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

Evidential probability (EP) (Kyburg and Teng, 2001) assigns a probability to an event given a knowledge base containing statements about logical relationships of classes of objects as well as statements about statistical frequencies pertaining to some of those classes. While the focus of evidential probability is the assignment of probability to a particular event, we might ask about the robustness of the inference backing an assignment. Evaluating robustness is straightforward in most probabilistic logics, since propositions are interpreted with respect to a full joint distribution and thus each proposition may be manipulated individually. An EP knowledge base includes logical relationships as well as probabilistic relationships, so alterations of an EP knowledge base include a qualitative belief revision/contraction problem. The classical modal system EMN provides a qualitative representation of (EP) knowledge bases (Kyburg and Teng, 2002), and since many monotone neighborhood structures may be simulated by bi-modal Kripke structures that admit a first-order correspondent, AGM revision and contraction operators are definable with respect to this (indirect) first-order correspondence. Specifically, system EMN has a first-order correspondent on which AGM revision and contraction operators are defined. Robustness of an EP assignment to a statement is a function of the variation of EP assignment to that event across varying evidence, where varying evidence is represented by a class of different contractions. Conditions and measures of robustness are proposed and discussed.

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 Corrigible acceptance

Behind *la connaissance risquée* is the theory of ε-acceptability. Briefly, a statement is accepted if its probability is higher than a threshold $1-\varepsilon$. Here the parameter ε is not an infinitesimal, but a fixed finite value analogous to the level of significance in statistical testing.[1] But although ε-acceptability is here conceived in terms of evidential probability, others have sought to explicate 'Lockean' (Foley 2009) threshold belief in terms of standard Bayesian machinery (Hawthorne and Bovens 1999). But this invites a world of trouble. An impossibility argument may be extracted from (Douven and Williamson 2006) to the effect that no coherent probabilistic modeling of rational acceptance of a sentence ϕ can be constructed on a logic satisfying axioms of System P (Kraus et al. 1990) for probabilities assigned to ϕ less than unity,[2] and recent explorations of a finite axiomatization of probabilistic consequence in terms

[1] This is in contrast to an approach taken by Ernest Adams (1975) and Judea Pearl (1988), where $1-\varepsilon$ denotes probabilities arbitrarily close to 1 that correspond to knowledge.

[2] Although Douven and Williamson do not mention System P nor Gabbay's (Gabbay 1985) result that identified *cumulative transitivity* and the corresponding cumulative consequence relations, the weakened form of transitivity they discuss in their footnote 2 is cumulative transitivity and the generality they mention suggests that they are discussing the class of cumulative non-monotonic logics.

of System O (Hawthorne and Makinson 2007), a weakening of System P, has been found not to admit a complete axiomatization by finite Horn rules (Paris and Simmonds 2009). In the wake of these results, some might think that one should give up on threshold-based views of belief altogether.

But this is a hasty conclusion, for evidential probability is *not* a species of Bayesian probability (Haenni et al. 2010, Seidenfeld 2007, Levi 2007), and there is a long history recording the clash between evidential probability and Bayesian methods.[3] EP sets out a logic, with a sound semantics, for assigning probability to a statement from an evidence base composed of logical formulae and statistical statements, but these statements do not necessarily collectively comport with the axioms of classical probability. For a start EP is interval-valued rather than point-valued. EP is a logic for approximate reasoning, one that is more similar to the theory of rough sets (Pawlak 1991) and systems of fuzzy logic (Dubois and Prade 1980) than to probabilistic logic (Haenni et al. 2010). There is much to say about what fragments of EP can be captured within other qualitative and quantitative uncertainty frameworks (Kyburg 1974, Levi 1977, Wheeler 2004, Wheeler and Damásio 2004, Kyburg et al. 2007, Seidenfeld 2007, Kyburg 2007, Haenni et al. 2010, Swift and Wheeler 2010), but some species of convex Bayesianism is not one of them. So, the recent discussion of attempts to axiomatize a *probabilistic* consequence relation is orthogonal to discussions of ε-acceptability with respect to evidential probability.

Evidential Probability

Evidential probability construes probability as a metalinguistic relation (on analogy of provability) and is defined rather than axiomatized. Following the treatment in (Kyburg and Teng 2001), let $\mathscr{L}$ be a guarded fragment of first-order logic (Andréka et al. 1998).[4] The domain of the probability function $\text{Prob}(\cdot,\cdot)$ is $\mathscr{L} \times \wp(\mathscr{L})$, 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 is accepted into Γ_δ if it runs a chance of error of no more than a small threshold value δ. Since evidential probability is interval-valued, a statement with probability $[l,u]$ is accepted into Γ_δ if the lower bound of its probability interval is above the acceptance threshold, that is, $l \geq 1-\delta$.

Categorical statements are logical formulas specifying what is accepted as true in the background knowledge, in particular relationships between the formulas mentioned in the statistical statements. *Statistical statements* in the language 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.[5] 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 $\%\vec{x}(\tau(\vec{x}),\rho(\vec{x}),[l,u])$ in Γ_δ where $\chi \leftrightarrow \tau(a)$ and $\rho(a)$ are both also in Γ_δ for some term a in $\mathscr{L}$. Given a statement χ, there are many statements of the form $\tau(a)$ known to be logically equivalent to χ,

[3] See in particular (Jeffrey 1956, Kyburg 1961, Carnap 1968, Jeffrey 1970, Kyburg 1970, Levi 1977, Seidenfeld 1979, Levi 1980, Harper 1982, Kyburg 1982, 1983, Seidenfeld 1992, Kyburg and Pittarelli 1996, Kyburg and Teng 2001, Harper and Wheeler 2007).

[4] A guarded fragment of first-order logic is a decidable fragment of first-order logic.

[5] The variables $\vec{x}$, ρ, τ, l, and u are all metalinguistic variables. A concrete example of a statistical statement might be $\%x(W(x),U_1(x),[.43,.51])$, which might express that the proportion of balls drawn from urn U_1 that are white is between .43 and .51.

where individual a is known to belong to some reference class ρ and where some statistical connection between ρ and τ is captured in the form of a statistical statement in Γ_δ. Each of these statistical statements is a candidate for determining the evidential probability of χ. The problem is that there may be many such candidate statistical statements, 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 $[l,u] \supseteq [l',u']$. Two statistical statements $\%\vec{x}(\tau(\vec{x}),\rho(\vec{x}),l,u)$ and $\%\vec{x}'(\tau'(\vec{x}'),\rho'(\vec{x}'),l',u')$ *conflict* iff their associated intervals $[l,u]$ and $[l',u']$ conflict.
>
> **[Cover]** The *cover* of a set of intervals is the shortest interval including all the intervals. The cover of a set of statistical statements is the cover of their associated 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'.
>
> Given a set of statistical statements K, a set of statistical statements K' is *closed under difference* with respect to K iff the intervals associated with K' are closed under difference with respect to the intervals associated with K.

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 while the second is based on the full joint distribution, 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; alternatively it is the intersection of all such covers.

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 Contracting Background Knowledge

Evidential probability assigns a unique probability to a sentence χ given background knowledge Γ_δ. Assessing the robustness of the probability assigned to χ involves examining the variability in probability assignments caused by varying the composition of Γ_δ. Here we explore variability caused by removing statements from Γ_δ, which calls for a contraction operation suitable for the structure of Γ_δ.

Although we know that the logical structure of Γ_δ is the classical modal system EMN, the AGM postulates for belief revision and contraction are devised for propositional languages. Thus, we propose to construct a contraction operator for Γ_δ by constructing a contraction operator for the class EMN of classical modal logics. 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 immediately via the Harper Identity.

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. But, since our translation technique is for classical propositional modal logics, we should note that we are focusing on a restricted fragment of EP that can be mapped to propositional logic. This restriction is not as strong as it may first appear, because evidential probability places restrictions on the kind of formulas that can take the place of reference formulas and target formulas in statistical formulas. Furthermore, our aim is robustness analysis of probability assignments from corrigible evidence, not the entire language. So, the mathematical statements in the knowledge base requiring the full expressive capacity of the first order language can be set aside.

3.1 Classical Model 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 $\mathscr{L}_\nabla(\Phi)$ is a pair $\mathbb{F} = (W, \mathscr{N})$ where

a) W is a non-empty set of worlds,

b) $\mathscr{N} : W \mapsto \wp(\wp(W))$ is a neighborhood function, i.e. $\mathscr{N}(w) \subseteq \wp(W)$, for each $w \in W$.

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

The satisfiability conditions for non-modal propositional formulas on neighborhood models are analogous to Kripke models, but modal necessity ($\nabla\varphi$) and possibility ($\Delta\varphi$) statements on neighborhood models are different. Like the normal modal logic (K) and its extensions, classical modal logics are based on the classical system (E) and the meaning of necessity statements in different classical systems is determined by the properties of neighborhood frames just as the meaning of necessity statements in different normal systems is determined by the properties of a Kripke frame. Let $\mathbb{M} = (W, \mathscr{N}, V)$ be a neighborhood model, w be a world in W, X a set of worlds, and $p \in \Phi$. Then:

- $\Vdash^{\mathbb{M}}_{w} \bot$ iff never
- $\Vdash^{\mathbb{M}}_{w} p$ iff $w \in V(p)$, for $p \in \Phi$
- $\nVdash^{\mathbb{M}}_{w} p$ iff $w \notin V(p)$
- $\Vdash^{\mathbb{M}}_{w} \varphi \vee \psi$ iff $w \in V(\varphi)$ or $w \in V(\psi)$
- $\Vdash^{\mathbb{M}}_{w} \nabla\varphi$ iff $(\exists X \in \mathscr{N}(w), \forall w^* \in X) : \Vdash^{\mathbb{M}}_{w^*} \varphi$
- $\Vdash^{\mathbb{M}}_{w} \Delta\varphi$ iff $(\forall X \in \mathscr{N}(w), \exists w^* \in X) : \Vdash^{\mathbb{M}}_{w^*} \varphi$

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),

(K) $\nabla(\varphi \to \psi) \to (\nabla\varphi \to \nabla\psi)$,

instances of (K) are not generally valid in non-standard (i.e., neighborhood) models (Chellas 1980). Specifically, while the following schemata are valid in all standard Kirpke models,

(C) $\nabla\varphi \wedge \nabla\psi \to \nabla(\varphi \wedge \psi)$,

(M) $\nabla(\varphi \wedge \psi) \to \nabla\varphi \wedge \nabla\psi$,

(N) $\nabla\top$,

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: $\mathscr{L}_\nabla$, a classical propositional monomodal language; $\mathscr{L}_\diamond$, a standard propositional polymodal language; and $\mathscr{L}^1_\nabla$, the final first-order translation language corresponding to $\mathscr{L}_\nabla$. 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 $\mathscr{L}_\nabla$, 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:

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

For the standard polymodal language $\mathscr{L}_\diamond(\Phi)$, a first-order correspondence language $\mathscr{L}^1_\nabla(\Phi)$ is generated from first-order variables $x, y, z, \ldots$, unary predicates $P_0, P_1, \ldots$ for each propositional atom $p_0, p_1, \ldots \in \Phi$, 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 $\mathscr{L}_\diamond$ by a standard bi-modal Kripke frame. Otherwise, the translation operations are homomorphic for non-modal formulas.

Define a bi-modal Kripke frame $\mathscr{F}^2 = (W \cup \wp(W), R_{\mathscr{N}}, R_{\ni}, \mathsf{pt})$. The neighborhood function $\mathscr{N} \in \mathbb{F}$ is represented within $\mathscr{F}^2$ by:

$$\begin{aligned} R_{\mathscr{N}} &= \{(w,X) \in W \times \wp(W) \mid X \in \mathscr{N}(w)\} \\ R_{\ni} &= \{(X,w) \in \wp(W) \times W \mid w \in X\} \\ \mathsf{pt} &= W. \end{aligned}$$

The $\mathscr{L}_\nabla$ to $\mathscr{L}_\diamond$ step. Define a translation τ between $\mathscr{L}_\nabla$ and $\mathscr{L}_\diamond$ as:

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

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

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

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

where $R_i ab$ 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 $\mathscr{L}^1_\nabla$ and $\mathscr{L}_\diamond$ 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, $*$, where $K = Cn(K)$, and φ, ψ are propositional formulas with respect to the propositional language $\mathscr{L}^{PL}$:

(K$*$1) $K * \phi$ is a belief set.

(K$*$2) $\phi \in (K * \phi)$.

(K∗3) $(K * \phi) \subseteq Cn(K \cup \{\phi\})$.

(K∗4) If $\neg\phi \notin K$, then $Cn(K \cup \{\phi\}) \subseteq (K * \phi)$.

(K∗5) $(K * \phi) = \mathscr{L}^{PL}$ only if $\phi \equiv \bot$.

(K∗6) If $\phi \equiv \psi$, then $(K * \phi) \equiv (K * \psi)$.

(K∗7) $K * (\phi \wedge \psi) \subseteq Cn((K * \phi) \cup \{\psi\})$.

(K∗8) If $\neg\psi \notin (K * \phi)$, then $Cn((K * \phi) \cup \{\psi\}) \subseteq K * (\phi \wedge \psi)$.

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

(K$*_{7,8}$) $Cn((K * \phi) \cup \{\psi\}) = K * (\phi \wedge \psi)$, when ψ 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 $\mathscr{L}^1_\nabla$ 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 $\mathscr{A}_{\mathscr{L}_\nabla}$ the (possibly empty) first-order characterization of classical monotonic modal system M.S, where S denotes a set (possibly empty) of modal schemata. Then:

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

We now have the following results.

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

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

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

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

3.4 Contracting knowledge bases

Assessing the robustness of the probability assigned to χ given Γ_δ involves varying the composition of Γ_δ to see the effect on the probability assigned to χ. We are focusing here on one type of variability in Γ_δ by considering the effect of missing data from Γ_δ. The contraction operator we defined is designed to handle non-statistical statements in Γ_δ, which are accepted statements whose probabilities are above the $1-\delta$ threshold for acceptance. These statements will correspond to ∇ formulae in the modal language $\mathscr{L}_\nabla$ 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.

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 $\mathscr{L}_\nabla$, 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

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

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 $\mathrm{Prob}(\cdot,\cdot)$ takes a set of non-modal, non-closed sentences in its second coordinate, but $\Lambda \mathbin{\dot{-}_{\mathsf{mn}}} \nabla\varphi$ 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

Recall our original question, which was whether learning that a piece of evidence is absent would impact the probability assigned to a target statement χ. It should be clear that not every contraction of the background knowledge will impact the evidential probability assignment of χ, even if that operation forces changes to the evidence: learning that Pluto is not a planet does not change the probability of Mrs. L.Q. Smith succumbing to sepsis.

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 a contraction 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, $\mathrm{Prob}(\chi,\, \Gamma_\delta \mathbin{\dot{-}_{ep}} \phi) = \mathrm{Prob}(\chi, \Gamma_\delta)$.

Such inconsequential statements contribute to the robustness of the inference. Indeed, this is a version of the first scenario we imagined in Section 1 in which retracting any one of the three evidence statements does not change the probability assignment to χ. Each of these statements is individually inconsequential to χ with respect to the background knowledge that includes the three statements.

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:

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

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.

$$\{\ [0.2, 0.4], [0.7, 0.9]\ \};$$
$$\{\ [0.2, 0.4], [0.3, 0.9]\ \};$$
$$\{\ [0.2, 0.4], [0.2, 0.9]\ \}.$$

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.

Andréka, H., van Benthem, J., and Németi, I. (1998). Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274.

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.

Carnap, R. (1968). On rules of acceptance. In Lakatos, I., editor, *The Problem of Inductive Logic*, pages 146–150. North-Holland Co., Amsterdam.

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 Philosophy of Science*, 57(4):755–779.

Dubois, D. and Prade, H. (1980). *Fuzzy Sets and Systems: Theory and Applications*. Kluwer, North Holland.

Foley, R. (2009). Beliefs, degrees of belief, and the lockean thesis. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 37–47. Synthese Library.

Gabbay, D. (1985). Theoretical foundations for nonmonotonic reasoning in expert systems. In Apt, K., editor, *Logics and Models of Concurrent Systems*. Springer-Verlag, Berlin.

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. and Wheeler, G., editors (2007). *Probability and Inference: Essays In Honor of Henry E. Kyburg, Jr.* King's College Publications, London.

Harper, W. L. (1977). Rational conceptual change. In *PSA 1976*, pages 462–494, Chicago. Philosophy of Science Association, University of Chicago Press.

Harper, W. L. (1982). Kyburg on direct inference. In Bogdan, R., editor, *Henry E. Kyburg and Isaac Levi*, pages 97–128. Kluwer Academic Publishers.

Hawthorne, J. and Bovens, L. (1999). The preface, the lottery, and the logic of belief. *Mind*, 108:241–264.

Hawthorne, J. and Makinson, D. (2007). The quantitative/qualitative watershed for rules of uncertain inference. *Studia Logica*, 86(2):247–297.

Jeffrey, R. (1956). Valuation and acceptance of scientific hypotheses. *Philosophy of Science*, 23(3):237–246.

Jeffrey, R. (1970). Dracula meets wolfman: Acceptance vs. partial belief. In Swain, M., editor, *Induction, Acceptance, and Rational Belief*, volume 26 of *Synthese Library*. Reidel, Dordrecht.

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. (1970). Two world views. In *Epistemology and Inference*, pages 18–27. Minnesota Press, Minneapolis.

Kyburg, Jr., H. E. (1974). *The Logical Foundations of Statistical Inference*. D. Reidel, Dordrecht.

Kyburg, Jr., H. E. (1982). Reply. In Bogdan, R., editor, *Henry E. Kyburg, Jr. and Isaac Levi*, pages 129–162. Kluwer Academic Publishers.

Kyburg, Jr., H. E. (1983). The reference class. *Philosophy of Science*, 50:374–397.

Kyburg, Jr., H. E. (2007). Bayesian inference with evidential probability. In Harper, W. and Wheeler, G., editors, *Probability and Inference: Essays in Honor of Henry E. Kyburg, Jr.*, pages 281–296. King's College, London.

Kyburg, Jr., H. E. and Pittarelli, M. (1996). Set-based Bayesianism. *IEEE Transactions on Systems, Man and Cybernetics A*, 26(3):324–339.

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.

Kyburg, Jr., H. E., Teng, C. M., and Wheeler, G. (2007). Conditionals and consequences. *Journal of Applied Logic*, 5(4):638–650.

Levi, I. (1977). Direct inference. *Journal of Philosophy*, 74:5–29.

Levi, I. (1980). *The Enterprise of Knowledge*. MIT Press, Cambridge, MA.

Levi, I. (2007). Probability logic and logical probability. In Harper, W. and Wheeler, G., editors, *Probability and Inference: Essays in Honor of Henry E. Kyburg, Jr.*, pages 255–266. College Publications.

Paris, J. and Simmonds, R. (2009). O is not enough. *The Review of Symbolic Logic*, 2:298–309.

Pawlak, Z. (1991). *Rough Sets: Theoretical Aspects of Reasoning about Data*. Kluwer, Dordrecht.

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.

Seidenfeld, T. (1979). *Philosophical Problems of Statistical Inference: Learning from R. A. Fisher*. D. Reidel, Dordrecht.

Seidenfeld, T. (1992). R. A. Fisher's fiducial argument and Bayes' theorem. *Statistical Science*, 7:358–368.

Seidenfeld, T. (2007). Forbidden fruit: When Epistemic Probability may not take a bite of the Bayesian apple. In Harper, W. and Wheeler, G., editors, *Probability and Inference: Essays in Honor of Henry E. Kyburg, Jr.* King's College Publications, London.

Spirtes, P., Glymour, C., and Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press, Cambridge, MA.

Swift, T. and Wheeler, G. (2010). Extending description logics to support statistical information. *International Journal of Approximate Reasoning*. to appear.

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. (2004). A resource bounded default logic. In Delgrande, J. and Schaub, T., editors, *NMR 2004*, pages 416–422.

Wheeler, G. (2010). AGM revision in classical modal logics. Technical report, CENTRIA - Center for Artificial Intelligence, Universidade Nova de Lisboa.

Wheeler, G. and Damásio, C. (2004). An implementation of statistical default logic. In Alferes, J. and Leite, J., editors, *Logics in Artificial Intelligence: Proceedings of the 9th European Conference (JELIA 2004)*, Lecture Notes in Artificial Intelligence, pages 121–133, Berlin. Springer-Verlag.

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.

Appendix

Theorems 3.1 The operator $*_{\mathsf{m}}$ is an AGM operator for the smallest monotonic logic, EM

Proof Let Λ be an EM-consistent monotonic modal theory, and ϕ, ψ and γ sentences in $\mathscr{L}_\nabla$. We show that $*_{\mathsf{m}}$ satisfies the AGM postulates. First, observe that M is the smallest classical monotonic modal system, which is equivalent to $\mathsf{EM.S}$, where $\mathsf{S} = \emptyset$. Hence, $\mathscr{A}_{\mathscr{L}_\nabla} = \emptyset$.

1. ($\Lambda*1$): $\Lambda *_{\mathsf{m}} \phi$ is a belief set.

   Since $\Lambda^t(w) * (\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla}))$ is closed under $\vdash$ by (K$*$1), then $\Lambda *_{\mathsf{m}} \phi$ is closed under $\vdash_{\mathsf{EM}}$.

2. ($\Lambda*2$): $\phi \in (\Lambda *_{\mathsf{m}} \phi)$.

   From ($K*2$) we have $\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla} \in \Lambda^t(w) * (\phi^t(w) \wedge \mathscr{A}_L)$. Since $\Lambda^t(w) * (\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla})$ is closed under $\vdash$, by (K$*$1), and $\vdash$ is reflexive, then $\Lambda^t(w) * (\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla}) \vdash \phi^t(w)$. So, $\phi \in (\Lambda *_{\mathsf{m}} \phi)$ by ($\Lambda*2$).

3. ($\Lambda*3,4$): If sentence ϕ is EM-consistent with Λ, then $\Lambda *_{\mathsf{m}} \phi$ is equal to the closure of $\{\Lambda \cup \{\phi\}\}$ under $\vdash_{\mathsf{EM}}$, written $C_{\mathsf{m}}(\Lambda \cup \{\phi\})$.

   First we make the following two observations.
   Observation 1. Recall that if Λ is an EM-consistent modal theory, then $\Lambda \nvdash_{\mathsf{EM}} \bot$ and there exists a monotone neighborhood model for Λ.
   Observation 2. If $\Lambda \cup \{\phi\}$ is consistent with respect to classical modal logic EM, then $\Lambda^t(w)$ is classically consistent with respect to its translation, $\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla}$. Since by hypothesis $\Lambda \cup \{\phi\}$ has a monotone neighborhood model, by Observation 1, there exists a classical first-order model of its translation, $\Lambda^t(w) \cup \{\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla}\}$.

   Suppose that Θ denotes the classical provability closure of the first-order translation from Observation 2, $\Lambda^t(w)*(\phi^t(w) \wedge \mathscr{A}_{\mathscr{L}_\nabla})$. We now show that if $\psi^t(w) \in \Theta$, then $\Lambda*_{\mathsf{m}}\phi \vdash \psi$.

   Suppose that $C_{\mathsf{m}}(\Lambda)$ is Λ closed under $\vdash_{\mathsf{EM}}$ and $\Lambda^t(w)$ is the first-order translation of Λ. We denote the corresponding $\mathscr{A}_{\mathscr{L}_\nabla}$-simulated closure in classical logic of the first-order translation by $Cn(\Lambda^t)$. There are two parts.

   (a) First, for any $\gamma \in \mathscr{A}_{\mathscr{L}_\nabla}$, if $\gamma^t \in Cn(\Lambda^t)$, then $\gamma \in \Lambda$. To see this, notice that $C_{\mathsf{m}}(\Lambda)$ is a maximally EM-consistent set, so $\gamma \in C_{\mathsf{m}}(\Lambda)$ iff $\Lambda \vdash_{\mathsf{EM}} \gamma$.
   Proof: Suppose that $\gamma \notin \Lambda$. Then, there is a classical monotone model satisfying $\Lambda \cup \{\neg\gamma\}$ and a translation of this into first-order logic. But on the first-order model for this translation $\gamma^t \notin Cn(\Lambda^t)$, which falsifies the hypothesis.

(b) Second, for a closed classical theory $Cn(\Lambda^t)$ s.t. $\mathscr{A}_{\mathscr{L}_\nabla} \subseteq Cn(\Lambda^t)$ and $\{\gamma : \gamma^t \in \Lambda^t\}$, then $\Lambda \vdash \gamma$ only if $\gamma^t \in Cn(\Lambda^t)$.
Proof: Suppose that $\gamma^t \notin Cn(\Lambda^t)$. Then there is a model of $\Lambda^t \cup \{\neg\gamma^t\}$, so there is classical monotone model satisfying $\Lambda \cup \{\neg\gamma\}$ which falsifies the hypothesis.

This concludes the proof.

4. $(\Lambda{*}5)$: $\Lambda *_{\mathsf{m}} \phi = \mathscr{L}_\nabla$ only if $\phi \equiv \bot$.

   Since Λ is an EM-consistent modal theory, $\Lambda \neq \mathscr{L}_\nabla$. So $\Lambda^t(w) \neq \mathscr{L}^1_\nabla$. So if $\Lambda^t(w) * \phi^t(w) = \mathscr{L}^1_\nabla$, then $\phi^t(w) = \bot$; thus $\phi \equiv \bot$.

5. $(\Lambda{*}6)$: If $\vdash_{\mathsf{EM}} \phi \equiv \psi$, then $\Lambda *_{\mathsf{m}} \phi \equiv \Lambda *_{\mathsf{m}} \psi$.

   If $\vdash_{\mathsf{EM}} \phi \equiv \psi$, then $\vdash \phi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}) \equiv \psi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}$. So, by (K$*$6), $\Lambda * (\phi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}) \equiv \Lambda * (\psi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla})$. Therefore, $\Lambda *_{\mathsf{m}} \phi \equiv \Lambda *_{\mathsf{m}} \psi$.

6. $(\Lambda{*}7,8)$: $\Lambda *_{\mathsf{m}} (\phi \wedge \psi) = C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\})$, when ψ is EM-consistent with $\Lambda *_{\mathsf{m}} \phi)$.

   Now we proceed in two parts.

   (a) $\Lambda *_{\mathsf{m}} (\phi \wedge \psi) \subseteq C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\})$: By $(\Lambda{*}1)$, $\Lambda *_{\mathsf{m}} (\phi \wedge \psi) = C_{\mathsf{m}}(\Lambda *_{\mathsf{m}} (\phi \wedge \psi))$. Suppose that $\gamma \in C_{\mathsf{m}}(\Lambda *_{\mathsf{m}} (\phi \wedge \psi))$. Then by the correspondence theorem $\gamma^t \in Cn(\Lambda^t * (\phi^t \wedge \psi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}))$. So $\gamma^t \in Cn(\Lambda^t * (\phi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}) \cup \{\psi^t\})$, by (K$*$7), and $\gamma \in C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\})$, by correspondence. Since γ is an arbitrary modal formula, $\Lambda *_{\mathsf{m}} (\phi \wedge \psi) \subseteq C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\})$.

   (b) $C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\}) \subseteq \Lambda *_{\mathsf{m}} (\phi \wedge \psi)$: Suppose that $\gamma \in C_{\mathsf{m}}(\Lambda *_{\mathsf{m}} \phi)$. Since γ is EM-consistent with $\Lambda *_{\mathsf{m}} \phi)$, $\gamma \in C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\})$. Thus, $\gamma^t \in Cn(\Lambda^t * (\phi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}) \cup \{\psi^t\})$, by the correspondence theorem, and $\gamma^t \in Cn(\Lambda^t * (\phi^t \wedge \psi^t \wedge \mathscr{A}_{\mathscr{L}_\nabla}))$, by (K$*$8). So, $\gamma \in C_{\mathsf{m}}(\Lambda *_{\mathsf{m}} (\phi \wedge \psi))$, by correspondence. Since γ is an arbitrary modal formula, $C_{\mathsf{m}}((\Lambda *_{\mathsf{m}} \phi) \cup \{\psi\}) \subseteq \Lambda *_{\mathsf{m}} (\phi \wedge \psi)$. □

Corollary 3.2. The operator $*_{\mathsf{mn}}$ is an AGM operator for the smallest monotonic logic, EMN

Proof. We observed that all instances of the modal schemas (N) and (M) are theorems the modal system EMN, and that instances of (N) and (M) are valid with respect to the class of neighborhood models satisfying the frame conditions (*m*) and (*n*) in Eq. (4).

$$\begin{array}{ll} (m) & \forall w \in W,\ \forall X_1, X_2 \subseteq W : (X_1 \subseteq X_2\ \wedge\ X_1 \in \mathscr{N}(w)) \rightarrow X_2 \in \mathscr{N}(w) \\ (n) & \forall w \in W : W \in \mathscr{N}(w) \end{array} \tag{4}$$

We just need to show how to effect an extension of $*_{\mathsf{m}}$ to $*_{\mathsf{mn}}$. But this case is handled by setting $\mathscr{A}_{\mathscr{L}_\nabla} = \{(n)\}$ in the proof of Thm 3.1. □