Conditional Syntax Splitting for Non-monotonic Inference Operators

Jesse Heyninck[1], Gabriele Kern-Isberner[2], Thomas Meyer[3,4], Jonas Philipp Haldimann[5], Christoph Beierle[5]

[1] Open Universiteit, the Netherlands
[2] Technische Universität Dortmund, Germany
[3]University of Cape Town and CAIR, South-Africa
[4] Centre for Artificial Intelligence Research (CAIR), South-Africa
[5] FernUniversität in Hagen, Germany
jesse.heyninck@ou.nl

Abstract

Syntax splitting is a property of inductive inference operators that ensures we can restrict our attention to parts of the conditional belief base that share atoms with a given query. To apply syntax splitting, a conditional belief base needs to consist of syntactically disjoint conditionals. This requirement is often too strong in practice, as conditionals might share atoms. 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 and system W satisfy conditional syntax splitting, and connect conditional syntax splitting to several known properties from the literature on non-monotonic reasoning, including the drowning effect.

1 Introduction

Inductive inference operators generate non-monotonic inference relations $\mid\!\sim_{\Delta}$ on the basis of a set of conditionals Δ of the form $(\psi|\phi)$, read as "if ϕ holds, then typically ψ holds". Examples of inductive inference operators include rational closure (also known as system Z) (Goldszmidt and Pearl 1996), lexicographic inference (Lehmann 1995), system W (Komo and Beierle 2022) and c-representations (Kern-Isberner 2002). For these systems, known complexity results point to a high computational complexity (Eiter and Lukasiewicz 2000). Syntax splitting for inference operators (Kern-Isberner, Beierle, and Brewka 2020) is a property requiring 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 from 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 decreasing the computational strain for many specific queries. However, this presupposes 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 (Kern-Isberner, Beierle, and Brewka 2020) 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 syntactically unrelated parts.*

An intuitively related problem that was surprisingly shown to be independent of syntax splitting in (Heyninck, Kern-Isberner, and Meyer 2022) is the so-called *drowning problem*. It consists of the fact that with some inductive inference relations, abnormal individuals do not inherit any properties. It is best illustrated using the Tweety-example:

Example 2 (The Drowning Problem, (Pearl 1990; Benferhat, Dubois, and Prade 1993))**.** *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 predicated on the fact that some inductive inference operators, such as system* Z*, do not infer that penguins typically have beaks* ($p \mid\!\sim^{Z}_{\Delta} b$)*, i.e. the abnormality of penguins w.r.t. flying* drowns *inferences about the beaks of penguins. It is well-known that lexicographic inference does not suffer from the drowning problem (Lehmann 1995).*

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, as long as we know we are talking about birds, having beaks has nothing to do with flying or having wings. However, (unconditional) syntax splitting does not allow capturing 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.

The definition of conditional independence is of major practical relevance. Over the last decades, conditional independence was shown to be a crucial concept supporting adequate modelling and efficient reasoning in probabilistic reasoning (Pearl 1988), as it is fundamental for network-based reasoning, arguably one of the most important factors in the rise of contemporary artificial intelligence. Even though many reasoning tasks on the basis of probabilistic information have a high worst-case complexity due to their semantic nature, network-based models allow for the efficient computation of many concrete instances of these reasoning tasks thanks to local reasoning techniques. The formulation of a concept of conditional independence for inductive inference operators is therefore an important step towards efficient implementations of non-monotonic reasoning.

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 a notion of conditional independence to inductive inference operators;
2. We show that lexicographic entailment and system W satisfy conditional syntax splitting;
3. We show that the drowning effect can be seen as a violation of conditional syntax splitting.

Outline of this paper: Section 2 states the necessary preliminaries, while Section 3 introduces the concept of conditional syntax splitting. We show that lexicographic inference and system W satisfy conditional syntax splitting in Section 4, and how avoidance of the drowning effect can be seen as a special case of conditional syntax splitting in Section 5. We discuss related work in Section 6 and conclude in Section 7.

2 Preliminaries

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

2.1 Propositional Logic

For a set Σ of atoms, here referred to as a *signature*, let $\mathcal{L}(\Sigma)$ 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 $\mathcal{L}(\Sigma)$ is a function $\omega : \Sigma \rightarrow \{\top, \bot\}$. Let $\Omega(\Sigma)$ denote the set of all interpretations for Σ. We simply write Ω if the set of atoms is implicitly given. An interpretation ω *satisfies* (or is a *model* of) an atom $a \in \Sigma$, 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 $a_1, \ldots, a_n \in \Sigma$ are those atoms that are assigned $\top$ by ω and $a_{n+1}, \ldots, a_m \in \Sigma$ 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 $X \subseteq \mathcal{L}(\Sigma)$ we also define $\omega \models X$ if and only if $\omega \models A$ for every $A \in X$. Define the set of models $\mathsf{Mod}(X) = \{\omega \in \Omega(\Sigma) \mid \omega \models X\}$ 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 (de Finetti 1937) who considered conditionals as *generalized indicator functions* for possible worlds resp. propositional interpretations ω:

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

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) or strict partial order (in short, SPO) $\preceq$ on possible worlds, representing relative plausibility, we define $A \preceq B$ iff for every $\omega' \in \min_{\preceq}(\mathsf{Mod}(B))$ there is an $\omega \in \min_{\preceq}(\mathsf{Mod}(A))$ such that $\omega \preceq \omega'$. This allows for expressing the validity of defeasible inferences via stating that $A \mathrel{|\!\sim}_{\preceq} B$ iff $(A \wedge B) \prec (A \wedge \neg B)$ (Makinson 1988). Analogously, we define the validity of defeasible inferences for strict partial orders (SPOs). As 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

$$\omega_1^\Theta \preceq_{|\Theta} \omega_2^\Theta \text{ iff } \omega_1^\Theta \preceq \omega_2^\Theta. \qquad (2)$$

Note that on the right hand side of the *iff* condition above $\omega_1^\Theta, \omega_2^\Theta$ are considered as propositions in the super-language $\mathcal{L}(\Omega)$; this marginalization of TPOs is a special case of the forgetful functor $Mod(\sigma)$ from Σ-models to Θ-models in (Beierle and Kern-Isberner 2012) where σ is the inclusion from Θ to Σ. Hence $\omega_1^\Theta \preceq \omega_2^\Theta$ is well defined (Kern-Isberner and Brewka 2017). Similarly, any inference relation $\mathrel{|\!\sim}$ on $\mathcal{L}(\Sigma)$ induces a *marginalized inference relation* $\mathrel{|\!\sim}_{|\Theta}$ on $\mathcal{L}(\Theta)$ by setting, for any $A, B \in \mathcal{L}(\Theta)$:

$$A \mathrel{|\!\sim}_{|\Theta} B \text{ iff } A \mathrel{|\!\sim} B \qquad (3)$$

An obvious implementation of total preorders are *ordinal conditional functions (OCFs)*, (also called *ranking functions*) $\kappa : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$. (Spohn

1988). 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 ((Kern-Isberner, Beierle, and Brewka 2020))**.** *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* $A \mathrel{|\!\sim}_{\Delta} B$*.*

As already indicated in Section 2.2, inference relations can be obtained based on SPOs, TPOs, and OCFs, respectively:

Definition 2. *An* SPO-based inductive inference operator *is a mapping* $\mathbf{C}^{spo}$ *that assigns to each conditional belief base* Δ *a strict partial order* $\prec_{\Delta}$ *on* Ω *s.t.* $A \mathrel{|\!\sim}_{\prec_{\Delta}} B$ *for every* $(B|A) \in \Delta$ *(i.e., s.t.* **DI** *is ensured). A* model-based inductive inference operator for total preorders $\mathbf{C}^{tpo}$ *is defined similarly, by using a TPO* $\preceq_{\Delta}$ *instead of an SPO* $\prec_{\Delta}$*.*

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 include System Z ((Goldszmidt and Pearl 1996), see Sec. 2.4) and c-representations ((Kern-Isberner 2002), whereas lexicographic inference ((Lehmann 1995), see Sec. 2.5) is an example of an inductive inference operator for TPOs and system W ((Komo and Beierle 2022), see Sec. 2.6) is an example of an inductive inference operator for SPOs.

To define the property of *syntax splitting* (Kern-Isberner, Beierle, and Brewka 2020), 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**), (Kern-Isberner, Beierle, and Brewka 2020))**.** *An inductive inference operator* $\mathbf{C}$ *satisfies (***Ind***) if for any* $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2$ *and for any* $A, B \in \mathcal{L}_i$, $C \in \mathcal{L}_j$ *(*$i, j \in \{1, 2\}$*,* $j \neq i$*),*

$$A \mathrel{|\!\sim}_{\Delta} B \text{ iff } AC \mathrel{|\!\sim}_{\Delta} B.$$

Definition 4 (Relevance (**Rel**), (Kern-Isberner, Beierle, and Brewka 2020))**.** *An inductive inference operator* $\mathbf{C}$ *satisfies (***Rel***) if for any* $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2$ *and for any* $A, B \in \mathcal{L}_i$ *(*$i \in \{1, 2\}$*),*

$$A \mathrel{|\!\sim}_{\Delta} B \text{ iff } A \mathrel{|\!\sim}_{\Delta^i} B.$$

Definition 5 (Syntax splitting (**SynSplit**), (Kern-Isberner, Beierle, and Brewka 2020))**.** *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 (Goldszmidt and Pearl 1996) 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 (or ordered partition) $OP(\Delta) = (\Delta_0, \ldots, \Delta_n)$ of Δ is defined as:

- $\Delta_0 = \{\delta \in \Delta \mid \Delta \text{ tolerates } \delta\}$;
- $OP(\Delta \setminus \Delta_0) = \Delta_1, \ldots, \Delta_n$.

For $\delta \in \Delta$ we define: $Z_{\Delta}(\delta) = i$ iff $\delta \in \Delta_i$ and $OP(\Delta) = (\Delta_0, \ldots, \Delta_n)$. 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* (Lehmann and Magidor 1992).

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. $\top \mathrel{|\!\sim}^Z_{\Delta} \neg p$ *and* $p \wedge f \not\mathrel{|\!\sim}^Z_{\Delta} b$*.*

2.5 Lexicographic Entailment

We recall lexicographic inference as introduced by (Lehmann 1995). 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 $OP(\Delta) = (\Delta_0, \ldots, \Delta_n)$, 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.

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* $p \wedge f \mid\!\sim^{\mathsf{lex}}_{\Delta} b$*.*

2.6 System W

System W is a recently introduced inductive inference operator (Komo and Beierle 2020, 2022) that takes into account the structural information of which conditionals are falsified.

Definition 6 (ξ^j, preferred structure $\prec^{\mathsf{w}}_{\Delta}$ on worlds (Komo and Beierle 2022))**.** *For a consistent belief base $\Delta = \{(B_i|A_i) \mid i \in \{1,\dots,n\}\}$ with $OP(\Delta) = (\Delta^0,\dots,\Delta^k)$ and for $j = 0,\dots,k$, the functions ξ^j are given by*

$$\xi^j(\omega) := \{(B_i|A_i) \in \Delta^j \mid \omega \models A_i\overline{B_i}\}.$$

The preferred structure on worlds *is given by the binary relation $\prec^{\mathsf{w}}_{\Delta} \subseteq \Omega \times \Omega$ defined by, for any $\omega, \omega' \in \Omega$,*

$$\begin{aligned} \omega \prec^{\mathsf{w}}_{\Delta} \omega' \ \ \textit{iff}\ \ & \textit{there exists an } m \in \{0,\dots,k\} \textit{ such that} \\ & \xi^i(\omega) = \xi^i(\omega') \quad \forall i \in \{m+1,\dots,k\} \textit{ and} \\ & \xi^m(\omega) \subset \xi^m(\omega'). \end{aligned}$$

I.e., $\omega \prec^{\mathsf{w}}_{\Delta} \omega'$ if and only if ω falsifies strictly fewer (in the set-theoretic sense) conditionals than ω' in the Δ^m with the biggest index m where the conditionals falsified by ω and ω' differ. Note that $\prec^{\mathsf{w}}_{\Delta}$ is a strict partial order (Komo and Beierle 2022, Lemma 3).

Definition 7 (system W, $\mid\!\sim^{\mathsf{w}}_{\Delta}$(Komo and Beierle 2022))**.** *Let Δ be a belief base and A, B be formulas. Then B is a* system W inference *from A (in the context of Δ), denoted $A \mid\!\sim^{\mathsf{w}}_{\Delta} B$, if for every $\omega' \in \mathsf{Mod}(A\overline{B})$ there is an $\omega \in \mathsf{Mod}(AB)$ such that $\omega \prec^{\mathsf{w}}_{\Delta} \omega'$.*

The relation $\prec^{\mathsf{w}}_{\Delta}$ is a strict partial order. Thus, using Definition 2, system W is an SPO-based inductive inference operator $C^{\mathsf{w}} : \Delta \mapsto \prec^{\mathsf{w}}_{\Delta}$. Note that system W satisfies **DI** because it captures system Z (see (Komo and Beierle 2020)) which in turn satisfies **DI**. In fact, system W strictly lies between system Z and lexicographic inference:

Proposition 1 ((Komo and Beierle 2020; Haldimann and Beierle 2022b))**.** *If A is consistent, then $A \mid\!\sim^{z}_{\Delta} B$ implies $A \mid\!\sim^{\mathsf{w}}_{\Delta} B$ and $A \mid\!\sim^{\mathsf{w}}_{\Delta} B$ implies $A \mid\!\sim^{lex}_{\Delta} B$, but not vice versa.*

Example 5 (Example 3 ctd.)**.** *The belief base Δ from Ex. 3 induces the $\prec^{\mathsf{w}}_{\Delta}$ below. We can entail $pb \mid\!\sim^{\mathsf{w}}_{\Delta} \overline{f}$ as the verifying world $pb\overline{f}$ is $\prec^{\mathsf{w}}_{\Delta}$-preferred to the only falsifying world $pb\overline{f}$, i.e., $pbf \prec^{\mathsf{w}}_{\Delta} pb\overline{f}$.*

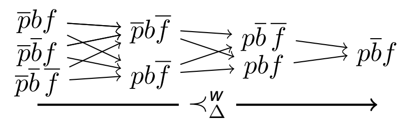

3 Conditional Syntax Splitting

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 8. *We say a conditional belief base Δ can be* split into subbases Δ_1,Δ_2 conditional on a sub-alphabet Σ_3, *if $\Delta_i = \Delta \cap (\mathcal{L}(\Sigma_i \cup \Sigma_3) \mid \mathcal{L}(\Sigma_i \cup \Sigma_3))$ 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:*

$$\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3 \tag{4}$$

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 $\Sigma_2 \cup \Sigma_3$. Thus, if (4) holds, we have $\Delta_1 \cup \Delta_2 = \Delta$ and $\Delta_1 \cap \Delta_2 = \Delta \cap (\mathcal{L}(\Sigma_3) \mid \mathcal{L}(\Sigma_3))$.

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 e.g. lexicographic inference. The underlying problem is that toleration might not be respected by conditional belief bases that conditionally split:

Example 6. *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. $\top \not\mid\!\sim^{\mathsf{lex}}_{\{(c|a\wedge b)\}} \neg(a \wedge b)$ whereas $\top \mid\!\sim^{\mathsf{lex}}_{\Delta} \neg(a \wedge b)$ (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, in 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 9. *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:*

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

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 (Proposition 2).

We argue here that safe splitting faithfully captures the independence of two conditional belief bases conditional on a sub-signature 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 $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$.

A simpler case of conditions 1 and 2 in Proposition 2 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 the conditional belief base from Example 2 has the form described in Proposition 2:

Example 7. *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 $\Delta = \{(f|b), (b|p), (\neg f|p)\} \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \{(e|b)\} \mid \Sigma_3$.

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

The bicycle example is also of this form:

Example 8. *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]*,*
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 not surprising that a purely syntactic condition is elusive.

Safe conditional splitting of a conditional belief base is consistent with toleration, in the sense that $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ 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 an 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. $\Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ implies (for any $i = 1, 2$) that Δ^i tolerates $(B|A) \in \Delta_i$ iff Δ tolerates $(B|A)$.*

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

Conditional independence (**CInd**) and 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 that we have full information on the "conditional pivot" Σ_3:

Definition 10. *An inductive inference operator $\mathbf{C}$ satisfies (**CInd**) if for any $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, 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 $D \in \mathcal{L}(\Sigma_3)$,*

$$AD \mathrel{|\!\sim}_{\Delta} B \text{ iff } ADC \mathrel{|\!\sim}_{\Delta} B$$

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.

[2] Even though this is a stronger version of condition 1 in Prop. 2, these two conditions are not sufficient to ensure safe splitting.

Definition 11. *An inductive inference operator* $\mathbf{C}$ *satisfies (*$\mathbf{CRel}$*) if for any* $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$*, and for any* $A, B \in \mathcal{L}(\Sigma_i)$ *(for* $i \in \{1,2\}$*) and a complete conjunction* $D \in \mathcal{L}(\Sigma_3)$*,*

$$AD \mathrel{|\!\sim}_{\Delta} B \text{ iff } AD \mathrel{|\!\sim}_{\Delta_i} B.$$

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 conditionals on the basis of that sub-language.

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

Definition 12. *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, SPOs, and OCFs respectively is equivalent to the following properties (for any $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ 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 $AD \preceq_{\Delta^i} BD$
$\mathbf{CInd}^{spo}$	$AD \prec_\Delta BD$ iff $ACD \prec_\Delta BCD$
$\mathbf{CRel}^{spo}$	$AD \prec_\Delta BD$ iff $AD \prec_{\Delta_i} BD$
$\mathbf{CInd}^{ocf}$	$\kappa_\Delta(AD) \leqslant \kappa_\Delta(BD)$ 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 (Kern-Isberner, Heyninck, and Beierle 2022):

Definition 13 ((Kern-Isberner, Heyninck, and Beierle 2022))**.** *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* $\omega^3 \in \Omega(\Sigma_3)$ *it holds that for all* $i, j \in \{1,2\}, i \neq j$*,*

$$\omega_1^i \omega_1^j \omega^3 \preceq \omega_2^i \omega_1^j \omega^3 \text{ iff } \omega_1^i \omega^3 \preceq \omega_2^i \omega^3. \tag{5}$$

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* $\Delta = \Delta_1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$*, it holds that* $\Sigma_1 \perp\!\!\!\perp_{\preceq_\Delta} \Sigma_2 | \Sigma_3$*.*

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* $\Delta = \Delta_1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$*, it holds that* $\preceq_{\Delta^i} = \preceq_{\Delta|\Sigma_i}$ *.*

Thanks to the close relationship between rankings and probabilities, there is a straightforward adaptation of conditional independence for OCFs (Spohn 2012, Chapter 7).

Definition 14. *Let* $\Sigma_1 \dot\cup \Sigma_2 \dot\cup \Sigma_3 \subseteq \Sigma$ *and let* κ *be an OCF.* Σ_1, Σ_2 *are* conditionally independent given Σ_3 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* $\omega^3 \in \Omega(\Sigma_3)$*,* $\kappa(\omega^1 | \omega^2 \omega^3) = \kappa(\omega^1 | \omega^3)$ *holds.*

As for probabilities, conditional independence for OCFs expresses that information on Σ_2 is redundant for Σ_1 if full information on Σ_3 is available and used. We can now characterize $\mathbf{CInd}$ and $\mathbf{CRel}$ for OCFs as follows:

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* $\Sigma_1 \perp\!\!\!\perp_{\kappa_\Delta} \Sigma_2 | \Sigma_3$*.*

Proposition 7. *An inductive inference operator for OCFs* $\mathbf{C}^{ocf} : \Delta \mapsto \kappa_\Delta$ *satisfies* **CRel** *iff for any* $\Delta = \Delta_1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$*, we have* $\kappa_{\Delta_i} = \kappa_\Delta |_{\Sigma_i \cup \Sigma_3}$ *for* $i \in \{1,2\}$*.*

4 Lexicographic Inference and System W Satisfy Conditional Syntax Splitting

We now show that two inductive inference operators, lexicographic inference and system W, which both satisfy syntax splitting (Heyninck, Kern-Isberner, and Meyer 2022; Haldimann and Beierle 2022a) also satisfy conditional syntax splitting, thus providing a proof of concept for the notion of conditional syntax splitting.

Theorem 1. C^{lex} *satisfies* **CRel** *and* **CInd** *(and thus* **CSynSplit***).*

A crucial result for this proof is the fact that the components of $\mathsf{lex}(\omega)$ can be simply combined by summation over disjoint sub-languages (taking into account double counting), i.e. for $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$ partitioned in $(\Delta_0, \ldots, \Delta_n)$, we have (for $1 \leqslant i \leqslant n$):

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

Theorem 2. *System W fulfils* **CRel** *and* **CInd** *(and thus* **CSynSplit***).*

The proof of Theorem 2 is based on the observations that for $\Delta = \Delta_1 \bigcup^{s}_{\Sigma_1,\Sigma_2} \Delta_2 \mid \Sigma_3$ the order $\omega \prec^{\mathsf{w}}_{\Delta} \omega'$ of two worlds coinciding on $\Sigma_2 \cup \Sigma_3$ depends only on Δ_1 and that the order $\omega \prec^{\mathsf{w}}_{\Delta_1} \omega'$ of worlds induced by Δ_1 does not change if we change the valuation over Σ_2 in the worlds. System Z does not satisfy **CSynSplit** as it does not satisfy **SynSplit** (Kern-Isberner, Beierle, and Brewka 2020).

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 (see e.g. (Pearl 1990; Benferhat, Dubois, and Prade 1993; Giordano and Gliozzi 2020; Benferhat et al. 1993), 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 9)*: 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 subalphabet 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 ($ADC \mathrel{|\!\!\sim}_{\Delta} B$). In other words, we can define the drowning effect in a general way, i.e. without recourse to a specific example, as follows:

Definition 15. *An inductive inference operator* $\mathbf{C}$ *shows the drowning problem if there is some* $\Delta = \Delta^1 \bigcup^{\mathsf{s}}_{\Sigma_1,\Sigma_2} \Delta^2 \mid \Sigma_3$, *some* $A, B \in \mathcal{L}(\Sigma_1)$, *some* $C \in \mathcal{L}(\Sigma_2)$ *and a complete conjunction* $D \in \mathcal{L}(\Sigma_3)$ *for which* $AD \mathrel{|\!\!\sim}_{\Delta} B$ *yet* $ADC \mathrel{|\!\!\not\sim}_{\Delta} B$.

In other words, an inductive inference operator shows the drowning problem if for some conditional belief base for which C is unrelated to AD, and for which AD typically implies B, ADC does *not* typically imply B.

Example 9 (Example 2 ctd.)**.** *We already saw in Example 7 that* $\Delta = \{(f|b), (b|p), (\neg f|p)\} \bigcup^{\mathsf{s}}_{\{f,p\},\{e\}} \{(e|b)\} \mid \{b\}$. *The drowning problem for this belief base consists in the fact that* $b \mathrel{|\!\!\sim}_{\Delta} e$ *yet* $b \wedge p \mathrel{|\!\!\not\sim}_{\Delta} e$. *It is not hard to see that any inductive inference operator* $\mathbf{C}$ *that satisfies (***DI***) and (***CInd***) avoids the drowning effect. In more detail, we have:*

$$b \mathrel{|\!\!\sim}_{\Delta} e \qquad \textit{by } \mathbf{DI} \tag{6}$$

$$b \wedge p \mathrel{|\!\!\sim}_{\Delta} e \qquad \textit{by } \mathbf{CInd} \textit{ and (5)} \tag{7}$$

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 SPOs, TPOs or OCFs (Kraus, Lehmann, and Magidor 1990), we obtain:*

$$p \mathrel{|\!\!\sim}_{\Delta} b \qquad \textit{by } \mathbf{DI} \tag{8}$$

$$p \mathrel{|\!\!\sim}_{\Delta} e \qquad \textit{by } \mathbf{Cut}\textit{, (7), and (8)} \tag{9}$$

Summarizing, we can express our findings as follows (notice that the proof of this proposition is trivial):

Proposition 8. *Any inductive inference operator that satisfies (***CInd***) does not show the drowning problem.*

As an immediate corollary, we obtain that lexicographic inference and system W avoid the drowning problem.

Notice that we do not claim here that the drowning problem is something that has to be avoided under any circumstances: in some situations, it might be useful to take a very cautious approach about assuming typicality. However, in many other situations, it is warranted to avoid the drowning problem. Our work gives a general description of what this means, and shows under which conditions on the inductive inference operator this is guaranteed.

6 Related Work

The phenomenon of syntax splitting has been observed as early as 1980 in (Shore and Johnson 1980) under the name of "system independence". *Syntax splitting* was coined in (Parikh 1999) who studied it in the context of belief revision. Later, it was studied for other forms of belief revision in (Aravanis, Peppas, and Williams 2019; Kern-Isberner and Brewka 2017), and for inductive inference operators in (Kern-Isberner, Beierle, and Brewka 2020). In (Heyninck, Kern-Isberner, and Meyer 2022) it was shown that lexicographic inference satisfies syntax splitting, and that the drowning effect is independent of (non-conditional) syntax splitting. We pick up where (Heyninck, Kern-Isberner, and Meyer 2022) stopped, showing the connection between the drowning effect and conditional syntax splitting.

Conditional independence for OCFs has been studied in (Spohn 1988), for belief revision in (Lynn, Delgrande, and Peppas 2022), and for conditional belief revision in (Kern-Isberner, Heyninck, and Beierle 2022). 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 (Spohn 1988; Kern-Isberner, Heyninck, and Beierle 2022) on the TPOs respectively OCFs underlying inductive inference operators (Definition 13) guarantee conditional syntax splitting.

7 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 and system W satisfy conditional syntax splitting; and (3) we show how the drowning effect can be seen as a violation of conditional syntax splitting.

Avenues for further work include investigating whether other inference operators, such as c-representations (Kern-Isberner 2002) and RC-extending inference relations (as described in (Casini, Meyer, and Varzinczak 2019)) satisfy conditional syntax splitting, developing algorithms for deciding whether and how a conditional belief base can be safely split and analyzing their complexity. Finally, we plan to compare the semantic notion of syntax splitting with related notions studied in logic programming, such as *splitting* (Lifschitz and Turner 1994) and modularity (Janhunen et al. 2009; Lierler and Truszczynski 2013).

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), grant BE 1700/10-1 awarded to Christoph Beierle as part of the priority program "Intentional Forgetting in Organizations" (SPP 1921). Jonas Haldimann was supported by this grant.

References

Aravanis, T.; Peppas, P.; and Williams, M.-A. 2019. Full Characterization of Parikh's Relevance-Sensitive Axiom for Belief Revision. *Journal of Artificial Intelligence Research*, 66: 765–792.

Beierle, C.; and Kern-Isberner, G. 2012. Semantical Investigations into Nonmonotonic and Probabilistic Logics. *Annals of Mathematics and Artificial Intelligence*, 65(2-3): 123–158.

Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax-based entailment. In *IJCAI*, volume 93, 640–645.

Benferhat, S.; Dubois, D.; and Prade, H. 1993. Possibilistic logic: From nonmonotonicity to logic programming. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, 17–24. Springer.

Casini, G.; Meyer, T.; and Varzinczak, I. 2019. Taking defeasible entailment beyond rational closure. In *European Conference on Logics in Artificial Intelligence*, 182–197. Springer.

de Finetti, B. 1937. La prévision, ses lois logiques et ses sources subjectives. English translation in *Studies in Subjective Probability*, ed. H. Kyburg and H.E. Smokler, 1974, 93–158. New York: Wiley & Sons.

Eiter, T.; and Lukasiewicz, T. 2000. Default reasoning from conditional knowledge bases: Complexity and tractable cases. *Artificial Intelligence*, 124(2): 169–241.

Giordano, L.; and Gliozzi, V. 2020. Reasoning about exceptions in ontologies: from the lexicographic closure to the skeptical closure. *Fundamenta Informaticae*, 176(3-4): 235–269.

Goldszmidt, M.; and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *AI*, 84(1-2): 57–112.

Haldimann, J.; and Beierle, C. 2022a. Inference with System W Satisfies Syntax Splitting. In Kern-Isberner, G.; Lakemeyer, G.; and Meyer, T., eds., *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel. July 31 - August 5, 2022*, 405–409.

Haldimann, J.; and Beierle, C. 2022b. Properties of System W and Its Relationships to Other Inductive Inference Operators. In Varzinczak, I., ed., *Foundations of Information and Knowledge Systems - 12th International Symposium, FoIKS 2022, Helsinki, Finland, June 20-23, 2022, Proceedings*, volume 13388 of *LNCS*, 206–225. Springer.

Heyninck, J.; Kern-Isberner, G.; and Meyer, T. A. 2022. Lexicographic Entailment, Syntax Splitting and the Drowning Problem. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, 2662–2668. ijcai.org.

Janhunen, T.; Oikarinen, E.; Tompits, H.; and Woltran, S. 2009. Modularity aspects of disjunctive stable models. *Journal of Artificial Intelligence Research*, 35: 813–857.

Kern-Isberner, G. 2002. Handling conditionals adequately in uncertain reasoning and belief revision. *Journal of Applied Non-Classical Logics*, 12(2): 215–237.

Kern-Isberner, G.; Beierle, C.; and Brewka, G. 2020. 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, 560–571.

Kern-Isberner, G.; and Brewka, G. 2017. Strong Syntax Splitting for Iterated Belief Revision. In Sierra, C., ed., *Proceedings International Joint Conference on Artificial Intelligence, IJCAI 2017*, 1131–1137. ijcai.org.

Kern-Isberner, G.; Heyninck, J.; and Beierle, C. 2022. Conditional Independence for Iterated Belief Revision. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, 2690–2696. ijcai.org.

Komo, C.; and Beierle, C. 2020. Nonmonotonic Inferences with Qualitative Conditionals Based on Preferred Structures on Worlds. In Schmid, U.; Klügl, F.; and Wolter, D., eds., *KI 2020: Advances in Artificial Intelligence - 43rd German Conference on AI, Bamberg, Germany, September 21-25, 2020, Proceedings*, volume 12325 of *LNCS*, 102–115. Springer.

Komo, C.; and Beierle, C. 2022. Nonmonotonic reasoning from conditional knowledge bases with system W. *Ann. Math. Artif. Intell.*, 90(1): 107–144.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence*, 44(1-2): 167–207.

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

Lehmann, D.; and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial intelligence*, 55(1): 1–60.

Lierler, Y.; and Truszczynski, M. 2013. Modular answer set solving. In *Proceedings of the 17th AAAI Conference on Late-Breaking Developments in the Field of Artificial Intelligence*, 68–70.

Lifschitz, V.; and Turner, H. 1994. Splitting a logic program. In *ICLP*, volume 94, 23–37.

Lynn, M. J.; Delgrande, J. P.; and Peppas, P. 2022. Using Conditional Independence for Belief Revision. In *Proceedings AAAI-22*.

Makinson, D. 1988. General theory of cumulative inference. In *International Workshop on Non-Monotonic Reasoning (NMR)*, 1–18. Springer.

Parikh, R. 1999. Beliefs, belief revision, and splitting languages. *Logic, language and computation*, 2(96): 266–268.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan kaufmann.

Pearl, J. 1990. System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of the 3rd conference on Theoretical aspects of reasoning about knowledge*, 121–135.

Shore, J.; and Johnson, R. 1980. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, IT-26: 26–37.

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

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