

Robustness of Evidential Probability

Choh Man Teng and Gregory Wheeler

Artificial Intelligence Center - CENTRIA

Department of Computer Science, Universidade Nova de Lisboa

2829-516 Caparica, Portugal

cm.teng,grw@fct.unl.pt

Abstract

In the Evidential Probability (EP) (Kyburg and Teng 2001) framework, interval-valued probabilities are assigned to events given a knowledge base containing statements about logical relationships of classes of objects as well as statements about relative frequencies pertaining to some of those classes. We are interested in the robustness of this inference with respect to perturbations in the given knowledge base. First we define how changes to a (non-deductively-closed) knowledge base are to be carried out by constructing revision and contraction operators that satisfy the AGM postulates, and then we discuss some conditions and measures of robustness. We appeal to a result in (Kyburg and Teng 2002) that allows a logic of risky knowledge based on EP to be characterized by the classical modal system EMN, and define AGM operators for EP through a mapping from monotone neighborhood modal structures to polymodal Kripke structures that admit a first-order correspondent.

1 Introduction

The logic of risky knowledge (Kyburg and Teng 2002) interprets the necessity operator ∇_ε in the classical modal system EMN (Chellas 1980, Hansen 2003) as an ε -accepted knowledge operator, where a sentence of the language is ε -accepted if the probability of its denial is at most ε for some small, fixed value ε . System EMN thus reveals the logical structure of a set of ε -accepted sentences.

But risky knowledge is based upon a body of evidence, itself uncertain, and while each item of evidence supporting an ε -accepted sentence faces a similar if smaller risk of error, the impact of spurious evidence may vary depending upon how that sentence is connected to the evidence. To illustrate, imagine that the ε -acceptability of a sentence ϕ rests upon three pieces of evidence, one of which is false. If there is no significant difference between the probability that ϕ given the full evidence and the probability that ϕ given any pair of the three evidence statements, then we might view the full evidence to provide robust grounds for the probability assigned to ϕ . If instead ϕ depends more on one piece of evidence than it does on the other two, then assessing what effect a false piece of evidence would have on the probability that ϕ will depend upon which piece is found to be in error. And if the probability that ϕ varies with every pair of evidence statements, then discovering there is an error would be sufficient to know there is an effect on ϕ even before learning its type or size.

So even if the chance of error for every item of evidence is uniformly small, the effect of false evidence on risky knowledge might vary significantly. But to investigate variability in evidence we must first be in a position to evaluate the evidential probability (EP) (Kyburg 1961, Kyburg and Teng 2001) of ϕ relative to various counter-to-accepted bodies of evidence.

An approach to ‘counterfactual’ evidential probability explored in (Haenni et al. 2010, Wheeler and Williamson 2009) treats individual evidence statements as random variables and constructs counter-to-accepted evidence sets by directly manipulating individual variables—much like Pearl’s treatment of counterfactual conditionals (Pearl 2000) and Spirtes *et. al.*’s (Spirtes et al. 2000) interventions within probabilistic causal models—except that rather than activate the theory of causal Bayes nets for a single positive joint distribution over the variables, the PROGICNET framework turns to credal networks (Levi 1980, Cozman 2000, Haenni et al. 2010) to work with sets of distributions associated with those variables.

A counterfactual evidence set on the PROGICNET approach is achieved by flipping the value of an individual random variable, and the question of variability in the evidence for a sentence ϕ is explored by wiggling every variable individually to see the effect on the evidential probability of ϕ . But EP evidence sets contain logical and probabilistic relationships, so the genuine counterfactual evidence sets for a sentence ϕ will typically be a proper subset of the *possibly relevant* evidence sets discussed in (Haenni et al. 2010, §4.3.2). Because individually flipping variables ignores logical relationships between sentences (variables) that occur in EP evidence sets, the practice of manipulating variables one at a time may create an unsatisfiable set of constraints that shows up as spurious variability in the evidence. To identify genuine counterfactual evidence sets we must account for the ramifications from altering an item of evidence instead of simply manipulating isolated sentences, and this demands that we view a proposed counterfactual evidence set centered around the alteration of some sentence as a contraction of the body of evidence by that sentence (Gärdenfors 1988).

In this essay we define an AGM contraction operator for the classical system EMN by adapting a result for constructing AGM revision operators for different classes of classical modal logics (Wheeler 2010). The contracted evidence set can then be used to evaluate the robustness of an evidential probability assignment.

The structure of the essay is as follows. In Section 2 we review evidential probability and briefly address resistance in some quarters to threshold accounts of rational belief by focusing on recent impossibility results that threaten to scupper the plan to threshold anything. In Section 3 we define AGM revision and contraction operators for the modal logic EMN, which characterizes the structure of risky knowledge based on high threshold evidential probability. In Section 4 we discuss a robustness measure of evidential probability and some conditions relating robustness and changes in the knowledge base. Proofs are set in an appendix.

2 Evidential Probability

Everyday reasoning is nonmonotonic. We draw conclusions that typically, usually, probably, customarily follow from a set of premises, even though the truth of these conclusions are not guaranteed given the truth of the premises. We accept the conclusions according to certain criteria and act as if these conclusions were true until further evidence casts doubt on their status. One criterion for acceptance that bridges probabilistic and logical reasoning is to accept as true statements whose probabilities are above a given threshold—statements whose probabilities are so high that they are practically certain.

In (Kyburg and Teng 2002) such a logic of risky knowledge was laid out, where acceptance is grounded in evidential probability (EP) (Kyburg 1961, Kyburg and Teng 2001). Let us first discuss a number of properties of acceptance, with respect to high probability thresholding in general, and with respect to risky knowledge in the evidential probability framework in particular.

Let \mathcal{L} be a guarded fragment of first-order logic. The domain of the probability function $\text{Prob}(\cdot, \cdot)$ is $\mathcal{L} \times \wp(\mathcal{L})$, and its range is intervals $[l, u]$. Evidential probability is based on observed relative frequencies, which are objective but approximate, and thus the probabilities are given by intervals. The probability of a statement χ is relativized to a given background knowledge Γ_δ , denoted by $\text{Prob}(\chi, \Gamma_\delta) = [l, u]$. The statements in Γ_δ represent a knowledge base, which includes categorical statements as well as statistical generalities. A statement with probability $[l, u]$ is accepted into Γ_δ if $l \geq 1 - \delta$.

Categorical statements are logical formulas specifying what is accepted as true in the background knowledge. *Statistical statements* are of the form

$$\% \vec{x}(\tau(\vec{x}), \rho(\vec{x}), [l, u]), \tag{1}$$

where \vec{x} is a sequence of logical variables, and the statement as a whole says that among the models

satisfying ρ , between a fraction l and a fraction u also satisfy τ . For example $\%x(W(x), U_1(x), [.43, .51])$ might express that the proportion of balls drawn from urn U_1 that are white is between .43 and .51. We call $\tau(\vec{x})$ the *target formula* and $\rho(\vec{x})$ the *reference formula* of the above statistical statement.

A *candidate statistical statement* concerning χ , given background knowledge Γ_δ , is a statistical statement $\%x(\tau(\vec{x}), \rho(\vec{x}), [l, u])$ in Γ_δ where $\chi \leftrightarrow \tau(a)$ and $\rho(a)$ are both also in Γ_δ for some term a in \mathcal{L} . There may be many such statistical statements for a given χ , each being a candidate for determining the evidential probability of χ , but each referring to an interval that does not necessarily agree with the others. This is the classic problem of the reference class (Reichenbach 1949).

Three principles are used to vet and combine such candidate statistical statements to obtain a single probability interval for the target statement χ . These principles make use of the following definitions.

[Conflict] Two intervals $[l, u]$ and $[l', u']$ *conflict* iff neither $[l, u] \subseteq [l', u']$ nor $[l, u] \supseteq [l', u']$.

[Cover] The *cover* of a set of intervals is the shortest interval including all the intervals.

[Closure under Difference] Given a set of intervals I , a set of intervals I' is *closed under difference* with respect to I iff $I' \subseteq I$ and I' contains every interval in I that conflicts with an interval in I' .

These definitions are extended to statistical statements by applying the definitions to their associated intervals.

The three principles for resolving conflict between candidate statistical statements pertaining to a target statement χ may be stated as follows.

1. **[Richness]** If two statistical statements conflict and the first is based on a marginal distribution of the joint distribution on which the second statistical statement is based, disregard the first.
2. **[Specificity]** If two conflicting statistical statements both survive the principle of richness, and the second employs a reference class that is known to be included in the first, disregard the first.
3. **[Strength]** Those statistical statements not disregarded by the principles of richness and specificity are called *relevant*. The probability of χ is the cover that is the shortest among all covers of non-empty sets of statistical statements closed under difference with respect to the set of relevant statistical statements.

When two statements conflict, the principle of richness specifies that more informative statistical statements should be favored over less informative statistical statements. The principle of specificity says that statistical statements associated with more specific reference classes should be preferred to those associated with less specific reference classes. The principle of strength takes the shortest interval that covers a set of statements closed under difference with respect to the surviving statements.

The probability of a statement χ with respect to Γ_δ , derived according to the above three principles applied exhaustively and sequentially to the set of candidate statistical statements in Γ_δ , is unique and consistent. For further discussion of these principles, see (Kyburg and Teng 2001, Teng 2007).

3 AGM Operators in Classical Modal Systems

Standard AGM constructions are not directly applicable in the evidential probability framework, nor are they applicable in the logic of risky knowledge. However, there is a method for constructing AGM revision and contraction operators for the classical modal system EMN, which we now describe.

We proceed in three steps. First, we review classical modal logic; second, we define a correspondence language suitable for translating sentences of EMN into a first-order correspondent which admits representation within a propositional language. Finally, we show that a revision operator defined on this first-order EMN correspondence language satisfies the AGM postulates. Contraction follows via the Harper Identity.

3.1 Classical Modal Logic

To begin, we highlight the difference between neighborhood structures and standard Kripke structures. Whereas Kripke frames are characterized by a binary accessibility relation defined over a set of worlds, a **neighborhood frame** for the propositional modal language $\mathcal{L}_\nabla(\Phi)$ is a pair $\mathbb{F} = (W, \mathcal{N})$ where

- a) W is a non-empty set of worlds,
- b) $\mathcal{N} : W \mapsto \wp(\wp(W))$ is a neighborhood function, i.e. $\mathcal{N}(w) \subseteq \wp(W)$, for each $w \in W$.

If $\mathbb{F} = (W, \mathcal{N})$ is a neighborhood frame, Φ a countable set of propositional variables, and $V : \Phi \mapsto \wp(W)$ is a valuation on \mathbb{F} , then $\mathbb{M} = (W, \mathcal{N}, V)$ is a **neighborhood model** based on \mathbb{F} .

The satisfiability conditions for non-modal propositional formulas on neighborhood models are analogous to Kripke models, except for modal necessity ($\nabla\varphi$) and possibility ($\Delta\varphi$) statements:

- $\Vdash_w^{\mathbb{M}} \nabla\varphi$ iff $(\exists X \in \mathcal{N}(w), \forall w^* \in X) : \Vdash_{w^*}^{\mathbb{M}} \varphi$,
- $\Vdash_w^{\mathbb{M}} \Delta\varphi$ iff $(\forall X \in \mathcal{N}(w), \exists w^* \in X) : \Vdash_{w^*}^{\mathbb{M}} \varphi$,

where $\mathbb{M} = (W, \mathcal{N}, V)$ is a neighborhood model, w a world in W , and X a set of worlds.

Nonstandard modal logics have been proposed as qualitative representations of rational acceptance (Kyburg and Teng 2002, Arló-Costa 2002). Whereas standard (i.e., Kripke) models enjoy strong distribution properties due to the validity of the schema (K) $\nabla(\varphi \rightarrow \psi) \rightarrow (\nabla\varphi \rightarrow \nabla\psi)$. Specifically, whereas (C) $\nabla\varphi \wedge \nabla\psi \rightarrow \nabla(\varphi \wedge \psi)$, (M) $\nabla(\varphi \wedge \psi) \rightarrow \nabla\varphi \wedge \nabla\psi$, and (N) $\nabla\top$ are all valid in standard Kripke models, none are generally valid on neighborhood models

The problematic axiom for logics of ε -acceptance is the schema (C), which is valid on the class of minimal models closed under intersections. But closure under conjunction is precisely the behavior to prohibit if we read $\nabla\varphi$ as ‘ φ has probability greater than $1 - \delta$ ’. So we shall be concerned with classes of models in which (C) is invalid and, correspondingly, systems of modal logic in which instances of (C) are not theorems. The logic of risky knowledge (Kyburg and Teng 2002) is identified with classical systems in which instances of (M) and (N) are theorems (i.e., system EMN).

3.2 Correspondence Languages

This section addresses the translation step by appealing to results from modal simulation theory, which identifies a class of neighborhood frames with some multi-modal Kripke frame, and correspondence theory, which here will characterize a bi-modal Kripke frame by sentences of first-order logic. This requires specifying three languages: \mathcal{L}_∇ , a classical propositional monomodal language; \mathcal{L}_\diamond , a standard propositional polymodal language; and \mathcal{L}_∇^1 , the final first-order translation language corresponding to \mathcal{L}_∇ . This technique does not cover all classical modal systems, but it does cover many of them, including EMN. Frame validity expresses a second-order property, because it quantifies over subsets of worlds, and this does not always admit expression by a first-order formula. In (Kracht 1993, Kracht and Wolter 1999) it was observed that a particular class of classical modal formulas in language \mathcal{L}_∇ , interpreted over bi-modal Kripke structures, correspond to Sahlqvist formulas, for which Sahlqvist correspondence holds via the Sahlqvist-van Benthem algorithm. This technique was extended to monotonic modal logic by Marc Pauly in an unpublished manuscript, which is described in (Hansen 2003).

Let $p \in \Phi$ and pt be a unary modal operator. A **classical monomodal** grammar and a **standard polymodal** grammar are generated by the following, respectively:

- $\mathcal{L}_\nabla(\Phi) : p \mid \neg\varphi \mid \varphi \vee \psi \mid \nabla\varphi$
- $\mathcal{L}_\diamond(\Phi) : p \mid \neg\varphi \mid \varphi \vee \psi \mid \diamond_1\varphi \mid \diamond_2\varphi \mid \text{pt}$

For the standard polymodal language $\mathcal{L}_\diamond(\Phi)$, a first-order correspondence language $\mathcal{L}_\nabla^1(\Phi)$ is generated from first-order variables x, y, z, \dots , unary predicates P_0, P_1, \dots for each propositional atom $p_0, p_1, \dots \in \Phi$, binary relation symbol(s) R_1, R_2 , and a unary relation symbol Q . The set of propositional atoms is constant, so we omit reference to Φ in the remainder.

First-order correspondence languages vary by the conditions imposed on the binary relations, and those conditions are determined by the interpretation supplied to \diamond_1 and \diamond_2 in \mathcal{L}_\diamond by a standard bi-modal Kripke frame. Otherwise, the translation operations are homomorphic for non-modal formulas. Define a bi-modal Kripke frame $\mathcal{F}^2 = (W \cup \wp(W), R_{\mathcal{N}}, R_\exists, \text{pt})$. The neighborhood function $\mathcal{N} \in \mathbb{F}$ is represented within \mathcal{F}^2 by:

$$\begin{aligned} R_{\mathcal{N}} &= \{(w, X) \in W \times \wp(W) \mid X \in \mathcal{N}(w)\} \\ R_\exists &= \{(X, w) \in \wp(W) \times W \mid w \in X\} \\ \text{pt} &= W. \end{aligned}$$

The \mathcal{L}_∇ to \mathcal{L}_\diamond step. Define a translation τ between \mathcal{L}_∇ and \mathcal{L}_\diamond as:

$$\begin{aligned} \perp^\tau &= \perp, \\ p^\tau &= p, \text{ for } p \in \Phi, \\ (\neg\varphi)^\tau &= \neg(\varphi^\tau), \\ (\varphi \vee \psi)^\tau &= (\varphi^\tau) \text{ or } (\psi^\tau), \\ (\nabla\varphi)^\tau &= \diamond_{\mathcal{N}} \Box_\exists (\varphi)^\tau. \end{aligned}$$

A global frame-validity preserving translation \blacklozenge between \mathcal{L}_∇ and \mathcal{L}_\diamond is defined by $\varphi^\blacklozenge = \text{pt} \rightarrow \varphi^\tau$ (Kracht and Wolter 1999, Hansen 2003).

The \mathcal{L}_\diamond to \mathcal{L}_∇^1 step. The local translation t between \mathcal{L}_\diamond and \mathcal{L}_∇^1 is defined in terms of the unary predicates $P_i \in \mathcal{L}_\nabla^1$, which are interpreted by their corresponding propositional variables $p_i \in \Phi$ as follows. $\Vdash_w^{\mathcal{M}} p^t = P(w)$ express that p is satisfied at world w in model \mathcal{M} , and this assertion is translated into first-order logic by $P(w)$. $p^t(w)$ abbreviates $\Vdash_w^{\mathcal{M}} p^t = P(w)$; $\neg(p^t(w))$ abbreviates $\nVdash_w^{\mathcal{M}} p^t$. Then:

$$\begin{aligned} (\perp)^t(w) &= x \neq x, \\ (p)^t(w) &= P(w), \\ (\neg\varphi)^t(w) &= \neg(\varphi^t(w)), \\ (\varphi \vee \psi)^t(w) &= \varphi^t(w) \vee \psi^t(w) \\ (\nabla\varphi)^t(w) &= \exists x(R_{\mathcal{N}}wx \wedge \forall y[R_\exists xy \rightarrow \varphi^t(y)]), \end{aligned}$$

where R_iab abbreviates $(a, b) \in R_i$.

The expression $(\varphi)^t(w)$ translates the assertion that φ is satisfied at world w within a model. To translate that φ is valid with respect to a class of models, a *global translation* function T translates the assertion that φ is satisfied at all worlds with respect to that class of models. The global translation T between \mathcal{L}_∇^1 and \mathcal{L}_\diamond defined by $(\varphi)^T(w) = \forall w(Q(w) \rightarrow (\varphi)^t(w))$ preserves frame validity (Hansen 2003).

3.3 THE AGM POSTULATES

To define AGM revision on this family of correspondence languages we adapt a strategy for normal monomodal logic (Gabbay et al. 2008) that requires (i) a sound and complete axiomatization of each classical modal system, (ii) a classical AGM revision operator.

Recall the AGM postulates for the revision operator, $*$. The AGM postulates hold for *any* consequence operation Cn defined on a classical propositional language that includes classical consequence, satisfies the Tarski closure conditions (idempotence, inclusion, and monotony), and satisfies disjunction in the premises (Alchourrón et al. 1985). Classical modal consequence is supraclassical, but obviously is not expressed within a purely propositional language. However, since *local* translations offer sufficient expressivity for revision and contraction operators, the first-order correspondent admits a propositional representation if the set of world $W \in \mathbb{F}$ is finite (Wheeler 2010). To ensure that closure and premise disjunction holds, we work with equivalent alternatives to (K*3) and (K*4), and (K*7) and (K*8).

$$(K*_{3,4}) \text{ If } K \cup \{\phi\} \text{ is consistent, then } K * \phi = Cn(K \cup \{\phi\});$$

$$(K*_{7,8}) Cn((K * \phi) \cup \{\psi\}) = K * (\phi \wedge \psi), \text{ when } \psi \text{ is consistent with } K * \phi.$$

Turn now to the definition of AGM revision in EM. Let $\Lambda^t(w)$ be the first-order local translation into \mathcal{L}_∇^1 of a classical monotonic modal theory, $\phi^t(w)$ and $\psi^t(w)$ first-order local translations of classical monotonic modal formulas ϕ and ψ , and $\mathcal{A}_{\mathcal{L}_\nabla}$ the (possibly empty) first-order characterization of classical monotonic modal system M.S, where S denotes a set (possibly empty) of modal schemata. Then:

$$\Lambda *_{\text{m}} \Psi = \{\phi : \Lambda^t(w) * \psi^t(w) \wedge \mathcal{A}_{\mathcal{L}_\nabla} \vdash \phi^t(w)\}.$$

We now have the following results.

Theorem 3.1. *The operator $*_{\text{m}}$ is an AGM revision operator.*

Corollary 3.2. *The operator $*_{\text{mn}}$ is AGM.*

The contraction operator $\dot{-}_{\text{mn}}$ can be defined from the revision operator $*_{\text{mn}}$ via the Harper identity (Harper 1977, Gärdenfors 1988). A theory Λ contracted by φ is given by

$$\Lambda \dot{-}_{\text{mn}} \varphi = \Lambda \cap (\Lambda *_{\text{mn}} \neg \varphi). \quad (2)$$

3.4 CONTRACTING KNOWLEDGE BASES IN THE EP FRAMEWORK

We focus here on one type of variability in Γ_δ , the effect of removing data from Γ_δ . Categorical statements in Γ_δ will correspond to ∇ formulae in the modal language \mathcal{L}_∇ interpreted by the EMN class of neighborhood models. Statistical statements of the form Eq. (1) are not included in the scope of this contraction operator, however, since statistical statements are specialized (Kyburg and Teng 2001). Contracting a knowledge base by a statistical statement is thus achieved by removing the statement from the knowledge base.

Evidential probability is formulated on a sorted first order language, with empirical predicates and mathematical predicates, whose domain of quantification consists of a finite set of empirical objects and the set of real numbers, respectively. Since our translation technique is for classical propositional modal logics, we note here that we are focusing on a restricted fragment of EP that can be mapped to propositional logic. This is analogous to the mapping of first order default theories to propositional default theories (Reiter 1980). Note also that evidential probability does not admit arbitrary formulas in the place of reference formulas and target formulas in the statistical statements. These restrictions are imposed to obtain a well-behaved formalism. Furthermore, our aim is robustness analysis of probability assignments from corrigible evidence, and the restrictions enable us to explore interesting and complex robustness issues in a more tractable setting.

Given a knowledge base Γ_δ consisting of categorical and statistical statements, let $\Gamma_\delta^{\%}$ be the set of statistical statements in Γ_δ . Let Λ be the formulas in $\Gamma_\delta \setminus \Gamma_\delta^{\%}$ expressed in the modal language \mathcal{L}_∇ ,

closed under the EMN consequence relation. Let $D_i(\Lambda)$ be a function that returns the set of formulas in Λ at modal depth i , stripped of their modal operators. The EP contraction operator $\dot{-}_{ep}$ is defined as follows. A knowledge base Γ_δ EP-contracted by a formula ϕ is given by

$$\Gamma_\delta \dot{-}_{ep} \phi = \begin{cases} \Gamma_\delta^{\%} \cup [\Gamma_\delta \cap D_1(\Lambda \dot{-}_{mn} \nabla \phi)] & \text{if } \phi \notin \Gamma_\delta^{\%}, \\ \Gamma_\delta \setminus \{\phi\} & \text{otherwise.} \end{cases} \quad (3)$$

EP contraction by a statistical statement is carried out by removing the statistical statement, if it exists, from the set. EP contraction by a categorical statement is carried out by first performing a contraction in EMN modal space with respect to the modal counterpart of the categorical statements, and then taking those singly-nested formulas (without their modal operators) in the contracted theory that exist in the original knowledge base Γ_δ and recombining them with the statistical statements in Γ_δ .

A remark on Eq. (3). The function $\text{Prob}(\cdot, \cdot)$ takes a set of non-modal, non-closed sentences in its second coordinate, but $\Lambda \dot{-}_{mn} \nabla \phi$ is a modal theory. The function D_1 returns a set of categorical statements, thus is in the right form, but may introduce new formulae that are modal theorems of EMN which were not originally in Γ_δ ; its meet with Γ_δ eliminates those. Finally, adding back the statistical statements $\Gamma_\delta^{\%}$ yields a full contracted knowledge base.

4 Robustness of Evidential Probability

Non-identical sets of candidate statistical statements may yield the same evidential probability assignment to a target statement χ . We are interested in characterizing the conditions under which the removal of a piece of evidence does or does not impact the probability assignment to χ , and if it does, the extent of change to the probability assignment. We say that a sentence ϕ is *inconsequential* to χ with respect to the background knowledge Γ_δ just in case the evidential probability of χ is invariant under contraction of Γ_δ by ϕ , that is, $\text{Prob}(\chi, \Gamma_\delta \dot{-}_{ep} \phi) = \text{Prob}(\chi, \Gamma_\delta)$. Such inconsequential statements contribute to the robustness of the inference.

4.1 Distance Measure

Before we investigate the robustness of evidential probability, first we need to construct a metric for measuring the distance between two evidential probability intervals. For point probabilities p_1 and p_2 , the distance is straightforwardly $d(p_1, p_2) = |p_2 - p_1|$. For interval probabilities, one possibility is the Hausdorff measure. Given two sets S_1 and S_2 and a distance measure d defined for points in the sets, the Hausdorff measure is the maximum distance of any point in one set to the closest point in the other set:

$$d(S_1, S_2) = \max\left\{ \sup_{p_1 \in S_1} \inf_{p_2 \in S_2} d(p_1, p_2), \sup_{p_2 \in S_2} \inf_{p_1 \in S_1} d(p_1, p_2) \right\}.$$

The Hausdorff measure is symmetrical, and for compact sets this value always exists. For one-dimensional intervals $I_1 = [l_1, u_1]$ and $I_2 = [l_2, u_2]$, this measure reduces to

$$d(I_1, I_2) = \max\{|l_1 - l_2|, |u_1 - u_2|\}.$$

The Hausdorff measure characterizes the worst case shortest distance between points in two sets. However, it is not entirely satisfactory for capturing the *change* between two intervals. For example, consider the following interval pairs.

$$\begin{aligned} & \{ [0.2, 0.4], [0.7, 0.9] \}; \\ & \{ [0.2, 0.4], [0.3, 0.9] \}; \\ & \{ [0.2, 0.4], [0.2, 0.9] \}. \end{aligned}$$

They all have the same Hausdorff distance, 0.5, but intuitively the changes between the interval pairs in the three cases are quite different. Thus, we will consider a less extreme measure than the worst case shortest distance for measuring the distance between two evidential probability intervals. We will employ a modified measure based on the Hausdorff measure.

$$d'(I_1, I_2) = (|l_1 - l_2| + |u_1 - u_2|)/2.$$

Instead of taking the maximum of the two component distances, we take the average of the two values. The new distances in the above cases are 0.5, 0.3 and 0.25 respectively, which captures a sense of the second interval being progressively “closer” to the first interval $[0.2, 0.4]$ in the pair in successive cases in the example.

4.2 Robustness Conditions and Change Conditions

Now we turn to the relationship between changes in the background knowledge Γ_δ and changes in the resulting evidential probability of a target statement χ . In general, as we add or remove statements from the background knowledge, the evidential probability interval of a statement may dilate or contract. It is not the case, for instance, that adding more statements to the background knowledge would necessarily lead to a new interval that is at least as tight as the original interval. There are however certain general conditions under which the evidential probability of a target statement is invariant or changes in a systematic way. We discuss some of these conditions here.

One of the simplest is that a statement that is not a member of the background knowledge is inconsequential. Contraction by such a statement $\phi \notin \Gamma_\delta$ does not impact the evidential probability of any statement with respect to Γ_δ , since $(\Gamma_\delta \dot{-}_{ep} \phi) = \Gamma_\delta$ and so $\text{Prob}(\chi, \Gamma_\delta \dot{-}_{ep} \phi) = \text{Prob}(\chi, \Gamma_\delta)$.

Below we consider the barebones setting in which robustness and change conditions are constructed based on the set of candidate and relevant statistical statements. Recall that given a target statement χ , a *candidate* statistical statement for χ is a statistical statement that could potentially determine the evidential probability of χ , and a *relevant* statistical statement for χ is a candidate statistical statement for χ that survives the three rules for resolving conflict in evidential probability.

Case 0: s_0 : a statistical statement that is not a candidate statistical statement for χ

Trivially such a statistical statement is inconsequential to χ with respect to the background knowledge since it does not bear on the probability of the target statement χ at all. Adding and removing s_0 , as well as changing its statistical interval, will not affect the evidential probability of χ .

Case 1: s_1 : a statistical statement whose associated interval does not conflict with the evidential probability interval of the target statement χ

Adding s_1 to the background knowledge would preserve the evidential probability interval of χ , since s_1 does not conflict with and thus cannot defeat any of the relevant statistical statements contributing to this evidential probability.

The evidential probability of χ may not be preserved, however, when such a statistical statement s_1 is removed from the corpus, since s_1 could have been used to defeat another statistical statement s'_1 that conflicts with the relevant statistical statements contributing to the evidential probability of χ . The new evidential probability would then have to be widened to accommodate the now undefeated s'_1 .

Case 2: s_2 : a candidate statistical statement whose associated interval is strictly within the intersection of the intervals associated with all the relevant statistical statements for the target statement χ

When adding s_2 to the background knowledge, the evidential probability of χ becomes the interval associated with s_2 . Since s_2 cannot conflict with any of the relevant statistical statements, it must be a relevant statistical statement itself and constitutes a single-element set closed under difference with respect to the relevant statistical statements. The cover of this set is shorter than and thus, by the principle of strength, is preferred to the covers of all other such sets.

Similarly, removing such an existing candidate statistical statement s_2 from the background knowledge would result in an expansion of the evidential probability interval of χ .

Case 3: s_3 : a statistical statement that does not conflict with any candidate statistical statement for χ

The statistical statement s_3 is obviously inconsequential to χ , since its interval cannot be a part of any set of relevant statistical statements that make up the minimal cover. Adding and removing s_3 , as well as changing its statistical interval in a way such that the new interval still does not conflict with any candidate statistical statement for χ , will not impact the probability of χ .

5 Conclusion

We have developed a contraction operator for evidential probability and discussed rudimentary robustness measures and robustness conditions. A more complete analysis would extend to the treatment of revision operators to handle variation of a knowledge base caused by altering statements within it, and identify tighter robustness conditions, in particular quantitative assessment of robustness under various circumstances.

Acknowledgement

We thank Senegal for motivating this paper.

References

- Adams, E. (1975). *The Logic of Conditionals*. Reidel, Dordrecht.
- Alchourrón, C., Gärdenfors, P., and Makinson, D. C. (1985). On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530.
- Arló-Costa, H. (2002). First order extensions of classical systems of modal logic; the role of the Barcan schemas. *Studia Logica*, 71(1):87–118.
- Chellas, B. (1980). *Modal Logic*. Cambridge University Press, Cambridge.
- Cozman, F. (2000). Credal networks. *Artificial Intelligence*, 120(2):199–233.
- Douven, I. and Williamson, T. (2006). Generalizing the lottery paradox. *The British Journal for the Philosophy of Science*, 57(4):755–779.
- Gabbay, D., Rodrigues, O., and Russo, A. (2008). Belief revision in non-classical logics. *The Review of Symbolic Logic*, 1(3):267–304.
- Gärdenfors, P. (1988). *Knowledge in Flux: modeling the dynamics of epistemic states*. MIT Press, Cambridge, Mass.
- Haenni, R., Romeyn, J.-W., Wheeler, G., and Williamson, J. (2010). *Probabilistic Logic and Probabilistic Networks*. Synthese Library. Springer, Dordrecht.

- Hansen, H. H. (2003). Monotonic modal logics. Master's thesis, University of Amsterdam, Institute for Logic, Language and Computation.
- Harper, W. L. (1977). Rational conceptual change. In *PSA 1976*, pages 462–494, Chicago. Philosophy of Science Association, University of Chicago Press.
- Hawthorne, J. and Makinson, D. (2007). The quantitative/qualitative watershed for rules of uncertain inference. *Studia Logica*, 86(2):247–297.
- Kracht, M. (1993). How completeness and correspondence theory got married. In de Rijke, M., editor, *Diamonds and Defaults*. Synthese Library.
- Kracht, M. and Wolter, F. (1999). Normal monomodal logics can simulate all others. *Journal of Symbolic Logic*, 64(1):99–138.
- Kraus, S., Lehman, D., and Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207.
- Kyburg, Jr., H. E. (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press, Middletown, CT.
- Kyburg, Jr., H. E. and Teng, C. M. (2001). *Uncertain Inference*. Cambridge University Press, Cambridge.
- Kyburg, Jr., H. E. and Teng, C. M. (2002). The logic of risky knowledge. In *Electronic Notes in Theoretical Computer Science*, volume 67. Elsevier Science.
- Levi, I. (1980). *The Enterprise of Knowledge*. MIT Press, Cambridge, MA.
- Paris, J. and Simmonds, R. (2009). O is not enough. *The Review of Symbolic Logic*, 2:298–309.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Francisco, CA.
- Pearl, J. (2000). *Causality*. Cambridge University Press, Cambridge.
- Reichenbach, H. (1949). *The Theory of Probability*. University of California Press, Berkeley.
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13:81–132.
- Seidenfeld, T. and Wasserman, L. (1993). Dilation for sets of probabilities. *The Annals of Statistics*, 21:1139–154.
- Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press, Cambridge, MA.
- Teng, C. M. (2007). Conflict and consistency. In Harper, W. L. and Wheeler, G., editors, *Probability and Inference: Essays in Honor of Henry E. Kyburg, Jr.* King's College Publications, London.
- Wheeler, G. (2010). AGM revision in classical modal logics. Technical report, CENTRIA - Center for Artificial Intelligence, Universidade Nova de Lisboa.
- Wheeler, G. and Williamson, J. (2009). Evidential probability and objective Bayesian epistemology. In Bandyopadhyay, P. and Forster, M., editors, *Handbook of the Philosophy of Statistics*. Elsevier Science.