Conditional Syntax Splitting, Lexicographic Entailment and the Drowning Effect

Jesse Heyninck[1,2,3,*], Gabriele Kern-Isberner[4] and Thomas Meyer[2]

[1] *Vrije Universiteit Brussel, Belgium*
[2] *University of Cape Town and CAIR, South-Africa*
[3] *Open Universiteit Heerlen, the Netherlands*
[4] *Technische Universität Dortmund*

Abstract

Lexicographic inference [1] is a well-behaved and popular approach to reasoning with non-monotonic conditionals. In recent work we have shown that lexicographic inference satisfies syntax splitting, which means we can restrict our attention to parts of the belief base that share atoms with a given query. In this paper, we introduce the concept of conditional syntax splitting, inspired by the notion of conditional independence as known from probability theory. We show that lexicographic inference satisfies conditional syntax splitting, and connect conditional syntax splitting to several known properties from the literature on non-monotonic reasoning, including the drowning effect.

Keywords

Non-monotonic Reasoning, lexicographic inference, defeasible reasoning, non-monotonic logic, syntax splitting

1. Introduction

Lexicographic inference [1] is a well-known and popular approach to reasoning with non-monotonic conditionals, which has been applied in description logics [2], probabilistic description logics [3] and richer preferential languages [4]. It is seen as a logic of very high-quality, as it extends rational closure (also known as system Z) [5] and avoids the so-called *drowning problem*. This high quality seems to come at a cost, as reasoning on the basis of lexicographic inference is P^{NP}-complete, even when restricted to belief bases consisting of Horn-literal rules, i.e. rule bases where every rule's antecedent is a conjunction of atoms and every rule's consequent is a literal [6]. In previous work [7], we have shown that lexicographic inference satisfies *syntax splitting* [8]. Syntax splitting is a property of inference operators that requires that, for a belief base which can be split syntactically into two parts (i.e. there exists two sub-signatures such that every conditional in the belief base is built up entirely one of the two sub-signatures), restricting attention to the sub-signature does not result in a loss or addition of inferences. In other words, syntax splitting ensures we can safely restrict our attention to parts of the belief base that share atoms with a given query, thus seriously lessening the computational strain for many concrete queries. However, this presupposed that parts of a conditional belief base are syntactically independent, meaning that no common atoms are allowed. This might be an overly strong requirement, as the two parts of the belief base might have common elements. Consider the following example:

Example 1. Usually, **b**ikes are **c**hain-driven $(c|b)$, usually chain-driven bikes have multiple **g**ears $(g|c)$, and usually a bike frame consists of **f**our pipes $(f|b)$. The form of the frame is independent of whether a bike is chain driven and how many gears it has. However, syntax splitting as defined in [8] does not allow us to restrict attention to $\{(f|b)\}$ when we want to make inferences about the form of a bike frame, as the common atom b prevents us from splitting the belief base into two independent parts.

An intuitively related problem that was somewhat surprisingly shown to be independent of syntax splitting in [7] is the so-called *drowning problem*. It consists in the fact that under some inductive inference relations, abnormal individuals do not inherit any properties. It is best illustrated using the canonical Tweety-example:

Example 2 (The Drowning Problem)**.** The drowning problem is illustrated by using the following conditional belief base $\Delta = \{(f|b), (b|p), (\neg f|p), (e|b)\}$, which represents the Tweety-example, i.e. that **b**irds typically **f**ly, **p**enguins are typically birds, and penguins typically don't fly, together with the additional conditional "birds typically have b**e**aks". The drowning problem is constituted by the fact that some inductive inference operators, such as system Z, do not allow to infer that penguins typically have beaks ($p \mathrel{|\!\sim}^{Z}_{\Delta} b$), i.e. the fact that penguins are abnormal when it comes to flying *drowns* inferences about penguins' beaks. It is well-known that lexicographic inference does not suffer from the drowning problem.

The drowning problem seems to be related to syntax splitting. Intuitively, $\{(e|b)\}$ is unrelated to the rest of the belief base, in the sense that having beaks has nothing to do with flying or having wings, as long as we know we are talking about birds. However, (unconditional) syntax splitting does not allow to capture this kind of independence, since the atom b prohibits the belief base from being split into information about flying and wings on the one hand, and information about beaks on the other hand. It is exactly this kind of *conditional* independencies between conditionals that we seek to formally capture and study in this paper. In more detail, the contributions of the paper are the following:

1. we introduce and study the notion of *conditional splitting* of a belief base, a property of conditional belief bases, and generalize the concept of syntax splitting, a property of inductive inference operators, to *conditional syntax splitting*, thus bringing the idea of conditional independence into the realm of inductive inference operators;
2. we show that lexicographic entailment satisfies conditional syntax splitting;
3. we argue that the drowning effect can be seen as a violation of conditional syntax splitting; and

NMR 2022: 20th International Workshop on Non-Monotonic Reasoning, August 07–09, 2022, Haifa, Israel
* Corresponding author.
jesse.heyninck@ou.nl (J. Heyninck);
gabriele.kern-isberner@cs.tu-dortmund.de (G. Kern-Isberner);
tmeyer@cair.org.za (T. Meyer)
0000-0002-3825-4052 (J. Heyninck)

4. we show how Lehmann's so-called *desirable closure properties* [1] can be derived from conditional syntax splitting.

Outline of this Paper: We first state all the necessary preliminaries in Section 2 on propositional logic (Section 2.1), reasoning with non-monotonic conditionals (Section 2.2), inductive inference (Section 2.3), System Z (Section 2.4) and lexicographic inference (Section 2.5). In Section 3 we define and study the concept of conditional syntax splitting. In Section 4, we show that lexicographic inference satisfies conditional syntax splitting. In Sections 5 and 6, we show how properties of inductive inference operators previously only discussed informally, namely the drowning effect (Section 5) and the properties introduced by Lehmann 1995 (Section 6) can be seen as special cases of conditional syntax splitting. Finally, we discuss related work in Section 7 and conclude in Section 8.

2. Preliminaries

In the following, we briefly recall some general preliminaries on propositional logic, and technical details on inductive inference.

2.1. Propositional Logic

For a set At of atoms let $\mathcal{L}(\mathsf{At})$ be the corresponding propositional language constructed using the usual connectives $\wedge$ (*and*), $\vee$ (*or*), $\neg$ (*negation*), $\rightarrow$ (*material implication*) and $\leftrightarrow$ (*material equivalence*). A (classical) *interpretation* (also called *possible world*) ω for a propositional language $\mathcal{L}(\mathsf{At})$ is a function $\omega : \mathsf{At} \rightarrow \{\top, \bot\}$. Let $\Omega(\mathsf{At})$ denote the set of all interpretations for At. We simply write Ω if the set of atoms is implicitly given. An interpretation ω *satisfies* (or is a *model* of) an atom $a \in \mathsf{At}$, denoted by $\omega \models a$, if and only if $\omega(a) = \top$. The satisfaction relation $\models$ is extended to formulas as usual. As an abbreviation we sometimes identify an interpretation ω with its *complete conjunction*, i. e., if $a_1, \ldots, a_n \in \mathsf{At}$ are those atoms that are assigned $\top$ by ω and $a_{n+1}, \ldots, a_m \in \mathsf{At}$ are those propositions that are assigned $\bot$ by ω we identify ω by $a_1 \ldots a_n \overline{a_{n+1}} \ldots \overline{a_m}$ (or any permutation of this). For $X \subseteq \mathcal{L}(\mathsf{At})$ we also define $\omega \models X$ if and only if $\omega \models A$ for every $A \in X$. Define the set of models $\mathsf{Mod}(X) = \{\omega \in \Omega(\mathsf{At}) \mid \omega \models X\}$ for every formula or set of formulas X. A formula or set of formulas X_1 *entails* another formula or set of formulas X_2, denoted by $X_1 \models X_2$, if $\mathsf{Mod}(X_1) \subseteq \mathsf{Mod}(X_2)$. Where $\theta \subseteq \Sigma$, and $\omega \in \Omega(\Sigma)$, we denote by ω^θ the restriction of ω to θ, i.e. ω^θ is the interpretation over Σ^θ that agrees with ω on all atoms in θ. Where $\Sigma_i, \Sigma_j \subseteq \Sigma$, $\Omega(\Sigma_i)$ will also be denoted by Ω_i for any $i \in \mathbb{N}$, and likewise $\Omega_{i,j}$ we denote $\Omega(\Sigma_i \cup \Sigma_j)$ (for $i, j \in \mathbb{N}$). Likewise, for some $X \subseteq \mathcal{L}(\Sigma_i)$, we define $\mathsf{Mod}_i(X) = \{\omega \in \Omega_i \mid \omega \models X\}$.

2.2. Reasoning with Nonmonotonic Conditionals

Given a language $\mathcal{L}$, conditionals are objects of the form $(B|A)$ where $A, B \in \mathcal{L}$. The set of all conditionals based on a language $\mathcal{L}$ is defined as: $(\mathcal{L}|\mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}$. We follow the approach of [9] who considered conditionals as *generalized indicator functions* for possible worlds resp. propositional interpretations ω:

$$((B|A))(\omega) = \begin{cases} 1 & : \quad \omega \models A \wedge B \\ 0 & : \quad \omega \models A \wedge \neg B \\ u & : \quad \omega \models \neg A \end{cases} \tag{1}$$

where u stands for *unknown* or *indeterminate*. In other words, a possible world ω *verifies* a conditional $(B|A)$ iff it satisfies both antecedent and conclusion ($(B|A)(\omega) = 1$); it *falsifies, or violates* it iff it satisfies the antecedence but not the conclusion ($(B|A)(\omega) = 0$); otherwise the conditional is *not applicable*, i. e., the interpretation does not satisfy the antecedent ($(B|A)(\omega) = u$). We say that ω *satisfies* a conditional $(B|A)$ iff it does not falsify it, i.e., iff ω satisfies its *material counterpart* $A \rightarrow B$. Given a total preorder (in short, TPO) $\preceq$ on possible worlds, representing relative plausibility, $A \preceq B$ iff $\omega \preceq \omega'$ for some $\omega \in \min_\preceq(\mathsf{Mod}(A))$ and some $\omega' \in \min_\preceq(\mathsf{Mod}(B))$. This allows for expressing the validity of defeasible inferences via stating that $A \mathrel{|\!\sim}_\preceq B$ iff $(A \wedge B) \prec (A \wedge \neg B)$ [10]. As is usual, we denote $\omega \preceq \omega'$ and $\omega' \preceq \omega$ by $\omega \approx \omega'$ and $\omega \preceq \omega'$ and $\omega' \not\preceq \omega$ by $\omega \prec \omega'$ (and similarly for formulas). We can *marginalize* total preorders and even inference relations, i.e., restricting them to sublanguages, in a natural way: If $\Theta \subseteq \Sigma$ then any TPO $\preceq$ on $\Omega(\Sigma)$ induces uniquely a *marginalized TPO* $\preceq_{|\Theta}$ on $\Omega(\Theta)$ by setting

$$\omega_1^\Theta \preceq_{|\Theta} \omega_2^\Theta \text{ iff } \omega_1^\Theta \preceq \omega_2^\Theta. \tag{2}$$

Note that on the right hand side of the *iff* condition above $\omega_1^\Theta, \omega_2^\Theta$ are considered as propositions in the superlanguage $\mathcal{L}(\Omega)$, hence $\omega_1^\Theta \preceq \omega_2^\Theta$ is well defined [11].

Similarly, any inference relation $\mathrel{|\!\sim}$ on $\mathcal{L}(\Sigma)$ induces a *marginalized inference relation* $\mathrel{|\!\sim}_{|\Theta}$ on $\mathcal{L}(\Theta)$ by setting

$$A \mathrel{|\!\sim}_{|\Theta} B \text{ iff } A \mathrel{|\!\sim} B \tag{3}$$

for any $A, B \in \mathcal{L}(\Theta)$.

An obvious implementation of total preorders are *ordinal conditional functions (OCFs)*, (also called *ranking functions*) $\kappa : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$. [12]. They express degrees of (im)plausibility of possible worlds and propositional formulas A by setting $\kappa(A) := \min\{\kappa(\omega) \mid \omega \models A\}$. A conditional $(B|A)$ is accepted by κ iff $A \mathrel{|\!\sim}_\kappa B$ iff $\kappa(A \wedge B) < \kappa(A \wedge \neg B)$.

2.3. Inductive Inference Operators

In this paper, we will be interested in inference relations $\mathrel{|\!\sim}_\Delta$ parametrized by a conditional belief base Δ. In more detail, such inference relations are *induced by* Δ, in the sense that Δ serves as a starting point for the inferences in $\mathrel{|\!\sim}_\Delta$. We call such operators *inductive inference operators*:

Definition 1 ([8]). An *inductive inference operator (from conditional belief bases)* is a mapping $\mathbf{C}$ that assigns to each conditional belief base $\Delta \subseteq (\mathcal{L}|\mathcal{L})$ an inference relation $\mathrel{|\!\sim}_\Delta$ on $\mathcal{L}$ that satisfies the following basic requirement of *direct inference*:

DI If Δ is a conditional belief base and $\mathrel{|\!\sim}_\Delta$ is an inference relation that is induced by Δ, then $(B|A) \in \Delta$ implies $A \mathrel{|\!\sim}_\Delta B$.

Examples of inductive inference operators include system P [13], System Z ([5], see Section 2.4), lexicographic inference ([1], see Section 2.5) and c-representations ([14].

As already indicated in the previous subsection, inference relations can be obtained on the basis of TPOs respectively OCFs:

Definition 2. A *model-based inductive inference operator for total preorders* (on Ω) is a mapping $\mathbf{C}^{tpo}$ that assigns to each conditional belief base Δ a total preorder $\preceq_\Delta$ on Ω s.t. $A \mathrel{|\!\sim}_{\preceq_\Delta} B$ for every $(B|A) \in \Delta$ (i.e. s.t. **DI** is ensured). A *model-based inductive inference operator for OCFs* (on Ω) is a mapping $\mathbf{C}^{ocf}$ that assigns to each conditional belief base Δ an OCF κ_Δ on Ω s.t. Δ is accepted by κ_Δ(i.e. s.t. **DI** is ensured).

Examples of inductive inference operators for OCFs System Z ([5], see Sec. 2.4) and c-representations ([14], whereas lexicographic inference ([1], see Sec. 2.5) is an example of an inductive inference operator for TPOs.

To define the property of *syntax splitting* [8], we assume a conditional belief base Δ that can be split into subbases Δ^1, Δ^2 s.t. $\Delta^i \subset (\mathcal{L}_i|\mathcal{L}_i)$ with $\mathcal{L}_i = \mathcal{L}(\Sigma_i)$ for $i = 1, 2$ s.t. $\Sigma_1 \cap \Sigma_2 = \emptyset$ and $\Sigma_1 \cup \Sigma_2 = \Sigma$, writing:

$$\Delta = \Delta^1 \bigcup_{\Sigma_1, \Sigma_2} \Delta^2$$

whenever this is the case.

Definition 3 (Independence (**Ind**), [8])**.** An inductive inference operator $\mathbf{C}$ satisfies (**Ind**) if for any $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2$ and for any $A, B \in \mathcal{L}_i$, $C \in \mathcal{L}_j$ $(i, j \in \{1, 2\}, j \neq i)$,

$$A \mid\!\sim_\Delta B \text{ iff } AC \mid\!\sim_\Delta B$$

Definition 4 (Relevance (**Rel**), [8])**.** An inductive inference operator $\mathbf{C}$ satisfies (**Rel**) if for any $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2$ and for any $A, B \in \mathcal{L}_i$ $(i \in \{1, 2\})$,

$$A \mid\!\sim_\Delta B \text{ iff } A \mid\!\sim_{\Delta^i} B.$$

Definition 5 (Syntax splitting (**SynSplit**), [8])**.** An inductive inference operator $\mathbf{C}$ satisfies (**SynSplit**) if it satisfies (**Ind**) and (**Rel**).

Thus, **Ind** requires that inferences from one sub-language are independent from formulas over the other sublanguage, if the belief base splits over the respective sublanguages. In other words, information on the basis of one sublanguage does not influences inferences made in the other sublanguage. **Rel**, on the other hand, restricts the scope of inferences, by requiring that inferences in a sublanguage can be made on the basis of the conditionals in a conditional belief base formulated on the basis of that sublanguage. **SynSplit** combines these two properties.

2.4. System Z

We present system Z defined in [5] as follows. A conditional $(B|A)$ is tolerated by a finite set of conditionals Δ if there is a possible world ω with $(B|A)(\omega) = 1$ and $(B'|A')(\omega) \neq 0$ for all $(B'|A') \in \Delta$, i.e. ω verifies $(B|A)$ and does not falsify any (other) conditional in Δ. The Z-partitioning $(\Delta_0, \ldots, \Delta_n)$ of Δ is defined as:

- $\Delta_0 = \{\delta \in \Delta \mid \Delta \text{ tolerates } \delta\}$;
- $\Delta_1, \ldots, \Delta_n$ is the Z-partitioning of $\Delta \setminus \Delta_0$.

For $\delta \in \Delta$ we define: $Z_\Delta(\delta) = i$ iff $\delta \in \Delta_i$ and $(\Delta_0, \ldots, \Delta_n)$ is the Z-partioning of Δ. Finally, the ranking function κ^Z_Δ is defined via: $\kappa^Z_\Delta(\omega) = \max\{Z(\delta) \mid \delta(\omega) = 0, \delta \in \Delta\} + 1$, with $\max \emptyset = -1$. The resulting inductive inference operator $C^{ocf}_{\kappa^Z_\Delta}$ is denoted by C^Z.

In the literature, system Z has also been called *rational closure* [15]. An inference relation $\mid\!\sim_\Delta$ based on Δ s.t. $A \mid\!\sim^Z_\Delta B$ implies $A \mid\!\sim_\Delta B$ is called *RC-extending* [16]. An *RC-extending* inference relation has also been called a *refinement of System Z* [17]. We call an inductive inference operator $\mathbf{C}$ *RC-extending* iff every $\mathbf{C}(\Delta)$ is RC-extending.

We now illustrate OCFs in general and System Z in particular with the well-known "Tweety the penguin"-example.

Example 3. Let $\Delta = \{(f|b), (b|p), (\neg f|p)\}$ be a sub-base of belief base used in Example 2. This conditional belief base has the following Z-partitioning: $\Delta_0 = \{(f|b)\}$ and $\Delta_1 = \{(b|p), (\neg f|p)\}$. This gives rise to the following κ^Z_Δ-ordering over the worlds based on the signature $\{b, f, p\}$:

ω	κ^Z_Δ	ω	κ^Z_Δ	ω	κ^Z_Δ	ω	κ^Z_Δ
pbf	2	$pb\overline{f}$	1	$p\overline{b}f$	2	$p\overline{b}\,\overline{f}$	2
$\overline{p}bf$	0	$\overline{p}b\overline{f}$	1	$\overline{p}\overline{b}f$	0	$\overline{p}\overline{b}\,\overline{f}$	0

As an example of a (non-)inference, observe that e.g. $\top \mid\!\sim^Z_\Delta \neg p$ and $p \wedge f \not\mid\!\sim^Z_\Delta b$.

2.5. Lexicographic Entailment

We recall lexicographic inference as introduced by [1]. For some conditional belief base Δ, the order $\preceq^{\mathsf{lex}}_\Delta$ is defined as follows: Given $\omega \in \Omega$ and $\Delta' \subseteq \Delta$, $V(\omega, \Delta') = |(\{(B|A) \in \Delta' \mid (B|A)(\omega) = 0\}|$. Given a set of conditionals Δ Z-partitioned in $(\Delta_0, \ldots, \Delta_n)$, the *lexicographic vector* for a world $\omega \in \Omega$ is the vector $\mathsf{lex}(\omega) = (V(\omega, \Delta_0), \ldots, V(\omega, \Delta_n))$. Given two vectors $(x_1, \ldots, x_n)$ and $(y_1, \ldots, y_n)$, $(x_1, \ldots, x_n) \preceq^{\mathsf{lex}} (y_1, \ldots, y_n)$ iff there is some $j \leqslant n$ s.t. $x_k = y_k$ for every $k > j$ and $x_j \leqslant y_j$. $\omega \preceq^{\mathsf{lex}}_\Delta \omega'$ iff $\mathsf{lex}(\omega) \preceq^{\mathsf{lex}} \mathsf{lex}(\omega')$. The resulting inductive inference operator $C^{tpo}_{\preceq^{\mathsf{lex}}}$ will be denoted by C^{lex} to avoid clutter.

In [1], lexicographic inference was shown to be RC-extending (for finite conditional belief bases):

Proposition 1 ([1, Theorem 3])**.** For any $A \in \mathcal{L}$ s.t. $\kappa^Z_\Delta(A)$ is finite, then $A \mid\!\sim^Z_\Delta B$ implies $A \mid\!\sim^{\mathsf{lex}}_\Delta B$.

Example 4 (Example 3 ctd.)**.** For the Tweety belief base Δ as in Example 3 we obtain the following $\mathsf{lex}(\omega)$-vectors:

ω	$\mathsf{lex}(\omega)$	ω	$\mathsf{lex}(\omega)$	ω	$\mathsf{lex}(\omega)$	ω	$\mathsf{lex}(\omega)$
pbf	(0,1)	$pb\overline{f}$	(1,0)	$p\overline{b}f$	(0,2)	$p\overline{b}\,\overline{f}$	(0,1)
$\overline{p}bf$	(0,0)	$\overline{p}b\overline{f}$	(1,0)	$\overline{p}\overline{b}f$	(0,0)	$\overline{p}\overline{b}\,\overline{f}$	(0,0)

The lex-vectors are ordered as follows:

$$(0,0) \prec^{\mathsf{lex}} (1,0) \prec^{\mathsf{lex}} (0,1) \prec^{\mathsf{lex}} (0,2).$$

Observe that e.g. $\top \mid\!\sim^{\mathsf{lex}}_\Delta \neg p$ (since $\mathsf{lex}(\top \wedge \neg p) = (0,0) \prec^{\mathsf{lex}} \mathsf{lex}(\top \wedge p) = (1,0)$) and $p \wedge f \mid\!\sim^{\mathsf{lex}}_\Delta b$.

3. Conditional Syntax Splitting

We now introduce a conditional version of syntax splitting. A first central idea is the syntactical notion of *conditional splitting*, a property of belief bases.

Definition 6. We say a conditional belief base Δ can be *split into subbases* Δ_1,Δ_2 *conditional on a sub-alphabet* Σ_3, if $\Delta_i \subset (\mathcal{L}(\Sigma_i \cup \Sigma_3) \mid \mathcal{L}(\Sigma_i \cup \Sigma_3))$ for $i = 1, 2$ s.t. Σ_1, Σ_2 and Σ_3 are pairwise disjoint and $\Sigma = \Sigma_1 \cup \Sigma_2 \cup \Sigma_3$, writing:

$$\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$$

Intuitively, a conditional belief base can be split into Σ_1 and Σ_2 conditional on Σ_3, if every conditional is built up from atoms in $\Sigma_1 \cup \Sigma_3$ or atoms in $\Sigma_2 \cup \Sigma_3$.

The above notion of conditional syntax splitting, however, is too strong, in the sense that it does not warrant satisfaction of conditional variants of relevance and independence (we will define them in formal detail below) for lexicographic inference. The underlying problem is that toleration might not be respected by conditional belief bases that conditionally split:

Example 5. Let $\Delta = \{(x|b), (\neg x|a), (c|a \wedge b)\}$. Then

$$\Delta = \{(x|b), (\neg x|a)\} \bigcup_{\{x\},\{c\}} \{(c|a \wedge b)\} \mid \{a, b\}$$

However, this notion of purely syntactical conditional independence is not reflected on the level of tolerance (and therefore entailment). Indeed, $\{(c|a \wedge b)\}$ (trivially) tolerates itself, i.e. $Z_{\{(c|a\wedge b)\}}(c|a \wedge b) = 0$, yet Δ does not tolerate $(c|a \wedge b)$, i.e. $Z_\Delta(c|a \wedge b) = 1$.

This means that for system Z and lexicographic entaiment, conditional relevance (now only introduced informally) is violated for this belief base. In more detail, even though $\Delta = \{(x|b), (\neg x|a)\} \bigcup_{\{x\},\{c\}} \{(c|a \wedge b)\} \mid \{a, b\}$, we have e.g. $\top \not\mathrel{|\!\sim}^{\mathsf{lex}}_{\{(c|a\wedge b)\}} \neg(a \wedge b)$ whereas $\top \mathrel{|\!\sim}^{\mathsf{lex}}_{\Delta} \neg(a \wedge b)$ (and likewise for system Z).

What happens here is that $(x|b)$ and $(\neg x|a)$ act as "constraints" on a and b being true together, which on its turn is needed for $(c|a \wedge b)$ to be tolerated. In other words, pure syntactic conditional splitting is not reflected on the semantic level (in contradistinction to unconditional splitting). We can exclude such cases by using the following weaker notion of safe conditional syntax splitting:

Definition 7. A conditional belief base $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ can be *safely split into subbases* Δ_1, Δ_2 *conditional on a sub-alphabet* Σ_3, writing:

$$\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3,$$

if for every $\omega^3 \in \Omega(\Sigma_i \cup \Sigma_3)$, there is a $\omega^j \in \Omega(\Sigma^j)$ s.t. $\omega^j\omega^3 \not\models \bigvee_{(F|E)\in\Delta^j} E \wedge \neg F$ (for $i, j = 1, 2$ and $i \neq j$).

The notion of safe splitting is explained as follows: Δ can be safely split into Δ_1 and Δ_2 conditional on Σ_3 if it can be split in Δ_1 and Δ_2 conditional on Σ_3, and additionally, for every world $\omega^i\omega^3$ in the subsignature $\Sigma_i \cup \Sigma_3$, we can find a world ω^j in the subsignature Σ_j ($i, j = 1, 2$ and $j \neq i$) s.t. no conditional $\delta \in \Delta^j$ is falsified by $\omega^i\omega^j\omega^3$ (or, equivalently, by $\omega^j\omega^3$). We will show some more syntactical formulated conditions that ensure safe splitting below.

We argue here that safe splitting faithfully captures independences of two conditional belief bases conditional on a subsignature Σ_3. Indeed, safe splitting requires that (1) all conditionals are built up from the sub-signatures $\Sigma_1 \cup \Sigma_3$ or $\Sigma_2 \cup \Sigma_3$ (i.e. $\Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$), and (2) that any information on $\Sigma_i \cup \Sigma_3$ is compatible with Δ^j, i.e. no world $\omega^i\omega^3$ causes a conditional in Δ^j to be violated. In other words, toleration with respect to Δ^j is independent of Δ^i.

We now delineate some more syntactic conditions that ensure safe syntax splitting. These conditions are typically easier to check, and might reasonably be expected to hold for certain natural language scenarios. For example, if it holds that (1) $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, (2) all antecedents and consequents (of conditionals in Δ) using elements of the common sub-alphabet Σ_3 are equivalent, and (3) all material versions of the conditional sub-base Δ^i are consistent with the set of consequents of the conditionals whose antecedent uses atoms in the common sub-alphabet Σ_3, then Δ can be safely split into Δ^1 and Δ^2 conditional on Σ_3.

Proposition 2. Let a conditional belief base $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ be given. If there is a $C \in \mathcal{L}(\Sigma_3)$ s.t. for every conditional in $(B|A) \in \Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$:

1. $B \in \mathcal{L}(\Sigma_1) \cup \mathcal{L}(\Sigma_2)$, or $B \equiv C$.
2. $A \in \mathcal{L}(\Sigma_1) \cup \mathcal{L}(\Sigma_2)$, or $A \equiv C$.
3. $\bigwedge_{(G|H)\in\Delta^i} H \to G \not\vdash \bigvee\{\neg F \mid (F|C') \in \Delta^i, C' \equiv C\}$ for $i = 1, 2$.[1]

Then $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$.

Proof. Consider some $\omega^3 \in \Omega(\Sigma_3)$. If $\omega^3 \not\models C$ we are done. Suppose therefore $\omega^3 \models C$. With the third condition, there is an $\omega^2\omega'^3 \in \mathsf{Mod}(\bigwedge_{(G|H)\in\Delta^2} H \to G)$ s.t. $\omega^2\omega'^3 \models F$ for every $(F|C') \in \Delta^2$ s.t. $C' \equiv C$. Notice that, in view of the first two conditions, for any $(B|A) \in \Delta^j$, either $\omega^2 \models A \to B$ or $B \equiv C$ or $A \equiv C$. Furthermore, in the case where $A \equiv C$, $\omega^2\omega'^3 \models B$. Since $B \in \mathcal{L}(\Sigma_2)$ or $B \equiv C$, also $\omega^2\omega^3 \models B$. Altogether, $\omega^2\omega^3 \models \bigwedge_{(G|H)\in\Delta^2} H \to G$. Thus, $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$. □

A simpler case of this condition is a belief base where all antecedents derive from the common alphabet Σ_3 and all consequents derive from either Σ_1 or Σ_2.

Notice that e.g. the conditional belief base from Example 2 has the form described in Proposition 2:

Example 6. Consider again Δ from Example 2, and let $\Sigma_1 = \{f, p\}$, $\Sigma_2 = \{e\}$ and $\Sigma_3 = \{b\}$. Observe that:

$$\Delta = \{(f|b), (b|p), (\neg f|p)\} \bigcup_{\Sigma_1,\Sigma_2} \{(e|b)\} \mid \Sigma_3.$$

Furthermore, the first two items in Proposition 2 are satisfied as every conditional is either completely on the basis of the alphabet $\{f, p\}$ or has as an antecedent or a consequent b. Finally, the last condition is satisfied as $\{b \to e\} \not\models \neg e$ and $\{b \to f, p \to b, p \to \neg f\} \not\models f \vee \neg b$. We thus see that $\Delta = \{(f|b), (b|p), (\neg f|p)\} \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \{(e|b)\} \mid \Sigma_3$.

The bicycle example is also of this form:

Example 7. Consider again Δ from Example 1. We see that:

$$\{(c|b), (g|c)\} \bigcup^{\mathsf{s}}_{\{g,c\},\{f\}} \{(f|b)\} \mid \{b\}.$$

Remark 1. Note that a weaker prerequisite such as taking only the first two conditions in Proposition 2 does not work: in more detail, requiring that there is a $C \in \mathcal{L}(\Sigma_3)$ s.t. for every conditional in $(B|A) \in \Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$:

1. $B \in \mathcal{L}(\Sigma_1) \cup \mathcal{L}(\Sigma_2)$,
2. $A \in \mathcal{L}(\Sigma_1) \cup \mathcal{L}(\Sigma_2)$, or $A \equiv C$.

In other words, these two conditions say that conditionals are either fully from the language based on either Σ_1 or Σ_2, or their antecedent is fully based on Σ_3, and there is only a single formula allowed to occur as such. However, this notion is not consistent with toleration. Consider $\Delta = \{(y|\top), (\neg y|a), (x|a)\}$. Then

$$\Delta = \{(y|\top), (\neg y|a)\} \bigcup_{\{y\},\{x\}} \{(x|a)\} \mid \{a\}$$

and $\{(x|a)\}$ tolerates itself (trivially), yet $\{(y|\top), (\neg y|a), (x|a)\}$ does not tolerate $(x|a)$. It is perhaps not surprising that such a purely syntactic condition is elusive.

[1] Or, equivalently, $\{H \to G \mid (G|H) \in \Delta^i\} \cup \{F \mid (F|C') \in \Delta^i, C' \equiv C\} \not\vdash \bot$.

Safe conditional splitting of a conditional belief base is consistent with toleration, in the sense that $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ implies that toleration of a conditional $(B|A)$ by Δ is equivalent to toleration of $(B|A)$ by the conditional sub-base Δ^i in which it occurs. This gives further evidence to the fact that safe conditional splitting adequately captures the notion of independence of sub-bases: toleration of a conditional is independent of a (conditionally) unrelated sub-base.

Proposition 3. Let a conditional belief base $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ be given. $\Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ implies (for any $i = 1, 2$) that Δ^i tolerates $(B|A) \in \Delta_i$ iff Δ tolerates $(B|A)$.

Proof. Suppose $\Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$. Suppose Δ_i tolerates $(B|A) \in \Delta_i$. Wlog let $i = 1$. This means there is an $\omega^1 \in \Omega(\Sigma_1)$ and $\omega^3 \in \Omega(\Sigma_3)$ s.t. $\omega^1\omega^3 \models A \wedge B$ and $\omega^1\omega^3 \models C \to D$ for every $(D|C) \in \Delta^1$. Since $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, $\omega^1\omega^2\omega^3 \models \bigwedge_{(F|E)\in\Delta_2} E \wedge \neg F$, i.e. $\omega^1\omega^2\omega^3 \models E \to F$ for every $(F|E) \in \Delta_j$. Thus, Δ tolerates $(B|A)$. The other direction is immediate. $\square$

We now move to the formulation of conditional syntax splitting, a property of inductive inference relations that expresses that the independencies between sub-bases of conditionals, as encoded in safe splitting, are respected by an inductive inference relation.

Conditional independence (**CInd**) and safe conditional relevance (**CRel**) are defined analogous to (**Ind**) and (**Rel**), but now assuming that a conditional belief base can be safely split and taking into account we have full information on the "conditional pivot" Σ_3:

Definition 8. An inductive inference operator $\mathbf{C}$ satisfies (**CInd**) if for any $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, and for any $A, B \in \mathcal{L}(\Sigma_i)$, $C \in \mathcal{L}(\Sigma_j)$ (for $i, j \in \{1, 2\}$, $j \neq i$) and a complete conjunction $D \in \mathcal{L}(\Sigma_3)$,

$$AD \mid\!\sim_{\Delta} B \text{ iff } ADC \mid\!\sim_{\Delta} B$$

Thus, an inductive inference operator satisfies conditional independence if, for any Δ that safely splits into Δ_1 and Δ_2 conditional on Σ_3, whenever we have all the necessary information about Σ_3, inferences from one sub-language are independent from formulas over the other sub-language.

Definition 9. An inductive inference operator $\mathbf{C}$ satisfies (**CRel**) if for any $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, and for any $A, B \in \mathcal{L}(\Sigma_i)$ (for $i \in \{1, 2\}$) and a complete conjunction $D \in \mathcal{L}(\Sigma_3)$,

$$AD \mid\!\sim_{\Delta} B \text{ iff } AD \mid\!\sim_{\Delta_i} B$$

Thus, **CRel** restricts the scope of inference by requiring that inferences in the sub-language $\Sigma_1 \cup \Sigma_3$ can be made on the basis of the conditionals on the basis of that sub-language.

Syntax splitting (**CSynSPlit**) combines the two properties (**CInd**) and (**CRel**):

Definition 10. An inductive inference operator $\mathbf{C}$ satisfies *conditional syntax splitting* (**CSynSPlit**) if it satisfies (**CInd**) and (**CInd**).

We now proceed with the study of conditional syntax splitting. We first analyse the properties of $\mathbf{CInd}$ and $\mathbf{CRel}$ for TPOs. We first notice that **CInd** and **CRel** for inductive inference operators for TPOs respectively for OCFs is equivalent to the following two properties (for any $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ and for $A, B \in \mathcal{L}(\Sigma_i)$, complete conjunction $D \in \mathcal{L}(\Sigma_3)$, $C \in \mathcal{L}(\Sigma_j)$, $i, j = 1, 2$ and $i \neq j$):

$\mathbf{CInd}^{tpo}$	$AD \preceq_\Delta BD$ iff $ACD \preceq_\Delta BCD$
$\mathbf{CRel}^{tpo}$	$AD \preceq_\Delta BD$ iff $AD \preceq_{\Delta^i} BD$
$\mathbf{CInd}^{ocf}$	$\kappa_\Delta(AD) \leqslant \kappa_\Delta(BD)$ iff $\kappa_\Delta(ACD) \leqslant \kappa_\Delta(BCD)$
$\mathbf{CRel}^{ocf}$	$\kappa_\Delta(AD) \leqslant \kappa_\Delta(BD)$ iff $\kappa_{\Delta^i}(AD) \leqslant \kappa_{\Delta^i}(BD)$

We now connect **CInd** to the notion of conditional independence of TPOs as known from belief revision. For this, we need the following notion taken from [18]:

Definition 11 ([18])**.** Let $\preceq$ be a total preorder on $\Omega(\Sigma)$, and let $\Sigma_1, \Sigma_2, \Sigma_3$ be three (disjoint) subsignatures of Σ. Then Σ_1 and Σ_2 are independent conditional on Σ_3, in symbols, $\Sigma_1 \perp\!\!\!\perp_{\preceq} \Sigma_2|\Sigma_3$, if for all $\omega_1^1, \omega_2^1 \in \Omega(\Sigma_1), \omega_1^2, \omega_2^2 \in \Omega(\Sigma_2)$, and $\omega^3 \in \Omega(\Sigma_3)$ holds that for all $i, j \in \{2, 3\}, i \neq j$,

$$\omega_1^i\omega_1^j\omega^3 \preceq \omega_2^i\omega_1^j\omega^3 \text{ iff } \omega_1^i\omega^3 \preceq \omega_2^i\omega^3. \quad (4)$$

Independence of two subsignatures Σ_i and Σ_j conditional on Σ_3 means that, in the context of fixed information about Σ_3, information about Σ_j is irrelevant for the ordering of worlds based on Σ_i: ω_1^j can be "cancelled out".

Proposition 4. An inductive inference operator for TPOs $\mathbf{C}^{tpo}$: $\Delta \mapsto \preceq_\Delta$ on $\mathcal{L}$ satisfies ($\mathbf{CInd}$) iff for any $\Delta = \Delta_1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$, it holds that $\Sigma_1 \perp\!\!\!\perp_{\preceq} \Sigma_2|\Sigma_3$.

Proof. For the $\Rightarrow$-direction, suppose that $\mathbf{C}^{tpo}$ satisfies ($\mathbf{CInd}$) and $\Delta = \Delta_1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$. Consider some $\omega_1^1, \omega_2^1 \in \Omega(\Sigma_1)$, $\omega^2 \in \Omega(\Sigma_2)$ and $\omega^3 \in \Omega(\Sigma_3)$ and suppose that $\omega_1^1\omega^3 \prec \omega_1^1\omega^3$. Thus, $\omega_1^1\omega^3 \vee \omega_2^1\omega^3 \mid\!\sim_{\preceq} \neg\omega_1^1$. Thus, with ($\mathbf{CInd}^{tpo}$), $\omega_1^1\omega^2\omega^3 \vee \omega_2^1\omega^2\omega^3 \mid\!\sim_{\preceq} \neg\omega_1^1$ and thus $\omega_1^1\omega^2\omega^3 \prec \omega_2^1\omega^2\omega^3$. The other direction of equation (4) is analogous.

For the $\Leftarrow$-direction, suppose $\Sigma_1 \perp\!\!\!\perp_{\preceq} \Sigma_2|\Sigma_3$ and suppose $\Delta = \Delta_1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$ and $AD \mid\!\sim_{\Delta} B$ for some $A, B \in \mathcal{L}(\Sigma_1)$ and some complete conjunction $D \in \mathcal{L}(\Sigma_3)$. Notice that since D is a complete conjunction, there is a unique world $\omega^3 \in \Omega(\Sigma_3)$ s.t. $\omega^3 \models D$. Then $ABD \prec A\overline{B}D$. Consider now some arbitrary $C \in \mathcal{L}(\Sigma_2)$. Notice that $A\overline{B}D \preceq A\overline{B}CD$. Take some $\omega_2^1\omega_2^2\omega^3 \in \min_{\preceq} \mathsf{Mod}(A\overline{B}CD)$. For any $\omega_1^1\omega_1^2\omega^3 \in \min_{\preceq} \mathsf{Mod}(ABD)$, $\omega_3^1\omega_1^2\omega^3 \in \mathsf{Mod}(A\overline{B}D)$ (as $B \in \mathcal{L}(\Sigma_1)$, and thus $\omega_1^1\omega_1^2\omega^3 \prec \omega_3^1\omega_1^2\omega^3$ (since $ABD \prec A\overline{B}D$). With independence, $\omega_1^1\omega^3 \prec \omega_2^1\omega^3$. Again with independence, $\omega_1^1\omega_2^2\omega^3 \prec \omega_2^1\omega_2^2\omega^3$. Since $\omega_2^2 \models C$, $\omega_1^1\omega_2^2\omega^3 \in \mathsf{Mod}(ABCD)$ and thus there is some $\omega_3^1\omega_3^2\omega^3 \in \min_{\preceq} \mathsf{Mod}(ABCD)$ with $\omega_3^1\omega_3^2\omega^3 \prec \omega_2^1\omega_2^2\omega^3$. Since $\omega_2^1\omega_2^2\omega^3 \in \min_{\preceq} \mathsf{Mod}(A\overline{B}CD)$, we have established that $ACD \mid\!\sim_{\preceq} B$. $\square$

Proposition 4 establishes a correspondence between the property **CInd** of inductive inference operators, and the notion of conditional independence for TPOs, as already known from belief revision.

Proposition 5. An inductive inference operator for TPOs $\mathbf{C}^{tpo}$: $\Delta \mapsto \preceq_\Delta$ on $\mathcal{L}$ satisfies ($\mathbf{CRel}$) iff for any $\Delta = \Delta_1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$, it holds that $\preceq_{\Delta^i} = \preceq_{\Delta|\Sigma_i}$.

Proof. For the $\Rightarrow$-direction, suppose that $\mathbf{C}^{tpo}$ satisfies ($\mathbf{CRel}^{tpo}$) and consider some $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$. Suppose $\omega_1^1\omega_1^3 \prec_{\Delta^i} \omega_2^1\omega_2^3$. Then $\omega_1^1\omega_1^3 \vee \omega_2^1\omega_2^3 \mid\!\sim_{\Delta_i} \neg\omega_2^1\omega_2^3$. With ($\mathbf{CRel}$), $\omega_1^1\omega_1^3 \vee \omega_2^1\omega_2^3 \mid\!\sim_{\Delta} \neg\omega_2^1\omega_2^3$ and thus $\omega_1^1\omega_1^3 \prec_\Delta \omega_2^1\omega_2^3$, i.e. $\omega_1^1\omega_1^3 \preceq_{\Delta|\Sigma_1} \omega_2^1\omega_2^3$. The other direction is analogous.

For the $\Leftarrow$-direction, suppose that $\Delta = \Delta^1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ and $\preceq_{\Delta_i} = \preceq_{\Delta|\Sigma_i}$. Suppose now $A \mid\!\sim_{\Delta} B$. Then $AB \prec_\Delta A\overline{B}$ and thus $AB \prec_{\Delta_i} A\overline{B}$, which implies $A \mid\!\sim_{\Delta^i} B$. $\square$

We now analyze conditional syntax splitting for inductive inference operators for OCFs. Thanks to the close relationship between rankings and probabilities, there is a straightforward adaptation of conditional independence for OCFs [19, Chapter 7].

Definition 12. Let $\Sigma_1 \dot\cup \Sigma_2 \dot\cup \Sigma_3 \subseteq \Sigma$ and let κ be an OCF. Σ_2, Σ_3 are *conditionally independent given* Σ_1 *with respect to* κ, in symbols $\Sigma_1 \perp\!\!\!\perp_\kappa \Sigma_2|\Sigma_3$, if for all $\omega^1 \in \Omega(\Sigma_1), \omega^2 \in \Omega(\Sigma_2)$, and $\omega^3 \in \Omega(\Sigma_3)$, $\kappa(\omega^1|\omega^1\omega^3) = \kappa(\omega^1|\omega^3)$ holds.

As for probabilities, conditional independence for OCFs expresses that information on Σ_3 is redundant for Σ_2 if full information on Σ_1 is available and used.

Proposition 6. An inductive inference operator for OCFs $\mathbf{C}^{ocf} : \Delta \mapsto \kappa_\Delta$ satisfies **CInd** iff for any $\Delta = \Delta_1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$ we have $\Sigma_1 \perp\!\!\!\perp_\kappa \Sigma_2|\Sigma_3$.

Proof. We first recall the following Lemma from [18]

Lemma 7. Let $\Sigma_1, \Sigma_2, \Sigma_3$ be disjoint subsignatures of Σ, let κ be an OCF. Then $\Sigma_1 \perp\!\!\!\perp_\kappa \Sigma_2|\Sigma_3$ iff for all $\omega^1 \in \Omega(\Sigma_1), \omega^2 \in \Omega(\Sigma_2)$, and $\omega^3 \in \Omega(\Sigma_3)$, we have $\kappa(\omega^1\omega^2\omega^3) = \kappa(\omega^1\omega^3) + \kappa(\omega^2\omega^3) - \kappa(\omega^3)$.

We now show that: (†) for any $A \in \mathcal{L}(\Sigma_i)$ and $\omega^3 \in \mathcal{L}(\Sigma_3)$, $\kappa(A\omega^3) = \min\{\kappa(\omega^1\omega^3) \mid \omega^i\omega^2\omega^3 \models A\}$. Wlog let $i = 1$ and $j = 2$. Observe that by definition, $\kappa(A\omega^3) = \min\{\kappa(\omega^1\omega^2\omega^3) \mid \omega^1\omega^2\omega^3 \models A\}$. Consider some $\omega^1\omega^2\omega^3 \in \mathsf{Mod}_{1,2,3}(A\omega^3)$ s.t. $\kappa(\omega^1\omega^2\omega^3) = \kappa(A\omega^3)$. With Lemma 7, $\kappa(\omega^1\omega^2\omega^3) = \kappa(\omega^1\omega^3) + \kappa(\omega^2\omega^3) - \kappa(\omega^3)$. Suppose now towards a contradiction that $\kappa(\omega^1\omega^3) > \kappa(A\omega^3)$, i.e. there is some $\omega^1_\star\omega^2_\star\omega^3$ s.t. $\omega^1_\star\omega^3 \models A\omega^3$ and $\kappa(\omega^1_\star\omega^3) + \kappa(\omega^2_\star\omega^3) < \kappa(\omega^1\omega^3) + \kappa(\omega^2\omega^3)$. Suppose first that $\kappa(\omega^2_\star\omega^3) > \kappa(\omega^2\omega^3)$. Then $\kappa(\omega^1_\star\omega^3) < \kappa(\omega^1\omega^3)$ and thus $\kappa(\omega^1_\star\omega^2\omega^3) < \kappa(\omega^1\omega^2\omega^3)$, contradiction. Suppose therefore that $\kappa(\omega^2_\star\omega^3) \leqslant \kappa(\omega^2\omega^3)$. Then we can derive that $\kappa(\omega^1_\star\omega^2\omega^3) < \kappa(\omega^1\omega^2\omega^3)$, contradiction.

From the †, it follows immediately that $\kappa(AB\omega^3) < \kappa(A\overline{B}\omega^3)$ iff $\kappa(ABC\omega^3) < \kappa(A\overline{B}C\omega^3)$ for any $A, B \in \mathcal{L}(\Sigma_i)$, $C \in \mathcal{L}(\Sigma_j)$ and $\omega^3 \in \mathcal{L}(\Sigma_3)$ (for $i, j = 1, 2$ and $i \neq j$). ☐

Proposition 8. An inductive inference operator for OCFs $\mathbf{C}^{ocf} : \Delta \mapsto \kappa_\Delta$ satisfies **CRel** iff for any $\Delta = \Delta_1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$, it holds that $\kappa_{\Delta_i} = \kappa_\Delta \mid_{\Sigma_1 \cup \Sigma_3}$.

Proof. Similar to the proof of Proposition 5. ☐

4. Lexicographic Inference Satisfies Conditional Syntax Splitting

In this section, we show that for any conditional belief base that safely splits conditionally, conditional syntax splitting is satisfied. We first need to show some intermediate results.

Fact 1. Where $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, $i \in \{1, 2\}$ and $(B|A) \in \Delta^i$, $Z_\Delta((B|A)) = Z_{\Delta^i}((B|A))$

Proof. Immediate from Proposition 3. ☐

The following Lemma shows that the components of vectors $\mathsf{lex}(\omega)$ can be simply combined by summation over disjoint sublanguages (taking into account double counting):

Lemma 9. Let a conditional belief base $\Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ with its corresponding Z-partition $(\Delta_0, \ldots, \Delta_n)$ be given. Then for every $0 \leqslant i \leqslant n$[2]:

$$\begin{aligned} V(\omega, \Delta_i) = &\; V(\omega^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^1_i) \\ = &\; V(\omega^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^2_i) \end{aligned}$$

Proof. Take some $0 \leqslant i \leqslant n$. Noice that $\Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, $(B|A) \in \Delta_i$ iff $(B|A) \in \Delta^1_i \setminus \Delta^2_i$ or $(B|A) \in \Delta^2_i \setminus \Delta^1_i$ or $(B|A) \in \Delta^1_i \cap \Delta^2_i$. Thus $V(\omega, \Delta_i)$

$$\begin{aligned} = &\; |\{(B|A) \in \Delta_i \mid \omega \models A \wedge \neg B\}| \\ = &\; |\{(B|A) \in \Delta^1_i \mid \omega^1\omega^3 \models A \wedge \neg B\}| \\ &+ |\{(B|A) \in \Delta^2_i \mid \omega^2\omega^3 \models A \wedge \neg B\}| \\ &- |\{(B|A) \in \Delta^1_i \cap \Delta^2_i \mid \omega^1\omega^2\omega^3 \models A \wedge \neg B\}|. \end{aligned}$$

Since $\Delta^1_i \cap \Delta^2_i = \Delta_i \cap (\mathcal{L}(\Sigma_3)|\mathcal{L}(\Sigma_3))$, and for any $(B|A) \in \Delta_i \cap (\mathcal{L}(\Sigma_3)|\mathcal{L}(\Sigma_3))$, $Z_\Delta((B|A)) = Z_{\Delta^1}((B|A)) = Z_{\Delta^2}((B|A))$ (with Fact 1) we have:

$$\begin{aligned} V(\omega, \Delta_i) = &\; V(\omega^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^1_i) \\ = &\; V(\omega^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^2_i). \end{aligned}$$

☐

Lemma 10. Where $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, $i \in \{1, 2\}$, $\phi \in \mathcal{L}(\Sigma_1)$ and $\omega^3 \in \Omega(\Sigma_3)$, $\omega \in \min_{\preceq^{\mathsf{lex}}_\Delta}(\mathsf{Mod}(\omega^3 \wedge \phi))$ iff $\omega^1 \in \min_{\preceq^{\mathsf{lex}}_{\Delta^1}}(\mathsf{Mod}_{1,3}(\omega^3 \wedge \phi))$ and $\omega^2 \in \min_{\preceq^{\mathsf{lex}}_{\Delta^2}} \mathsf{Mod}_{2,3}(\omega^3)$.

Proof. For the ⇒-direction, suppose $\omega \in \min_{\preceq^{\mathsf{lex}}_\Delta}(\mathsf{Mod}(\omega^3))$. Suppose now towards a contradiction that either (a) $\omega^2\omega^3 \notin \min_{\preceq^{\mathsf{lex}}_{\Delta^2}}(\mathsf{Mod}_{2,3}(\omega^3))$ or (b) $\omega^1 \notin \min_{\preceq^{\mathsf{lex}}_{\Delta^1}}(\mathsf{Mod}_{1,3}(\omega^3 \wedge \phi))$.

ad. (b) Suppose that $\omega^2\omega^3 \in \min_{\preceq^{\mathsf{lex}}_{\Delta^2}}(\mathsf{Mod}_{2,3}(\omega^3))$ (the case where also $\omega^2\omega^3 \notin \min_{\preceq^{\mathsf{lex}}_{\Delta^2}}(\mathsf{Mod}_{2,3}(\omega^3))$ is similar). $\omega^1 \notin \min_{\preceq^{\mathsf{lex}}_{\Delta^1}}(\mathsf{Mod}_{1,3}(\omega^3 \wedge \phi))$ implies that there is some $\omega'^1\omega^3 \in \mathsf{Mod}_{1,3}(\omega^3 \wedge \phi)$ s.t. $\omega'^1\omega^3 \prec^{\mathsf{lex}}_{\Delta^1} \omega^1\omega^3$, i.e. there is some $i \geqslant 0$ s.t. for every $j > i$, $V(\omega'^1\omega^3, \Delta^1_j) = V(\omega^1\omega^3, \Delta^1_j)$ and $V(\omega'^1\omega^3, \Delta^1_i) < V(\omega^1\omega^3, \Delta^1_i)$ (with Lemma 9) this implies:

$$\begin{aligned} &V(\omega'^1\omega^2\omega^3, \Delta_j) \\ = &\; V(\omega'^1\omega^3, \Delta^1_j) + V(\omega^2\omega^3, \Delta^2_j) - V(\omega^3, \Delta^1_j) \\ = &\; V(\omega^1\omega^2\omega^3, \Delta_j) \\ = &\; V(\omega^1\omega^3, \Delta^1_j) + V(\omega^2\omega^3, \Delta^2_j) - V(\omega^3, \Delta^1_j) \end{aligned}$$

and

$$\begin{aligned} &V(\omega'^1\omega^2\omega^3, \Delta_i) \\ = &\; V(\omega'^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^1_i) \\ < &\; V(\omega^1\omega^2\omega^3, \Delta_i) \\ = &\; V(\omega^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^1_i) \end{aligned}$$

which (since $\omega'^1\omega^2\omega^3 \in \mathsf{Mod}(\omega^3)$) contradicts $\omega \in \min_{\preceq^{\mathsf{lex}}_\Delta}(\omega^3)$.

ad. (a) Similar.

The ⇐-direction is similar.

☐

[2]Notice that it follows from Fact 1 that, given the Z-partition $(\Delta_1, \ldots, \Delta_n)$ of Δ and $\Sigma_i \subseteq \Sigma$, $\Delta^i_j = \Delta_j \cap (\mathcal{L}_i|\mathcal{L}_i)$ for any $0 \leqslant j \leqslant n$.

Proposition 11. C^{lex} satisfies **CInd**.

Proof. Suppose that $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$. We show that Σ_1 and Σ_2 are independent w.r.t. $\preceq^{\text{lex}}_\Delta$ conditional on Σ_3, i.e. for any $i \in \{1,2\}$, for any $\omega_1^i, \omega_2^i \in \Omega_i$, $\omega^3 \in \Omega_3$, $j \in \{1,2\}$, $j \neq i$, $\omega_1^i \omega^3 \preceq^{\text{lex}}_\Delta \omega_2^i \omega^3$ iff $\omega_1^i \omega^j \omega^3 \preceq^{\text{lex}}_\Delta \omega_2^i \omega^j \omega^3$ for all $\omega^j \in \Omega_j$. With Proposition 4 this is sufficient to show the proposition. For simplicity, we let $i = 1$ and $j = 2$, the other case follows by symmetry.

For the $\Rightarrow$-direction, suppose that $\omega_1^1 \omega^3 \preceq^{\text{lex}}_\Delta \omega_2^1 \omega^3$. We show that $\omega_1^1 \omega^2 \omega^3 \preceq^{\text{lex}}_\Delta \omega_1^1 \omega^2 \omega^3$ for all $\omega^2 \in \Omega_2$. We first make the following observation that follows in view of Lemma 10. For any $\omega^2_\star \in \min_{\preceq^{\text{lex}}_{\Delta^2}}(\mathsf{Mod}_{2,3}(\omega^3))$ and $i = 1, 2$:

$$\omega_i^1 \omega^3 \approx^{\text{lex}}_\Delta \omega_i^1 \omega^2_\star \omega^3.$$

Thus, for any $\omega^2 \in \min_{\preceq^{\text{lex}}_{\Delta^2}}(\mathsf{Mod}_{2,3}(\omega^3))$:

$$\omega_1^1 \omega^2 \omega^3 \preceq^{\text{lex}}_\Delta \omega_2^1 \omega^2 \omega^3.$$

This means (with Lemma 9) that there is some $i \geqslant 0$ s.t. $V(\omega_1^1\omega^3, \Delta^1_j) = V(\omega_2^1\omega^3, \Delta^1_j)$ for every $j > i$ and $V(\omega_1^1\omega^3, \Delta^1_i) \leqslant V(\omega_2^1\omega^3, \Delta^1_i)$. Thus, for any $\omega^2 \in \Omega_2$, it holds that (for any $j > i$):

$$\begin{aligned} & V(\omega_1^1\omega^3, \Delta^1_j) + V(\omega^2\omega^3, \Delta^2_j) - V(\omega^3, \Delta^1_j) \\ = \; & V(\omega_2^1\omega^3, \Delta^1_j) + V(\omega^2\omega^3, \Delta^2_j) - V(\omega_2^3, \Delta^1_j) \end{aligned}$$

and:

$$\begin{aligned} & V(\omega_1^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^1_j) \\ \leqslant \; & V(\omega_2^1\omega^3, \Delta^1_i) + V(\omega^2\omega^3, \Delta^2_i) - V(\omega^3, \Delta^1_j). \end{aligned}$$

With Lemma 9 we have that $V(\omega_k^1\omega^2\omega^3, \Delta_l) = V(\omega_k^1\omega^3, \Delta^1_l) + V(\omega^2\omega^3, \Delta^2_l) - V(\omega^3, \Delta^3_l)$ for $k = 1, 2$ and $l \leqslant n$ and thus $\omega_1^1\omega^2\omega^3 \preceq^{\text{lex}}_\Delta \omega_1^2\omega^2\omega^3$ (for any $\omega^2 \in \Omega_2$).

The $\Leftarrow$-direction is similar. □

Proposition 12. C^{lex} satisfies **Rel**.

Proof. This follows immediately from Lemma 10. Indeed, we have that (for $A, B \in \mathcal{L}_1$ and a complete conjunction $D \in \mathcal{L}_3$): $AD \mathrel{|\!\sim}^{\text{lex}}_\Delta B$ iff for all $\omega \in \min_{\preceq^{\text{lex}}_\Delta}(AD)$, $\omega^1\omega^3 \models B$. With Lemma 10, $\omega \in \min_{\preceq^{\text{lex}}_\Delta}(AD)$ iff $\omega^1 \in \min_{\preceq^{\text{lex}}_{\Delta^1}}(AD)$. Since $B \in \mathcal{L}_1$, this implies that if $\omega_1\omega^3 \in \min_{\preceq^{\text{lex}}_{\Delta^1}}(AD)$ then $\omega^1\omega^3 \models B$. □

From Proposition 11 and 12, we can immediately obtain the main result of this section:

Theorem 1. C^{lex} satisfies **SynSplit**.

5. The Drowning Effect as Conditional Independence

As mentioned in the introduction, the drowning effect, illustrated by Example 2, is intuitively related to syntax splitting. In more detail, the drowning effect is constituted by the fact that according to some inductive inference operators (e.g. system Z), exceptional subclasses (e.g. penguins) do not inherit any properties of the superclass (e.g. birds), even if these properties are unrelated to the reason for the subclass being exceptional (e.g. having beaks). To the best of our knowledge, discussion of the drowning effect in the literature has been restricted to informal discussions on the basis of examples such as the Tweety-example, but no generic formal description has been given.

In this paper, we have developed the necessary tools to talk about the drowning effect in a formally precise manner. Indeed, the first crucial notion is that of unrelatedness of propositions. This notion is formally captured by *safe splitting into subbases (Definition 7)*: given a belief base Δ, a proposition A is unrelated to a proposition C iff Δ can be safely split into subbases Δ_1, Δ_2 conditional on a sub-alphabet Σ_3, i.e. $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, and $A \in \mathcal{L}(\Sigma_2)$ and $C \in \mathcal{L}(\Sigma_1 \cup \Sigma_3)$. This means that the abstract situation of the drowning problem can be precisely described by conditional syntax splitting. We see that the drowning effect is nothing else than a violation of the postulate of conditional independence (**CInd**): if we know that a typical property B of AD-individuals ($AD \mathrel{|\!\sim}_\Delta B$) is unrelated to an exceptional subclass C of AD, then we can also derive that if something is ADC is typically B ($ADC \mathrel{|\!\sim}_\Delta B$). We illustrate this with the Tweety-example:

Example 8 (Example 2 ctd.)**.** We already established in Example 6 that $\Delta = \{(f|b), (b|p), (\neg f|p)\} \bigcup_{\{f,p\},\{e\}} \{(e|b)\} \mid \{b\}$. It is now not hard to see that any inductive inference operator **C** that satisfies (**DI**) and (**CInd**) avoids the drowning effect. In more detail, we have:

$$b \mathrel{|\!\sim}_\Delta e \qquad \text{by } \mathbf{DI} \tag{5}$$
$$b \wedge p \mathrel{|\!\sim}_\Delta e \qquad \text{by } \mathbf{CInd} \text{ and (5)} \tag{6}$$

For any inductive inference operator that additionally satisfies **Cut** (i.e. from $A \mathrel{|\!\sim} B$ and $A \wedge B \mathrel{|\!\sim} C$ derive $A \mathrel{|\!\sim} C$), a postulate that holds for any inductive inference operator based on TPOs [13], we obtain:

$$p \mathrel{|\!\sim}_\Delta b \qquad \text{by } \mathbf{DI} \tag{7}$$
$$p \mathrel{|\!\sim}_\Delta e \qquad \text{by } \mathbf{Cut}\text{, (6) and (7)} \tag{8}$$

Summarizing, we can express our findings as follows:

Proposition 13. Any inductive inference operator that satisfies (**CInd**) does not show the drowning problem (for $\Delta = \{(f|b), (b|p), (\neg f|p)\} \bigcup_{\{f,p\},\{e\}} \{(e|b)\} \mid \{b\}$).

We note that the drowning effect has not been defined in the literature in a general sense. Above, we provided, to the best of our knowledge, the first general definition of the drowning problem. As we defined the drowning effect as a violation of the postulate of conditional independence, it is trivial that any inductive inference operator that satisfies (**Cind**) does not show the drowning problem interpreted in this general, formal sense.

6. Lehmann's desirable closure properties

Lehmann 1995 remarks that in addition to the properties encoded by rational closure, other properties for inductive inference operators might be desirable. In particular, he lists four properties: *presumption of typicality*, *presumption of independence*, *priority of typicality* and *respect for specificity*. All of these properties are only explained informally and illustrated using examples in [1], and to the best of our knowledge, no attempts to formalize or generalize these notions has been made. In this section, we show how the desired behaviour for all but one of the examples given by [1] can be straightforwardly derived by assuming conditional inference relations satisfy (conditional) syntax splitting. We now describe the four properties from [1] and their relation with conditional syntax splitting:

Presumption of Typicality For a conditional belief base for which $(x|p) \in \Delta$ implies, for any inference relation $\mid\!\sim_\Delta$ that satisfies rational monotonicity (i.e. to derive from $A \mid\!\sim B$ and $A \not\mid\!\sim C$ that $AC \mid\!\sim B$), $p \wedge q \mid\!\sim_\Delta x$ or $p \mid\!\sim_\Delta \neg q$. The *presumption of typicality* obliges us, "in absence of a convincing reason to accept the latter" [1], to derive $p \wedge q \mid\!\sim x$. Lehmann does not elaborate on what constitutes "a convincing reason", but does state that for $\Delta_1 = \{(x|p)\}$, $p \wedge q \mid\!\sim_{\Delta_1} x$ should hold. It is clear that this behaviour follows from $(\mathbf{Ind})$:

Fact 2. Let an inductive inference operator satisfying $\mathbf{Ind}$ and $\Delta_1 = \{(x|p)\}$ be given. Then $p \wedge q \mid\!\sim_{\Delta_1} x$

Proof. Let $\Sigma_1 = \{p, x\}$ and $\Sigma_2 = \{q\}$. Then clearly $\Delta_1 = \Delta_1 \bigcup_{\Sigma_1,\Sigma_2} \emptyset$ and thus with $(\mathbf{Ind})$ we have: $p \mid\!\sim_{\Delta_1} x$ iff $p \wedge q \mid\!\sim_{\Delta_1} x$. Since $(x|p) \in \Delta_1$, with $(\mathbf{DI})$, $p \mid\!\sim_{\Delta_1} x$ and thus $p \wedge q \mid\!\sim_{\Delta_1} x$. □

Presumption of independence The presumption of independence states that "even if typicality is lost with respect to one consequent, we may still presume typicality with respect to another, unless there is a reason to the contrary" [1]. It is illustrated using the conditional belief base $\Delta_2 = \{(x|p), (\neg q|p)\}$. Then presumption of independence justifies us in deriving $p \wedge q \mid\!\sim_{\Delta_2} x$.

Fact 3. Let the inductive inference operator satisfying $\mathbf{CInd}$ and $\Delta_2 = \{(x|p), (\neg q|p)\}$ be given. Then $p \wedge q \mid\!\sim_{\Delta_2} x$.

Proof. It can be easily verified that $\Delta_2 = \{(x|p)\} \bigcup^{s}_{\{x\},\{q\}} \{(\neg q|p)\} \mid \{p\}$. With $\mathbf{DI}$, we have $p \mid\!\sim_{\Delta_2} x$. With $\mathbf{CInd}$ we have $p \wedge q \mid\!\sim_{\Delta_2} x$. □

Priority of typicality Priority of typicality gives, in situations where the presumption of typicality and the presumption of independence clash, priority to the former. It is illustrated using the conditional belief base $\Delta_3 = \{(x|p), (\neg x|p \wedge q)\}$. The presumption of typicality justifies us in deriving $p \wedge q \wedge r \mid\!\sim_{\Delta_3} \neg x$ (since no reason can be found for accepting $p \wedge q \mid\!\sim_{\Delta_3} \neg r$), whereas we can derive both $p \wedge q \wedge r \mid\!\sim \neg x$ and $p \wedge q \wedge r \mid\!\sim_{\Delta_3} x$ with the presumption of independence. Priority of typicality demands that priority is given to $p \wedge q \wedge r \mid\!\sim_{\Delta_3} \neg x$.

Fact 4. Let the inductive inference operator satisfying $\mathbf{Ind}$ and $\Delta_3 = \{(x|p), (\neg x|p \wedge q)\}$ be given. Then $p \wedge q \wedge r \mid\!\sim_{\Delta_3} \neg x$.

Proof. Let $\Sigma_1 = \{p, q, x\}$ and $\Sigma_2 = \{r\}$. Then clearly $\Delta_3 = \Delta_3 \bigcup_{\Sigma_1,\Sigma_2} \emptyset$ and thus with $(\mathbf{Ind})$ we have: $p \wedge q \mid\!\sim_{\Delta_3} \neg x$ iff $p \wedge q \wedge r \mid\!\sim_{\Delta_3} \neg x$. Since $(\neg x|p \wedge q) \in \Delta_3$, with $(\mathbf{DI})$, we have $p \wedge q \mid\!\sim_{\Delta_3} \neg x$. □

Lehmann 1995 uses a second example to illustrate priority of typicality:

Example 9 (Example 5 in [1]). Let $\Delta_4 = \{(x|p), (q|\top), (\neg x|q)\}$ and argues that here, $p \wedge q \mid\!\sim_{\Delta_4} x$ should hold. Here, Lehmann argues that $(q|\top)$ allows us to infer $p \mid\!\sim_{\Delta_4} q$ and thus p is "defeasibly more specific than q". Therefore, $p \wedge q \mid\!\sim_{\Delta_4} x$ should hold instead of $p \wedge q \mid\!\sim_{\Delta_4} \neg x$. However, we argue that this belief base and the resulting desirable inferences cannot be explained in terms of independence, as Δ_4 cannot be safely split into $\{(x|p)\}$ and $\{(q|\top), (\neg x|q)\}$.

Respect for specificity The final property, *respect for specificity*, gives guidelines on how to decide when the two presumptions clash: in that case the inference based on the assertion with a more specific antecedent should be used to guide the inferential process. These properties are illustrated in [1] using Δ_3 and Δ_4 as well.

This section thus shows that all four properties proposed in [1] are subsumed by (conditional) syntax splitting. Hence, any inference relation satisfying conditional syntax splitting also satisfies these properties.

7. Related Work

The phenomenon of syntax splitting has been observed as early as 1980 in [20] under the name of "system independence". The name *syntax splitting* was coined in [21] who studied it in the context of belief revision. Later, it was studied for other forms of belief revision in [22, 11], and for inductive inference operators in [8]. Our paper is a direct continuation of the work done in [7], where we have shown that lexicographic inference satisfies syntax splitting, and that the drowning effect is independent of syntax splitting. This work thus solves an important open question, namely whether generalization of syntax splitting as studied in [8, 7] can say something about the drowning effect.

Conditional independence for ranking functions has been studied in [12], for belief revision in [23], and for conditional belief revision in [18]. To the best of our knowledge, it has not been considered for inductive inference operators. We connect inductive inference operators with these works, as we show that the same conditions of conditional independence as studied in [12, 18] on the total preorders respectively OCFs underlying inductive inference operators (Definition 11) guarantee conditional syntax splitting.

In [16], the class of RC-extending inference relations is described. In future work, we plan to give a complete characterisation of the subclass of RC-extending inference relations satisfying syntax splitting.

8. Conclusion

The main contributions of this paper are the following: (1) we define the concept of conditional syntax splitting for inductive inference operators, thus bringing a notion of conditional independence between sub-signatures to the realm of inductive inference operators; (2) we show that lexicographic inference satisfies conditional syntax splitting; (3) we show how the drowning effect can be seen as a violation of conditional syntax splitting, and (4) we show how Lehman's desirable properties can be derived from (conditional) syntax splitting.

There are several main avenues for further work. Firstly, it will be interesting to investigate whether other inductive inference operators that satisfy (non-conditional) syntax splitting, such as c-representations ([14, 8] and system W [24, 25], also satisfy conditional syntax splitting. Secondly, we want to develop algorithms for deciding whether and how a conditional belief base can be safely split, and investigate their computational complexity. Thirdly, we plan to apply our results to applications of lexicographic inference to (probabilistic) description logics, and take advantage of them for the development of efficient implementations of lexicographic inference.

Acknowledgments

The work of Jesse Heyninck was partially supported by Fonds Wetenschappelijk Onderzoek – Vlaanderen (project G0B2221N).

References

[1] D. Lehmann, Another perspective on default reasoning, Annals of mathematics and artificial intelligence 15 (1995) 61–82.

[2] G. Casini, U. Straccia, Lexicographic closure for defeasible description logics, in: Proc. of Australasian Ontology Workshop, volume 969, 2012, pp. 28–39.

[3] T. Lukasiewicz, Expressive probabilistic description logics, Artificial Intelligence 172 (2008) 852–883.

[4] R. Booth, G. Casini, T. A. Meyer, I. J. Varzinczak, On the entailment problem for a logic of typicality, in: Twenty-Fourth International Joint Conference on Artificial Intelligence, 2015.

[5] M. Goldszmidt, J. Pearl, Qualitative probabilities for default reasoning, belief revision, and causal modeling, AI 84 (1996) 57–112.

[6] T. Eiter, T. Lukasiewicz, Default reasoning from conditional knowledge bases: Complexity and tractable cases, Artificial Intelligence 124 (2000) 169–241.

[7] J. Heyninck, G. Kern-Isberner, T. Meyer, Lexicographic entailment, syntax splitting and the drowning problem, in: Accepted for IJCAI 2022, 2022.

[8] G. Kern-Isberner, C. Beierle, G. Brewka, Syntax splitting= relevance+ independence: New postulates for nonmonotonic reasoning from conditional belief bases, in: Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning, volume 17, 2020, pp. 560–571.

[9] B. de Finetti, Theory of probability (2 vols.), 1974.

[10] D. Makinson, General theory of cumulative inference, in: International Workshop on Non-Monotonic Reasoning (NMR), Springer, 1988, pp. 1–18.

[11] G. Kern-Isberner, G. Brewka, Strong syntax splitting for iterated belief revision, in: C. Sierra (Ed.), Proceedings International Joint Conference on Artificial Intelligence, IJCAI 2017, ijcai.org, 2017, pp. 1131–1137.

[12] W. Spohn, Ordinal conditional functions: A dynamic theory of epistemic states, in: Causation in decision, belief change, and statistics, Springer, 1988, pp. 105–134.

[13] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artificial intelligence 44 (1990) 167–207.

[14] G. Kern-Isberner, Handling conditionals adequately in uncertain reasoning and belief revision, Journal of Applied Non-Classical Logics 12 (2002) 215–237.

[15] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artificial intelligence 55 (1992) 1–60.

[16] G. Casini, T. Meyer, I. Varzinczak, Taking defeasible entailment beyond rational closure, in: European Conference on Logics in Artificial Intelligence, Springer, 2019, pp. 182–197.

[17] M. Ritterskamp, G. Kern-Isberner, Preference-based default reasoning., in: FLAIRS Conference, 2008, pp. 672–677.

[18] G. Kern-Isberner, C. Beierle, J. Heyninck, Conditional independence for iterated belief revision, in: Accepted for IJCAI 2022, 2022.

[19] W. Spohn, The laws of belief: Ranking theory and its philosophical applications, Oxford University Press, 2012.

[20] J. Shore, R. Johnson, Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy, IEEE Transactions on Information Theory IT-26 (1980) 26–37.

[21] R. Parikh, Beliefs, belief revision, and splitting languages, Logic, language and computation 2 (1999) 266–268.

[22] T. Aravanis, P. Peppas, M.-A. Williams, Full characterization of parikh's relevance-sensitive axiom for belief revision, Journal of Artificial Intelligence Research 66 (2019) 765–792.

[23] M. J. Lynn, J. P. Delgrande, P. Peppas, Using conditional independence for belief revision, in: Proceedings AAAI-22, 2022.

[24] C. Komo, C. Beierle, Nonmonotonic reasoning from conditional knowledge bases with system w, Annals of Mathematics and Artificial Intelligence (2021) 1–38.

[25] J. Haldimann, C. Beierle, Inference with system w satisfies syntax splitting, arXiv preprint arXiv:2202.05511 (2022).