability in the light of qualitative belief change. *Journal of Philosophical Logic*, **40**, 121–153, 2011.
- [54] J. Minker.. On indefinite data bases and the closed world assumption. In *Proc. of the 6th Conf. on Automated Deduction (CADE-82), volume 138 of Lecture Notes in Computer Science*, Ewing Lusk, Ross Overbeek, eds., Springer-Verlag, 1982.
- [55] P. Peppas.. Belief revision. In *Handbook of Knowledge Representation, volume 3 of Foundations of Artificial Intelligence*, F. van Harmelen, V. Lifschitz and B. W. Porter, eds, pp. 317–359. Elsevier, 2008.
- [56] G. Qi and K. Wang.. Conflict-based belief revision operators in possibilistic logic. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22–26, 2012*, J. Hoffmann and B. Selman, eds, pp. 800–806. AAAI Press, 2012.
- [57] R. Ramachandran, A. Ramer and A. C. Nayak.. Probabilistic belief contraction. *Minds and Machines*, **22**, 325–351, 2012.
- [58] G. Rens, T. A. Meyer, G. Kern-Isberner and A. Nayak.. Probabilistic belief revision via similarity of worlds modulo evidence. In *KI 2018: Advances in Artificial Intelligence—41st German Conference on AI, Berlin, Germany, September 24–28, 2018, Proceedings, volume 11117 of Lecture Notes in Computer Science*, F. Trollmann and A.-Y. Turhan, eds, pages 343–356. Springer, 2018.
- [59] G. B. Rens and T. Meyer.. A new approach to probabilistic belief change. In *Proceedings of the Twenty-Eighth International Florida Artificial Intelligence Research Society Conference, FLAIRS 2015, Hollywood, Florida, USA, May 18–20, 2015*, I. Russell and W. Eberle, eds, pp. 581–587. AAAI Press, 2015.
- [60] H. Rott.. Two dogmas of belief revision. *Journal of Philosophy*, **97**, 503, 2000.

- [61] H. Rott and M. Pagnucco.. Severe withdrawal (and recovery). *Journal of Philosophical Logic*, **28**, 501–547, 1999.
- [62] K. Satoh.. *Nonmonotonic Reasoning by Minimal Belief Revision*. In *Fifth Generation Computer Systems*, Institute for New Generation Computer Technology (ICOT) in Tokyo, Japan, 1988.
- [63] N. Schwind, S. Konieczny and R. P. Pérez.. On paraconsistent belief revision in LP. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22–March 1, 2022*, Sven Koenig et al., eds., pp 5879–5887. AAAI Press, 2022.
- [64] C. E. Shannon.. A mathematical theory of communication a mathematical theory of information. *The Bell System Technical Journal*, **27**, 379–423, 1948.
- [65] W. Spohn.. *The Laws of Belief: Ranking Theory and its Philosophical Applications*. Oxford University Press, 2012.
- [66] R. Stalnaker.. Iterated belief revision. *Erkenntnis*, **70**, 189–209, 2009.
- [67] U. Straccia.. How much knowledge is in a knowledge base? Introducing knowledge measures (preliminary report). In *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI-20)*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pp. 905–912. Giuseppe De Giacomo, Alejandro Catala, Bistra Dilkina, Michela Milano, Senén Barro, Alberto Bugarín, and Jérôme Lang. eds., IOS Press, 2020.
- [68] M. Thimm and J. P. Wallner.. On the complexity of inconsistency measurement. *Artificial Intelligence*, **275**, 411–456, 2019.
- [69] M. Ulbricht, M. Thimm and G. Brewka.. Handling and measuring inconsistency in non-monotonic logics. *Artificial Intelligence*, **286**, 103344, 2020.
- [70] M. Winslett.. *Updating Logical Databases*. *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1990.
- [71] Z. Zhuang, J. P. Delgrande, A. C. Nayak and A. Sattar.. A unifying framework for probabilistic belief revision. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017* C. Sierra, ed, pp. 1370–1376., IJCAI Organization, 2017.

Appendix A. Proofs

PROPOSITION 1.

Consider any formulae ϕ and ψ w.r.t. Σ , and any probability distribution \mathbb{P} over Σ . If $\phi \triangleleft_{\mathbb{P}} \psi$ then $\mathbb{P}(\phi) \triangleleft \mathbb{P}(\psi)$, for $\triangleleft \in \{\leq, <\}$.

PROOF. Consider the case $\phi \leq_{\mathbb{P}} \psi$. By definition of \mathbb{P} -entailment, $[\phi \wedge \neg \mathbb{P}0] \subseteq [\psi]$ holds. Therefore, as we rule out 0 probability worlds, $\mathbb{P}(\phi) = \mathbb{P}(\phi \wedge \neg \mathbb{P}0) \leq \mathbb{P}(\psi)$.

The strict case follows from the fact that we rule out in $\phi \wedge \neg \mathbb{P}0$ all worlds with 0 probability (see also Remark 3). \square

PROPOSITION 4.

For any formula ϕ , given the uniform probability distribution \mathbb{P}_{Σ}^u over Σ , $\kappa_S(\phi) = \kappa_h(\phi)$.

PROOF. If ϕ is unsatisfiable then $\mathbb{P}^u(\phi) = 0$ and, thus, by definition $\kappa_S(\phi) = \infty = \kappa_h(\phi)$. Otherwise, $\mathbb{P}^u = 1/2^{|\Sigma|}$ and $\mathbb{P}^u(\phi) = |[[\phi]]|/2^{|\Sigma|}$ holds. So, $\kappa_S(\phi) = -\log_2 \mathbb{P}(\phi) = -\log_2 \frac{|[\phi]|}{2^{|\Sigma|}} = |\Sigma| - \log_2 |[\phi]| = \kappa_h(\phi)$. \square

PROPOSITION 6.

Consider any probability distribution \mathbb{P} over \mathcal{W}_Σ and any formulae ϕ and α such that $\phi \leq_{\mathbb{P}} \alpha$ and $\top \not\leq_{\mathbb{P}} \alpha$. Then $\psi \in \phi \perp \alpha$ iff

$$\psi \equiv F([\phi]_\Sigma^+ \cup \{w\})$$

for some $w \in [-\alpha]_\Sigma^+$. Furthermore, $\mathbb{P}(\psi) = \mathbb{P}(F([\phi]_\Sigma^+)) + \mathbb{P}(w)$. On the other hand, if $\phi \not\leq_{\mathbb{P}} \alpha$ or $\top \leq_{\mathbb{P}} \alpha$ then $\phi \perp \alpha = \{F([\phi]^+)\}$.

PROOF. Assume $\phi \leq_{\mathbb{P}} \alpha$ and $\top \not\leq_{\mathbb{P}} \alpha$. Clearly, $F([\phi]_\Sigma^+ \cup \{w\})$ is a possible reminder as $\phi \leq_{\mathbb{P}} F([\phi]_\Sigma^+ \cup \{w\})$ and $F([\phi]_\Sigma^+ \cup \{w\}) \not\leq_{\mathbb{P}} \alpha$ holds. Now, assume to the contrary that $F([\phi]_\Sigma^+ \cup \{w\})$ is not a reminder. Therefore, there is another $\psi \in \phi \perp \alpha$ s.t. $\psi <_{\mathbb{P}} F([\phi]^+ \cup \{w\})$. But, that to be the case, it must be $[\phi]^+ \subset [\psi]^+ \subset ([\phi]^+ \cup \{w\})$, that cannot be the case. Therefore, each reminder corresponds to a formula $F([\phi]_\Sigma^+ \cup \{w\})$ with $w \in [-\alpha]_\Sigma^+$, which concludes this part. Additionally, as $w \notin [\phi]_\Sigma^+$, $\mathbb{P}(\psi) = \mathbb{P}(F([\phi]_\Sigma^+)) + \mathbb{P}(w)$ holds.

Eventually, if $\phi \not\leq_{\mathbb{P}} \alpha$ or $\top \leq_{\mathbb{P}} \alpha$ holds then $\phi \perp \alpha = \{F([\phi]^+)\}$ by the definition of the reminder set. \square

PROPOSITION 7.

Consider any probability distribution \mathbb{P} over \mathcal{W} , any KM κ and formulae ϕ and α . Then

$$\phi_\alpha^{\dot{\kappa}} \equiv \begin{cases} F([\phi]^+ \cup \min_\kappa([-\alpha]^+)) & \text{if } \phi \leq_{\mathbb{P}} \alpha \\ F([\phi]^+) & \text{otherwise,} \end{cases}$$

where $\min_\kappa([-\alpha]^+) = \{w \in \mathcal{W} \mid w \in [-\alpha]^+ \text{ and there is no } w' \in \mathcal{W} \text{ s.t. } w' \in [-\alpha]^+ \text{ and } \kappa(w') < \kappa(w)\}$.

PROOF. If $\phi \not\leq_{\mathbb{P}} \alpha$ then by Proposition 6, $\phi \perp \alpha = \{F([\phi]^+)\}$ and, thus, $\phi_\alpha^{\dot{\kappa}} = F([\phi]^+)$. Now, let us assume $\phi \leq_{\mathbb{P}} \alpha$ instead. There are two cases.

Case $\top \leq_{\mathbb{P}} \alpha$. If $\top \leq_{\mathbb{P}} \alpha$, then $\min_\kappa([-\alpha]^+) = \emptyset$. Again, by Proposition 6, $\phi \perp \alpha = \{F([\phi]^+)\}$ and, thus, $\phi_\alpha^{\dot{\kappa}} = F([\phi]^+) = F([\phi]^+ \cup \emptyset)$.

Case $\top \not\leq_{\mathbb{P}} \alpha$. Let us prove that $F([\phi]^+ \cup \min_\kappa([-\alpha]^+)) \equiv \bigvee \min_\kappa(\phi \perp \alpha)$. It is sufficient to prove that each formula in $\min_\kappa(\phi \perp \alpha)$ corresponds to a formula $F([\phi]^+ \cup \{w\})$ with $w \in \min_\kappa([-\alpha]^+)$.

At first we prove that a formula $F([\phi]^+ \cup \{w\})$ with $w \in \min_\kappa([-\alpha]^+)$ is in $\min_\kappa(\phi \perp \alpha)$. Now, $F([\phi]^+ \cup \{w\})$ is a reminder by Proposition 6, i.e. $F([\phi]^+ \cup \{w\}) \in \phi \perp \alpha$, and we have to prove that $F([\phi]^+ \cup \{w\}) \in \min_\kappa(\phi \perp \alpha)$. Assume it is not the case. That implies that there is $\psi \in \phi \perp \alpha$ s.t. $\kappa(\psi) < \kappa(F([\phi]^+ \cup \{w\}))$ and, thus, $\mathbb{P}(\psi) > \mathbb{P}(F([\phi]^+ \cup \{w\})) = \mathbb{P}(F([\phi]^+) + \mathbb{P}(w))$ (by Proposition 6). But, $\psi \in \phi \perp \alpha$ means that, by Proposition 6, $\psi \equiv F([\phi]^+ \cup \{w'\})$ for some $w' \in [-\alpha]^+$, $\mathbb{P}(\psi) = \mathbb{P}(F([\phi]^+) + \mathbb{P}(w'))$ and, thus, $\mathbb{P}(w') > \mathbb{P}(w)$, that cannot be the case as $w \in \min_\kappa([-\alpha]^+)$ and, thus, $\mathbb{P}(w) \geq \mathbb{P}(w')$ has to hold.

To prove that every formula in $\min_\kappa(\phi \perp \alpha)$ indeed corresponds to some formula $F([\phi]^+ \cup \{w\})$, with $w \in \min_\kappa([-\alpha]^+)$, consider that every formula in $\min_\kappa(\phi \perp \alpha)$ must correspond to a formula $F([\phi]^+ \cup \{w\})$ with $w \in [-\alpha]^+$, and for it to be in $\min_\kappa(\phi \perp \alpha)$ it must be one of such formulae with the highest probability w.r.t. \mathbb{P} , hence w must be in $\min_\kappa([-\alpha]^+)$.

Therefore, $\min_\kappa(\phi \perp \alpha) = \{F([\phi]^+ \cup \{w\}) \mid w \in \min_\kappa([-\alpha]^+)\}$. That is, $F([\phi]^+ \cup \min_\kappa([-\alpha]^+)) \equiv \bigvee \min_\kappa(\phi \perp \alpha)$. \square

COROLLARY 1.

Consider any probability distribution \mathbb{P} over \mathcal{W} , any KM κ and formulae ϕ and α . Then

$$IC^{\dot{\div}}(\phi, \alpha) = \begin{cases} -\log_2 \left(1 + \frac{\mathbb{P}(F \min_{\kappa}([\neg\alpha]^+))}{\mathbb{P}(\phi)} \right) & \text{if } \phi \leq_{\mathbb{P}} \alpha \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. Consider Proposition 7 and Equation 19. If $\phi \not\leq_{\mathbb{P}} \alpha$ then $IC^{\dot{\div}}(\phi, \alpha) = \kappa(F([\phi]^+)) - \kappa(\phi) = \kappa(\phi) - \kappa(\phi) = 0$. Otherwise ($\phi \leq_{\mathbb{P}} \alpha$), let us note that

$$\begin{aligned} \mathbb{P}(F([\phi]^+ \cup \min_{\kappa}([\neg\alpha]^+))) &= \mathbb{P}(\phi) + \mathbb{P}(\min_{\kappa}([\neg\alpha]^+)) \\ &= \mathbb{P}(\phi) \cdot \left(1 + \frac{\mathbb{P}(\min_{\kappa}([\neg\alpha]^+))}{\mathbb{P}(\phi)} \right) \end{aligned}$$

and, thus,

$$\begin{aligned} IC^{\dot{\div}}(\phi, \alpha) &= \kappa(\phi_{\alpha}^{\dot{\div}}) - \kappa(\phi) \\ &= -\log_2 \mathbb{P}(F([\phi]^+ \cup \min_{\kappa}([\neg\alpha]^+))) + \log_2 \mathbb{P}(\phi) \\ &= -\log_2 \mathbb{P}(\phi) \cdot \left(1 + \frac{\mathbb{P}(\min_{\kappa}([\neg\alpha]^+))}{\mathbb{P}(\phi)} \right) + \log_2 \mathbb{P}(\phi) \\ &= -\log_2 \mathbb{P}(\phi) - \log_2 \left(1 + \frac{\mathbb{P}(\min_{\kappa}([\neg\alpha]^+))}{\mathbb{P}(\phi)} \right) + \log_2 \mathbb{P}(\phi) \\ &= -\log_2 \left(1 + \frac{\mathbb{P}(\min_{\kappa}([\neg\alpha]^+))}{\mathbb{P}(\phi)} \right). \end{aligned}$$

□

PROPOSITION 8.

For all probability distributions \mathbb{P} and KM-contraction operator $\dot{\div}_{\kappa}$ based on \mathbb{P} , $\dot{\div}_{\kappa}$ is an AGM contraction operator, i.e. it satisfies postulates $(\dot{\div}1)$ - $(\dot{\div}7)$.

PROOF. We have to prove that all AGM postulates are satisfied:

- $(\dot{\div}1)$ $\phi \leq_{\mathbb{P}} \psi$ for every $\psi \in \phi \perp \alpha$, hence $\phi \leq_{\mathbb{P}} \bigvee \min_{\kappa}(\phi \perp \alpha)$.
- $(\dot{\div}2)$ If $\phi \not\leq_{\mathbb{P}} \alpha$, then $\min_{\kappa}(\phi \perp \alpha)$ contains only $F([\phi]^+)$, which is \mathbb{P} -equivalent to ϕ .
- $(\dot{\div}3)$ Assume $\top \not\leq_{\mathbb{P}} \alpha$. Then $\psi \not\leq_{\mathbb{P}} \alpha$ for all $\psi \in \phi \perp \alpha$, which, by the monotonicity of $\leq_{\mathbb{P}}$, implies $\bigvee \min_{\kappa}(\phi \perp \alpha) \not\leq_{\mathbb{P}} \alpha$.
- $(\dot{\div}4)$ It is immediate from the definition of $\phi \perp \alpha$.
- $(\dot{\div}5)$ It suffices to prove that $\phi_{\alpha}^{\dot{\div}_{\kappa}} \leq_{\mathbb{P}} \alpha \rightarrow \phi$. To this end, we prove that for all $\psi \in \min_{\kappa}(\phi \perp \alpha)$, $\psi \leq_{\mathbb{P}} \alpha \rightarrow \phi$. So, assume that it is not the case. Clearly, $\psi \wedge (\alpha \rightarrow \phi) \leq_{\mathbb{P}} \psi$. But, $\psi \not\leq_{\mathbb{P}} \alpha \rightarrow \phi$ and, thus, $\psi \not\leq_{\mathbb{P}} \psi \wedge (\alpha \rightarrow \phi)$. Therefore, by Proposition 2, $\kappa(\psi \wedge (\alpha \rightarrow \phi)) > \kappa(\psi)$. As $\psi \in \min_{\kappa}(\phi \perp \alpha)$, $\psi \wedge (\alpha \rightarrow \phi) \notin \phi \perp \alpha$ follows. But, $\phi \leq_{\mathbb{P}} \psi \wedge (\alpha \rightarrow \phi)$, so $\psi \wedge (\alpha \rightarrow \phi) \leq_{\mathbb{P}} \alpha$. That is, $\psi \leq_{\mathbb{P}} (\alpha \rightarrow \phi) \rightarrow \alpha$. Since $(\alpha \rightarrow \phi) \rightarrow \alpha \equiv \alpha$, we conclude $\psi \leq_{\mathbb{P}} \alpha$, against the assumption that $\psi \in \phi \perp \alpha$. Therefore, for all $\psi \in \min_{\kappa}(\phi \perp \alpha)$, $\psi \leq_{\mathbb{P}} \alpha \rightarrow \phi$, and consequently $\phi_{\alpha}^{\dot{\div}_{\kappa}} \leq_{\mathbb{P}} \alpha \rightarrow \phi$.

- (÷6) By definition $\phi_{\alpha \wedge \beta}^{\ddot{\kappa}} = \bigvee \min_{\kappa}(\phi \perp (\alpha \wedge \beta))$, while $\phi_{\alpha}^{\ddot{\kappa}} = \bigvee \min_{\kappa}(\phi \perp \alpha)$ and $\phi_{\beta}^{\ddot{\kappa}} = \bigvee \min_{\kappa}(\phi \perp \beta)$. So, if $\psi \in \phi \perp (\alpha \wedge \beta)$ then either $\psi \not\leq_{\mathbb{P}} \alpha$ or $\psi \not\leq_{\mathbb{P}} \beta$, and that there is no $\psi' <_{\mathbb{P}} \psi$ s.t. $\phi \leq_{\mathbb{P}} \psi'$, $\psi' \leq_{\mathbb{P}} \psi$, and either $\psi' \not\leq_{\mathbb{P}} \alpha$ or $\psi' \not\leq_{\mathbb{P}} \beta$, respectively. That is, $\psi \in \phi \perp \alpha$ or $\psi \in \phi \perp \beta$. Hence, for all $\psi \in \phi \perp (\alpha \wedge \beta)$, either $\psi \in \phi \perp \alpha$ or $\psi \in \phi \perp \beta$. This leaves us with three possible options: $\min_{\kappa}(\phi \perp (\alpha \wedge \beta)) = \min_{\kappa}(\phi \perp \alpha)$, $\min_{\kappa}(\phi \perp (\alpha \wedge \beta)) = \min_{\kappa}(\phi \perp \beta)$ or $\min_{\kappa}(\phi \perp (\alpha \wedge \beta)) = \min_{\kappa}(\phi \perp \alpha) \cup \min_{\kappa}(\phi \perp \beta)$, depending whether, given any $\psi \in \min_{\kappa}(\phi \perp \alpha)$ and any $\psi' \in \min_{\kappa}(\phi \perp \beta)$, $k(\psi') > \kappa(\psi)$, $\kappa(\psi') < k(\psi)$ or $\kappa(\psi) = \kappa(\psi')$, respectively. In any of such three cases, (÷6) is satisfied.
- (÷7) If $\phi_{\alpha \wedge \beta}^{\ddot{\kappa}} \not\leq_{\mathbb{P}} \alpha$, then there is a $\gamma \in \min_{\kappa}(\phi \perp (\alpha \wedge \beta))$ s.t. $\gamma \not\leq_{\mathbb{P}} \alpha$. This implies that all the $\gamma \in \min_{\kappa}(\phi \perp \alpha)$ are in $\min_{\kappa}(\phi \perp (\alpha \wedge \beta))$. That is, $\min_{\kappa}(\phi \perp \alpha) \subseteq \min_{\kappa}(\phi \perp (\alpha \wedge \beta))$. This in turn implies that $\bigvee \min_{\kappa}(\phi \perp \alpha) \models \bigvee \min_{\kappa}(\phi \perp (\alpha \wedge \beta))$. Therefore, $\phi_{\alpha}^{\ddot{\kappa}} \leq_{\mathbb{P}} \phi_{\alpha \wedge \beta}^{\ddot{\kappa}}$ holds. \square

PROPOSITION 10.

For any formula ϕ , set of worlds \mathcal{W} and ϕ -faithful ranking r over \mathcal{W} , there exists an r -faithful probability distribution \mathbb{P} over \mathcal{W} .

PROOF. It suffices to determine \mathbb{P} given any r . For example, we may define a function $f : \mathcal{W} \mapsto \mathbb{N}$ with $f(w) = r(w) + 1$ and then define $\mathbb{P}(w)$ as follows: let $m = 1 + \max_{w' \in \mathcal{W}} r(w')$. Then $\mathbb{P}(w) = \frac{m - r(w)}{\sum_{w' \in \mathcal{W}} f(w')}$. It is easy to check that \mathbb{P} is an r -faithful probability distribution, i.e. $\sum_{w \in \mathcal{W}} \mathbb{P}(w) = 1$ and that the two conditions for r -faithfulness are satisfied. \square

PROPOSITION 11.

Let ϕ be any formula and let \div be an AGM-contraction. Then there is a KM contraction \div_{κ} s.t., for all α , $\phi_{\alpha}^{\ddot{\kappa}} \equiv \phi_{\alpha}^{\ddot{\kappa}}$ holds.

PROOF. Given ϕ and \div , by Proposition 9 we know that there is a faithful ranking r over a set of worlds \mathcal{W} s.t. the contraction \div_r associated with r (see Equation 25) corresponds to \div , i.e. for any α , $\phi_{\alpha}^{\ddot{\kappa}} \equiv \phi_{\alpha}^{\ddot{\kappa}r}$. We now define an r -faithful probability distribution \mathbb{P} over \mathcal{W} , and on top of it we generate a KM κ . Proposition 10 guarantees that such a probability distribution exists.

Now, we have to prove that \div_r and \div_{κ} correspond to each other. If $\phi \not\leq_{\mathbb{P}} \alpha$ then $\phi \not\leq_{\mathbb{P}} \alpha$, as there are no zero probability worlds, and, thus, by the vacuity postulate $\phi_{\alpha}^{\ddot{\kappa}r} \equiv \phi \equiv_{\mathbb{P}} \phi_{\alpha}^{\ddot{\kappa}}$. Otherwise, clearly, for any $w, w' \in \mathcal{W}$, $r(w) \leq r(w')$ iff $\mathbb{P}(w) \geq \mathbb{P}(w')$ iff $\kappa(w) \leq \kappa(w')$, by (KM2). Consequently, $\phi_{\alpha}^{\ddot{\kappa}r} \equiv F([\phi] \cup \min_r([\neg\alpha]))$ means that $\phi_{\alpha}^{\ddot{\kappa}r} \equiv F([\phi]^+ \cup \min_{\kappa}([\neg\alpha]^+)) \equiv \phi_{\alpha}^{\ddot{\kappa}}$, by Proposition 7. \square

COROLLARY 2.

Consider any probability distribution \mathbb{P} over \mathcal{W} and any formulae ϕ and α . Then

$$IC^{\ddot{\kappa}}(\phi, \alpha) = \begin{cases} -\log_2 \frac{\mathbb{P}(F\sigma(\neg\alpha))}{\mathbb{P}(\phi)} & \text{if } \top \not\leq_{\mathbb{P}} \alpha \\ 0 & \text{otherwise.} \end{cases}$$

PROOF. The case $\top \leq_{\mathbb{P}} \alpha$ is straightforward. So, let us consider the case $\top \not\leq_{\mathbb{P}} \alpha$. By definition of the loss defined as Equation 30 and Proposition 13, we have $IC^{\ddot{\kappa}}(\phi, \alpha) = -\log_2 \mathbb{P}(F\sigma(\neg\alpha)) + \log_2 \mathbb{P}(\phi) = -\log_2 \frac{\mathbb{P}(F\sigma(\neg\alpha))}{\mathbb{P}(\phi)}$. \square

PROPOSITION 14.

Consider any probability distribution \mathbb{P} and any KM-contraction operator \div_{κ} based on \mathbb{P} . Then the function \div_{κ} is a severe withdrawal operator, i.e. it satisfies postulates $(\div 1)$ - $(\div 4)$, $(\div 6a)$ and $(\div 7)$.

PROOF. We have to prove that all postulates of severe withdrawal are satisfied by \div_{κ} :

$(\div 1)$. Follows immediately from Proposition 13.

$(\div 2)$. If $\phi \not\leq_{\mathbb{P}} \alpha$ the $\sigma(\neg\alpha) = [\phi]^+$. Therefore, if $\phi \not\leq_{\mathbb{P}} \alpha$ or $\top \leq_{\mathbb{P}} \alpha$, then, by Proposition 13, $[\phi_{\alpha}^{\div_{\kappa}}]^+ = [\phi]^+$ and, thus, $\phi_{\alpha}^{\div_{\kappa}} \equiv_{\mathbb{P}} \phi$.

$(\div 3)$. If $\top \not\leq_{\mathbb{P}} \alpha$ then, by Proposition 13, $[\phi_{\alpha}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha) \neq \emptyset$ and, thus, $\phi_{\alpha}^{\div_{\kappa}} \not\leq_{\mathbb{P}} \alpha$.

$(\div 4)$. If $\alpha \equiv_{\mathbb{P}} \beta$, then $[\alpha]^+ = [\beta]^+$, i.e. $[\neg\alpha]^+ = [\neg\beta]^+$. Therefore, $\top \leq_{\mathbb{P}} \alpha$ iff $\top \leq_{\mathbb{P}} \beta$. So, if $\top \leq_{\mathbb{P}} \alpha$ then $\top \leq_{\mathbb{P}} \beta$ and, thus, by Proposition 13, $[\phi_{\alpha}^{\div_{\kappa}}]^+ = [\phi]^+ = [\phi_{\beta}^{\div_{\kappa}}]^+$. On the other hand, if $\top \not\leq_{\mathbb{P}} \alpha$ then $\top \not\leq_{\mathbb{P}} \beta$ and, thus, by Proposition 13, $[\phi_{\alpha}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha) = \sigma(\neg\beta) = [\phi_{\beta}^{\div_{\kappa}}]^+$. Therefore, in any case $[\phi_{\alpha}^{\div_{\kappa}}]^+ = [\phi_{\beta}^{\div_{\kappa}}]^+$ and, thus, $\phi_{\alpha}^{\div_{\kappa}} \equiv_{\mathbb{P}} \phi_{\beta}^{\div_{\kappa}}$.

$(\div 6a)$. Assume $\top \not\leq_{\mathbb{P}} \alpha$ and, thus, $\top \not\leq_{\mathbb{P}} \alpha \wedge \beta$. By Proposition 13 we have $[\phi_{\alpha}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha)$ and $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha \vee \neg\beta)$. If $\top \leq_{\mathbb{P}} \beta$, then $\sigma(\neg\alpha \vee \neg\beta) = \sigma(\neg\alpha)$ and, thus, $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = [\phi_{\alpha}^{\div_{\kappa}}]^+$. In particular, $\phi_{\alpha \wedge \beta}^{\div_{\kappa}} \leq_{\mathbb{P}} \phi_{\alpha}^{\div_{\kappa}}$ holds. Assume now $\top \not\leq_{\mathbb{P}} \beta$ instead. We have two cases: (i) $\sigma(\neg\alpha) \subseteq \sigma(\neg\beta)$; or (ii) $\sigma(\neg\beta) \subset \sigma(\neg\alpha)$. By Proposition 13, in case (i) we have $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha) = [\phi_{\alpha}^{\div_{\kappa}}]^+$. In case (ii) we have $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = \sigma(\neg\beta) \subset \sigma(\neg\alpha) = [\phi_{\alpha}^{\div_{\kappa}}]^+$. Therefore, in any case $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ \subseteq [\phi_{\alpha}^{\div_{\kappa}}]^+$ holds, i.e. $\phi_{\alpha \wedge \beta}^{\div_{\kappa}} \leq_{\mathbb{P}} \phi_{\alpha}^{\div_{\kappa}}$.

$(\div 7)$. Assume $\phi_{\alpha \wedge \beta}^{\div_{\kappa}} \not\leq_{\mathbb{P}} \alpha$. Therefore, $\top \not\leq_{\mathbb{P}} \alpha$ and, thus, $\top \not\leq_{\mathbb{P}} \alpha \wedge \beta$. By Proposition 13, $[\phi_{\alpha}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha)$ and $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha \vee \neg\beta)$. If $\top \leq_{\mathbb{P}} \beta$, then $\sigma(\neg\alpha \vee \neg\beta) = \sigma(\neg\alpha)$ and, thus, $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = [\phi_{\alpha}^{\div_{\kappa}}]^+$. So, assume now $\top \not\leq_{\mathbb{P}} \beta$ instead. From $\phi_{\alpha \wedge \beta}^{\div_{\kappa}} \not\leq_{\mathbb{P}} \alpha$ we also know that $\sigma(\neg\beta) \subset \sigma(\neg\alpha)$ cannot be the case as otherwise $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = \sigma(\neg\beta) \subseteq [\alpha]^+$, violating our assumption. Therefore, $\sigma(\neg\alpha) \subseteq \sigma(\neg\beta)$ has to hold and, thus, $[\phi_{\alpha \wedge \beta}^{\div_{\kappa}}]^+ = \sigma(\neg\alpha) = [\phi_{\alpha}^{\div_{\kappa}}]^+$. Consequently, in any case $\phi_{\alpha \wedge \beta}^{\div_{\kappa}} \leq_{\mathbb{P}} \phi_{\alpha}^{\div_{\kappa}}$ holds. \square

PROPOSITION 15.

Given any formula ϕ , for every formula α , ϕ_{α}^{\div} is a severe withdrawal operation iff it is a KM-severe withdrawal operation based on some probability distribution \mathbb{P} .

PROOF. Proposition 14 proves that every KM-severe withdrawal function \div_{κ} is a severe withdrawal operator, i.e. it satisfies the postulates $(\div 1)$ - $(\div 4)$, $(\div 6a)$ and $(\div 7)$.

We need to prove the opposite direction, and the proof is quite straightforward, given the results in [61], in particular Observation 15, which proves that for every severe withdrawal operator there is a correspondent withdrawal function defined on a system of spheres in the way shown in Figure 2. Given such a system of spheres, it is sufficient to define a probability distribution over the worlds in the same way we have done in the proof of Proposition 10, and it is easy to prove that such a probability distribution and the correspondent KM would generate the same withdrawal function. \square

PROPOSITION 18.

Consider any probability distribution \mathbb{P} over \mathcal{W} , and any KM κ and formulae ϕ and α . Then

$$\phi_{\alpha}^{\star} = \begin{cases} F(\min_{\kappa}([\alpha]^+)) & \text{if } \phi \leq_{\mathbb{P}} \neg\alpha \\ F([\phi]^+ \cap [\alpha]^+) & \text{otherwise.} \end{cases}$$

PROOF. We have to consider two cases:

- (1) if $\phi \leq_{\mathbb{P}} \neg\alpha$, then using Proposition 7

$$\begin{aligned}\phi_{\alpha}^{\star} &= \left(\bigvee_{\kappa} \min(\phi \perp \neg\alpha) \right) \wedge \alpha \\ &\equiv F([\phi]^+ \cup \min_{\kappa}[\alpha]^+) \cap [\alpha]^+ \\ &\equiv F(\min_{\kappa}([\alpha]^+)).\end{aligned}$$

- (2) If $\phi \not\leq_{\mathbb{P}} \neg\alpha$, then by Proposition 6, $\phi \perp \neg\alpha = \{F([\phi]^+)\}$ and, thus, $\phi_{\alpha}^{\star} = (\bigvee \min_{\kappa}(\phi \perp \neg\alpha)) \wedge \alpha \equiv F([\phi]^+ \cap [\alpha]^+)$. \square

COROLLARY 3.

Consider any probability distribution \mathbb{P} over \mathcal{W} , any KM κ and any formulae ϕ and α . Then

$$IC^+(\phi, \alpha) = -\log_2 \mathbb{P}(\alpha \mid \phi)$$

$$I^{\star}(\phi, \alpha) = \begin{cases} \log_2 \frac{\mathbb{P}(\phi)}{\mathbb{P}(F(\min_{\kappa}([\alpha]^+)))} & \text{if } \phi \leq_{\mathbb{P}} \neg\alpha \\ IC^+(\phi, \alpha) & \text{otherwise.} \end{cases}$$

PROOF. By definition of IC^+ we have

$$\begin{aligned}IC^+(\phi, \alpha) &= \kappa(\phi + \alpha) - \kappa(\phi) \\ &= -\log_2 \mathbb{P}(\phi \wedge \alpha) + \log_2 \mathbb{P}(\phi) \\ &= -\log_2(\mathbb{P}(\alpha \mid \phi) \cdot \mathbb{P}(\phi)) + \log_2 \mathbb{P}(\phi) \\ &= -\log_2 \mathbb{P}(\alpha \mid \phi) - \log_2 \mathbb{P}(\phi) + \log_2 \mathbb{P}(\phi) \\ &= -\log_2 \mathbb{P}(\alpha \mid \phi).\end{aligned}$$

By definition of IC^{\star} we have the following two cases. If $\phi \not\leq_{\mathbb{P}} \neg\alpha$ then, by Proposition 18, ϕ_{α}^{\star} is equivalent to $\phi \wedge \alpha$ and, thus, is equivalent to $\phi + \alpha$. Therefore, $IC^{\star}(\phi, \alpha) = IC^+(\phi, \alpha)$. On the other hand, assume $\phi \leq_{\mathbb{P}} \neg\alpha$ instead. Then, by Proposition 18, $\phi_{\alpha}^{\star} = F(\min_{\kappa}([\alpha]^+))$ and, thus,

$$\begin{aligned}IC^{\star}(\phi, \alpha) &= -\log_2 \mathbb{P}(F(\min_{\kappa}([\alpha]^+))) + \log_2 \mathbb{P}(\phi) \\ &= \log_2 \frac{\mathbb{P}(\phi)}{\mathbb{P}(F(\min_{\kappa}([\alpha]^+))}).\end{aligned}$$

\square

PROPOSITION 19.

Let \star be a KM-based revision operator. Then \star satisfies (C1), (C3) and (C4).

PROOF. We prove that each postulate holds.

(C1). If $\psi \leq_{\mathbb{P}} \alpha$, we have the following possible cases:

Case $\phi \not\leq_{\mathbb{P}} \neg\psi$. Therefore, $\phi \not\leq_{\mathbb{P}} \neg\alpha$ holds. Then, by Proposition 18, $[\phi_{\alpha}^*] = [\phi]^+ \cap [\alpha]^+$ and $[(\phi_{\alpha}^*)_{\psi}^*] = (([\phi]^+ \cap [\alpha]^+) \cap [\psi]^+)$. Now, $\phi \not\leq_{\mathbb{P}} \neg\psi$, by Proposition 18, $[\phi_{\psi}^*] = [\phi]^+ \cap [\psi]^+$. Eventually, as by hypothesis $[\psi]^+ \subseteq [\alpha]^+$, $(([\phi]^+ \cap [\alpha]^+) \cap [\psi]^+) = [\phi]^+ \cap [\psi]^+$ follows, i.e. $(\phi_{\alpha}^*)_{\psi}^* \equiv_{\mathbb{P}} \phi_{\psi}^*$.

Case $\phi \leq_{\mathbb{P}} \neg\alpha$. Therefore, $\phi \leq_{\mathbb{P}} \neg\psi$ holds. Then, by Proposition 18, $[\phi_{\alpha}^*] = \min_{\kappa}([\alpha]^+)$. Since by hypothesis $[\psi]^+ \subseteq [\alpha]^+$, we have two possibilities: (i) $\min_{\kappa}([\alpha]^+ \cap [\psi]^+) \neq \emptyset$, and in such a case $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\alpha]^+) \cap [\psi]^+$ and, since $[\psi]^+ \subseteq [\alpha]^+$, we have also $\min_{\kappa}([\alpha]^+) \cap [\psi]^+ = \min_{\kappa}([\psi]^+)$; (ii) $\min_{\kappa}([\alpha]^+) \cap [\psi]^+ = \emptyset$, and also in such a case $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$. On the other hand, given $\phi \leq_{\mathbb{P}} \neg\psi$ we have $[\phi_{\psi}^*] = \min_{\kappa}([\psi]^+)$. That is, $(\phi_{\alpha}^*)_{\psi}^* \equiv_{\mathbb{P}} \phi_{\psi}^*$.

Case $\phi \leq_{\mathbb{P}} \neg\psi$ and $\phi \not\leq_{\mathbb{P}} \neg\alpha$. Then, by Proposition 18, $[\phi_{\alpha}^*] = [\phi]^+ \cap [\alpha]^+$ and, since $[\phi]^+ \subseteq [\neg\psi]^+$ by hypothesis, $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$. On the other hand, again since $[\phi]^+ \subseteq [\neg\psi]^+$, $[\phi_{\psi}^*] = \min_{\kappa}([\psi]^+)$. That is, $(\phi_{\alpha}^*)_{\psi}^* \equiv_{\mathbb{P}} \phi_{\psi}^*$.

We can conclude that in any case $(\phi_{\alpha}^*)_{\psi}^* \equiv_{\mathbb{P}} \phi_{\psi}^*$ holds, as desired.

(C3). If $\phi_{\psi}^* \leq_{\mathbb{P}} \alpha$, we have the following possible cases:

Case $\phi \leq_{\mathbb{P}} \neg\alpha$. Therefore, $\phi \leq_{\mathbb{P}} \neg\psi$ holds. Then, by Proposition 18, $[\phi_{\alpha}^*] = \min_{\kappa}([\alpha]^+)$. Since by hypothesis $[\psi]^+ \subseteq [\alpha]^+$, we have two possibilities: (i) $\min_{\kappa}([\alpha]^+) \cap [\psi]^+ \neq \emptyset$, and in such a case $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\alpha]^+) \cap [\psi]^+$ and, since $[\psi]^+ \subseteq [\alpha]^+$, we have also $\min_{\kappa}([\alpha]^+) \cap [\psi]^+ = \min_{\kappa}([\psi]^+)$; (ii) $\min_{\kappa}([\alpha]^+) \cap [\psi]^+ = \emptyset$, and also in such a case $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$. On the other hand, given $\phi \leq_{\mathbb{P}} \neg\psi$ we have $[\phi_{\psi}^*] = \min_{\kappa}([\psi]^+)$. That is, $(\phi_{\alpha}^*)_{\psi}^* \equiv_{\mathbb{P}} \phi_{\psi}^*$.

Case $\phi \leq_{\mathbb{P}} \neg\alpha$. Since $\phi_{\psi}^* \leq_{\mathbb{P}} \alpha$, by Proposition 18, it must be $[\phi_{\psi}^*] = \min_{\kappa}([\psi]^+)$. Also, by Proposition 18 we have $[\phi_{\alpha}^*] = \min_{\kappa}([\alpha]^+)$. Now we have two possibilities: (i) $\min_{\kappa}([\alpha]^+) \subseteq [\neg\psi]^+$ and in such a case, by Proposition 18, $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$, i.e. $[(\phi_{\alpha}^*)_{\psi}^*] = [\phi_{\psi}^*]$, and consequently by hypothesis $(\phi_{\alpha}^*)_{\psi}^* \leq_{\mathbb{P}} \alpha$; (ii) $\min_{\kappa}([\alpha]^+) \not\subseteq [\neg\psi]^+$ and in such a case $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\alpha]^+) \cap [\psi]^+$, and consequently by hypothesis $(\phi_{\alpha}^*)_{\psi}^* \leq_{\mathbb{P}} \alpha$.

Case $\phi \not\leq_{\mathbb{P}} \neg\alpha$. Then, by Proposition 18 $[\phi_{\alpha}^*] = [\phi]^+ \cap [\alpha]^+$. Now we have two possibilities: (i) $[\phi]^+ \cap [\alpha]^+ \subseteq [\neg\psi]^+$ and in such a case we have, by Proposition 18, $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$, i.e. $[(\phi_{\alpha}^*)_{\psi}^*] = [\phi_{\psi}^*]$, and consequently by hypothesis $(\phi_{\alpha}^*)_{\psi}^* \leq_{\mathbb{P}} \alpha$. (ii) $[\phi]^+ \cap [\alpha]^+ \not\subseteq [\neg\psi]^+$ and in such a case, by Proposition 18, $[(\phi_{\alpha}^*)_{\psi}^*] = [\phi]^+ \cap [\alpha]^+ \cap [\psi]^+$, and consequently $(\phi_{\alpha}^*)_{\psi}^* \leq_{\mathbb{P}} \alpha$.

We can conclude that in any case $(\phi_{\alpha}^*)_{\psi}^* \leq_{\mathbb{P}} \alpha$, as desired.

(C4). If $\phi_{\psi}^* \not\leq_{\mathbb{P}} \neg\alpha$, we have the following possible cases:

Case $\phi \leq_{\mathbb{P}} \neg\alpha$. In such a case, by Proposition 18, $[\phi_{\alpha}^*] = \min_{\kappa}([\alpha]^+)$. Now we have two possibilities: (i) $\min_{\kappa}([\alpha]^+) \subseteq [\neg\psi]^+$. In such a case, by Proposition 18, $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$, i.e. $[(\phi_{\alpha}^*)_{\psi}^*] = [\phi_{\psi}^*]$, and consequently by hypothesis $(\phi_{\alpha}^*)_{\psi}^* \not\leq_{\mathbb{P}} \neg\alpha$. (ii) $\min_{\kappa}([\alpha]^+) \not\subseteq [\neg\psi]^+$. In such a case, by Proposition 18, $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\alpha]^+) \cap [\psi]^+$, and consequently $(\phi_{\alpha}^*)_{\psi}^* \not\leq_{\mathbb{P}} \neg\alpha$.

Case $\phi \not\leq_{\mathbb{P}} \neg\alpha$. Then, by Proposition 18, $[\phi_{\alpha}^*] = [\phi]^+ \cap [\alpha]^+$. Now we have two possibilities: (i) $[\phi]^+ \cap [\alpha]^+ \subseteq [\neg\psi]^+$. In such a case we have, by Proposition 18, $[(\phi_{\alpha}^*)_{\psi}^*] = \min_{\kappa}([\psi]^+)$, i.e. $[(\phi_{\alpha}^*)_{\psi}^*] = [\phi_{\psi}^*]$, and consequently by hypothesis $(\phi_{\alpha}^*)_{\psi}^* \not\leq_{\mathbb{P}} \neg\alpha$. (ii)

$[\phi]^+ \cap [\alpha]^+ \not\subseteq [\neg\psi]^+$. In such a case $[(\phi_\alpha^*)_\psi^*] = [\phi]^+ \cap [\alpha]^+ \cap [\psi]^+$, and consequently $(\phi_\alpha^*)_\psi^* \not\subseteq_{\mathbb{P}} \neg\alpha$.

We can conclude that in any case $(\phi_\alpha^*)_\psi^* \not\subseteq_{\mathbb{P}} \neg\alpha$, as desired. \square

Appendix B. On Dalal's and Satoh's BC operators

We consider here more in detail two well-known BC procedures mentioned in Section 6.

Dalal [15]. Semantically, Dalal defines BC in terms of a *Dalal distance* d_D , i.e. the Hamming distance between worlds, considering each world as a set of literals (see Section 2). Specifically, given two worlds w, w' w.r.t. Σ ,

$$d_D(w, w') = |\{L \mid L \in w \text{ and } L \notin w'\}|.$$

That is, the distance between two worlds corresponds to the number of literals the two worlds disagree on. The function d_D can now be extended to define the distance between a formula and a possible world in the following way:

$$d_D(\phi, w') = \min_{w \in [\phi]} (d_D(w, w')).$$

Now, d_D determines a total preference pre-order \leq_ϕ between worlds w.r.t. a formula ϕ : namely

$$w \leq_\phi w' \text{ iff } d_D(\phi, w) \leq d_D(\phi, w').$$

Obviously, since it is based on a comparison between natural numbers, \leq_ϕ is a total pre-order. Eventually, *Dalal contraction* \div_D and *revision* $*_D$ are defined as follows (see also [7]):

$$\phi_\alpha^{\div D} = [\phi] \cup \min_{\leq_\phi}([\neg\alpha]) \quad (\text{B1})$$

$$\phi_\alpha^{*D} = \min_{\leq_\phi}([\alpha]). \quad (\text{B2})$$

We recall an example from [7].

EXAMPLE B1.

[7, Example 1] Let $\Sigma = \{b, f, h\}$, interpreted as *bird*, *flies* and *has a beak*, respectively. We are informed that Tweety is a bird, and we assume that it can fly ($\phi = b \wedge f$); but later, getting closer, we realize that it is a penguin and it does not fly ($\alpha = \neg f$). Hence we need to revise our KB into ϕ_α^{*D} . We have that $[\phi] = \{bf\bar{h}, bfh\}$, while $[\alpha] = \{\bar{b}\bar{f}\bar{h}, \bar{b}\bar{f}h, b\bar{f}\bar{h}, b\bar{f}h\}$. Comparing these two sets of worlds, the function d_D has the following values:

	$bf\bar{h}$	bfh	$d_D(\phi, w)$
$\bar{b}\bar{f}\bar{h}$	2	3	2
$\bar{b}\bar{f}h$	3	2	2
$b\bar{f}\bar{h}$	1	2	1
$b\bar{f}h$	2	1	1

We obtain $\min_{\leq_\phi}([\alpha]) = \{\bar{b}\bar{f}\bar{h}, b\bar{f}h\}$, which is the set of worlds defining ϕ_α^{*D} , i.e. we end up believing that Tweety is a bird that does not fly.

Please note that \div_D and $*_D$ satisfy indeed all the classical AGM postulates. Consequently, by Proposition 9, we know that, given a formula ϕ , we can define a ϕ -faithful ranking based on d_D that defines \div_D and $*_D$. e.g. in this particular case, \div_D and $*_D$ would be defined using the following ranking r .

r	<i>worlds</i>
2	$\overline{bfh} \ \overline{bfh}$
1	$\overline{bfh} \ \overline{bfh} \ \overline{bfh} \ \overline{bfh}$
0	$bfh \ bfh$

Hence, by Proposition 11 we can easily conclude the following.

PROPOSITION 20.

Let ϕ be any formula and let \div_D be the Dalal contraction operator. Then there is a KM contraction \div_κ s.t., for all α , $\phi_\alpha^{\div_D} \equiv \phi_\alpha^{\div_\kappa}$ holds.

Analogously, let ϕ be any formula and let $*_D$ be the Dalal revision operator. Then there is a KM revision \star_κ s.t., for all α , $\phi_\alpha^{*_D} \equiv \phi_\alpha^{\star_\kappa}$ holds.

Proposition A12 does not hold in the opposite direction: given a formula ϕ , Dalal's approach can define only a single contraction and a single revision, both based on the Dalal distance. However, we can easily define a probability distribution that do not respect the Dalal distance and, thus, we may yield different results that may not be emulated by Dalal's contraction operator. For example, consider the following probability distribution:

<i>worlds</i>	\mathbb{P}	κ_S
$\overline{bfh} \ \overline{bfh} \ \overline{bfh}$	1/16	4
$\overline{bfh} \ \overline{bfh} \ \overline{bfh}$	2/16	3
$bfh \ bfh$	3.5/16	2.193

This ranking imposes e.g. that $\phi_\alpha^{\star_{\kappa_S}} = F(\overline{bfh})$, which is not what we obtain with the Dalal approach.

Anyway, despite Proposition A12 indicates a formal correspondence, the rationale motivating the Dalal operations and ours is quite different, as explained in Section 6.

Satoh [62].

Satoh proposes a form on non-monotonic reasoning that is defined on top of a notion of *minimal revision*. Its proposal is close to Dalal's one, since a pre-order between worlds is determined by referring to the differences in the literals satisfied by pairs of worlds. However, such an order between worlds is based on set-inclusion instead of cardinality, i.e.

$$\Delta_{w,w'} = \{L \mid L \in w \text{ and } L \notin w'\}.$$

Given a world v , we can define a pre-order \leq_v over the worlds as follows:

$$w \leq_v w' \text{ iff } \Delta_{v,w} \subseteq \Delta_{v,w'}.$$

Now, given two formulae ϕ and α , we define the set $\min_\phi([\alpha])$ as follows:

$$\min_\phi([\alpha]) = \{w \mid w \in \min_{\leq_v}([\alpha]) \text{ for some } v \in [\phi]\}.$$

That is, $\min_\phi([\alpha])$ returns the worlds in $[\alpha]$ that differ less w.r.t. some ϕ model. The correspondent contraction and revision operations, \div_S and \star_S , are defined according to Satoh as follows:

$$\phi_\alpha^{\div_S} = [\phi] \cup \min_\phi([\neg\alpha]) \quad (\text{B3})$$

$$\phi_\alpha^{\star_S} = \min_\phi([\alpha]). \quad (\text{B4})$$

Satoh's approach is not based on a total order of the worlds, and in fact his operations do not satisfy all the AGM postulates [13,section 2]. Consequently, as our operators satisfy the AGM postulates, in general Satoh's operators may not be modelled in our framework and, thus, we do not investigate it further here.

Appendix C. A short recap on knowledge measures as from [67]

We recap here succinctly the notion of KMs as introduced in [67].

Given an alphabet Σ , a *substitution* θ is a set $\theta = \{p_1/L_1, \dots, p_{|\Sigma|}/L_{|\Sigma|}\}$ such that each propositional letter in Σ occurs exactly once in $\{p_1, \dots, p_{|\Sigma|}\}$ as well as in $\{L_1, \dots, L_{|\Sigma|}\}$. The intuition is that propositional letters may be renamed by literals. With (i) $w\theta$ we indicate the interpretation obtained from w by replacing every occurrence of p_i in w with L_i (with the convention that double negations are normalized²⁷); and (ii) for a set of interpretations \mathcal{W} , with $\mathcal{W}\theta$ we denote the set $\mathcal{W}\theta = \{w\theta \mid w \in \mathcal{W}\}$. If \mathcal{W}_1 and \mathcal{W}_2 are two sets of interpretations, we write $\mathcal{W}_1 \leq_s \mathcal{W}_2$ if there is a substitution θ such that $\mathcal{W}_1\theta \subseteq \mathcal{W}_2$. If the subset relation is strict, we write $\mathcal{W}_1 <_s \mathcal{W}_2$. With $[\phi]_\Sigma$ we denote the set of models of a formula ϕ w.r.t. the alphabet Σ . Then, we say that a formula ϕ *s-entails* a formula ψ , denoted $\phi \leq_s \psi$, if $[\phi]_\Sigma \leq_s [\psi]_\Sigma$ holds. We write $\phi <_s \psi$ if $\phi \leq_s \psi$ and $\psi \not\leq_s \phi$.²⁸ Note that if $\phi \models \psi$ then $\phi \leq_s \psi$, but not vice versa [67]. We also write $\phi \equiv_s \psi$ if $\phi \leq_s \psi$ and $\psi \leq_s \phi$ hold.

EXAMPLE C1. ([67]).

It is easily verified that using substitution $\theta = \{p/\bar{q}, q/p\}$, $p \leq_s \bar{q}$, but $p \not\leq \bar{q}$, i.e. $p \not\equiv \bar{q}$. Similarly, it is easily verified that $\bar{q} \leq_s p$ and, thus, $p \equiv_s \bar{q}$.

We anticipate that the intuition behind s-entailment is to account for the idea that all propositional literals carry the same 'amount of information'. For instance, in Example C1 above, p and \bar{q} will carry the same amount of information as $p \equiv_s \bar{q}$ [67].

Now, according to [67], a KM κ is a function mapping formulae into $\mathbb{R}^+ \cup \{0, \infty\}$ satisfying the following:

- (T): $\kappa(\top) = 0$ and $\kappa(\perp) = \infty$;
- (E): if $\phi <_s \psi$ then $\kappa(\psi) < \kappa(\phi)$, for $< \in \{\leq, <\}$;
- (L): $\kappa(\phi)$ does depend on Σ_ϕ only, i.e. $\kappa(\phi) = \kappa(F[\phi]_{\Sigma_\phi}) = \kappa(F[\phi]_\Sigma)$;²⁹
- (M): if ϕ is satisfiable then
 - a. $0 \leq \kappa(\phi) \leq |\Sigma_\phi|$; and
 - b. if $|\phi]_{\Sigma_\phi}| = 1$ then $\kappa(\phi) = |\Sigma_\phi|$.

²⁷ $\neg\neg p \mapsto p$.

²⁸ We use the same relation symbol whether ranging over sets of interpretations or formulae. We will disambiguate it whenever required.

²⁹ Recall that $\phi \equiv F[\phi]_{\Sigma_\phi} \equiv \phi \equiv F[\phi]_\Sigma$ (see also Equation 3).

Let us briefly explain the above (TELM) axioms [67]. Concerning axiom (T), it is the same as (KM1), so see (KM1)'s explanation.

Concerning axiom (E), it is a stronger version of (KM2): if $\phi \models \psi$ then ϕ has fewer models than ψ , which means that there is less uncertainty about which is the actual model for ϕ compared with ψ . Moreover, whatever is entailed by ψ is also entailed by ϕ . Combining the two facts means that ϕ represents more information about what is the actual world than ψ . Concerning s-entailment, e.g. for $\phi := p$ and $\psi := q$ neither $\phi \models \psi$ nor $\psi \models \phi$ hold. But, according to axiom (E), we have that $\kappa(\phi) = \kappa(\psi)$ instead, i.e. p and q carry the same amount of knowledge, i.e. according to [67], a KM is insensitive to symbol names, i.e. a symbol p carries as much knowledge as another symbol q .

Concerning axiom (L), this axiom says that, to what concerns KMs, we may restrict our attention to Σ_ϕ , i.e. the set of all propositional letters occurring in a formula ϕ and, thus, symbols not occurring in ϕ do not contribute to represent additional information. Note, however, that such an assumption may not be considered valid e.g. under the *Closed World Assumption* [54], where the truth of a letter not occurring in a formula is always false.

Concerning axiom (M), this axiom tells us that KMs are bounded in the sense that a satisfiable formula may not represent more information than the number of symbols it relies on and the bound is reached only if the formula exactly describes the actual world.

It has been shown in [67] that a KM satisfying the axioms above is

$$\begin{aligned} \kappa_h(\phi) &= |\Sigma| - \log_2 |[\phi]_\Sigma| \\ &= |\Sigma_\phi| - \log_2 |[\phi]_{\Sigma_\phi}| \end{aligned} \quad (C1)$$

where $\kappa_h(\phi) = \infty$, if ϕ is not satisfiable. That is, we count the models of ϕ w.r.t. Σ_ϕ , take the logarithm and subtract it from the maximal amount of knowledge a formula may carry according to axiom (M). So, the more models ϕ has, the less we know about which is the actual world, i.e. the less information ϕ represents.

EXAMPLE C2. (Running example cont).

Consider the KB \mathcal{K} in Example 1. Then we have $|\Sigma_{\mathcal{K}}| = 5$, and $|[\mathcal{K}]_{\Sigma_{\mathcal{K}}}| = 6$ and, thus,

$$\kappa_h(\mathcal{K}) = 5 - \log_2 6 \approx 2.415.$$

REMARK C1.

In [67] some other related measures have been introduced: namely,

$$\begin{aligned} \bar{\alpha}(\phi) &= \frac{\kappa(\phi)}{|\Sigma_\phi|} && \text{accuracy} \\ \bar{\gamma}(\phi) &= \frac{\kappa(\phi)}{|\phi|} && \text{conciseness} \\ \pi(\phi) &= \arg \max_{\{\psi \mid \psi \equiv \phi\}} \bar{\alpha}(\psi) \cdot \bar{\gamma}(\psi) && \text{Pareto optimality} \end{aligned}$$

where $|\phi|$ is the ‘length’ of ϕ [67]. The first one defines how precise a KB is in describing the actual world, the second one defines how succinct a KB is w.r.t. the knowledge it represents and the last one establishes when we may not increase accuracy without decreasing conciseness (or vice versa). Note that, according to [67], for a satisfiable formula ϕ , ϕ is Pareto optimal iff for all ψ such that $\psi \equiv \phi$ we have that $|\Sigma_\phi| \cdot |\phi| \leq |\Sigma_\psi| \cdot |\psi|$.

These measures can trivially be extended to information-theoretic KMs, so we will not address them further here.

Received 21 May 2025