Lexicographic Entailment, Syntax Splitting and the Drowning Problem

Jesse Heyninck[1,2,3,4] , **Gabriele Kern-Isberner**[4] , **Thomas Meyer**[2]
[1] Open Universiteit Heerlen, the Netherlands
[2] University of Cape Town and CAIR, South-Africa
[3] Vrije Universiteit Brussel, Belgium
[4] Technische Universität Dortmund, Germany
{jesse.heyninck, gabriele.kern-isberner}@tu-dortmund.de, tmeyer@cair.org.za

Abstract

Lexicographic inference is a well-known and popular approach to reasoning with non-monotonic conditionals. It is a logic of very high-quality, as it extends rational closure and avoids the so-called drowning problem. It seems, however, this high quality comes at a cost, as reasoning on the basis of lexicographic inference is of high computational complexity. In this paper, we show that lexicographic inference satisfies syntax splitting, which means that we can restrict our attention to parts of the belief base that share atoms with a given query, thus seriously restricting the computational costs for many concrete queries. Furthermore, we make some observations on the relationship between c-representations and lexicographic inference, and reflect on the relation between syntax splitting and the drowning problem.

1 Introduction

Lexicographic inference [Lehmann, 1995] is a well-known and popular approach to reasoning with non-monotonic conditionals, which has been applied in (probabilistic) description logics [Casini and Straccia, 2012; Giugno and Lukasiewicz, 2002] as well as richer preferential languages [Booth *et al.*, 2015]. It is seen as a logic of very high-quality, as it extends rational closure and avoids the so-called *drowning problem*. It seems, however, this high quality comes 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 [Eiter and Lukasiewicz, 2000]. In this paper, we show that lexicographic inference satisfies *syntax splitting* [Kern-Isberner *et al.*, 2020], which means that we can restrict our attention to parts of the belief base that share atoms with a given query, thus seriously restricting the computational strain for many concrete queries. To the best of our knowledge, this is the first result that shows that lexicographic inference satisfies this property, and thus we show that there exists an inductive inference relation that both extends rational closure and satisfies syntax splitting. In more detail, the contributions of this paper are the following:

1. We show that lexicographic inference satisfies syntax splitting. This shows that lexicographic inference is an inference relation that is of the highest quality, and at least in many cases, its computational complexity can be circumvented.
2. We show that lexicographic inference is closely related but not identical to c-representations [Kern-Isberner, 2002], the only other inference relation that was shown to satisfy syntax splitting.
3. We show that lexicographic inference is not the only extension of rational closure that satisfies syntax splitting, and that syntax splitting is independent of the drowning problem.

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), lexicographic inference (Section 2.5) and c-representations (Section 2.6). In Section 3 we show that lexicographic inference satisfies syntax splitting. In Section 4 we compare lexicographic inference with the only other inductive inference relation known to satisfy syntax splitting: c-representations. In Section 5 we make some further observations on inference relations extending rational closure, syntax splitting and the drowning effect. In Section 6 we discuss related work, and we conclude in Section 7.

2 Preliminaries

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

2.1 Propositional Logic

For a set Σ of atoms 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 $X_1 \vdash X_2$, 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 Σ^θ, i.e. ω^θ is the interpretation over Σ^θ that agrees with ω on all atoms in Σ^θ. $\Omega(\Sigma_i)$ will also be denoted by Ω_i for any $i \in \mathbb{N}$. 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, 1974] 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 $\preceq$ 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 $(A \wedge B) \prec (A \wedge \neg B)$ [Makinson, 1988]. 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 eqn. (2) 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 [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

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

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

An obvious implementation of total preorders are *ordinal conditional functions (OCFs)*, (also called *ranking functions*) $\kappa : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ [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

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 *et al.*, 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 $\mathrel{|\!\sim}_\Delta \subseteq \mathcal{L} \times \mathcal{L}$ 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 on Δ, then $(B|A) \in \Delta$ implies $A \mathrel{|\!\sim}_\Delta B$.

Examples of inductive inference operators include system P [Kraus *et al.*, 1990] and the inference relations below.

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 are System Z ([Goldszmidt and Pearl, 1996], see Sec. 2.4) and c-representations ([Kern-Isberner, 2002], see Sec. 2.6), whereas lexicographic inference ([Lehmann, 1995], see Sec. 2.5) is an example of an inductive inference operator for TPOs.

To define the property of *syntax splitting* [Kern-Isberner *et al.*, 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 *et al.*, 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 *et al.*, 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 *et al.*, 2020]). An inductive inference operator **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 $(\Delta_0, \ldots, \Delta_n)$ 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* [Lehmann and Magidor, 1992]. An inference relation $\mathrel{|\!\sim}_\Delta$ based on Δ s.t. $A \mathrel{|\!\sim}^Z_\Delta B$ implies $A \mathrel{|\!\sim}_\Delta B$ is called *RC-extending* [Casini *et al.*, 2019]. An *RC-extending* inference relation has also been called a *refinement of System Z* [Ritterskamp and Kern-Isberner, 2008]. We call an inductive inference operator **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 1. Let $\Delta = \{(f|b), (b|p), (\neg f|p)\}$, which expresses that birds (b) typically fly (f), penguins ((p)) are typically birds, and penguins typically don't fly. 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 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 [Lehmann, 1995], lexicographic inference was shown to be RC-extending (for finite conditional belief bases):

Proposition 1 ([Lehmann, 1995, Theorem 3]). For any $A \in \mathcal{L}$ s.t. $\kappa^Z_\Delta(A)$ is finite, then $A \mathrel{|\!\sim}^Z_\Delta B$ implies $A \mathrel{|\!\sim}^{\mathsf{lex}}_\Delta B$.

Example 2 (Example 1 ctd.). For the Tweety belief base Δ as in Example 1 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. $\top \mathrel{|\!\sim}^{\mathsf{lex}}_\Delta \neg p$ (since $\mathsf{lex}(\top \wedge \neg p) = (0,0) \prec^{\mathsf{lex}} \mathsf{lex}(\top \wedge p) = (1,0)$) and $p \wedge f \mathrel{|\!\sim}^{\mathsf{lex}}_\Delta b$.

2.6 C-Representations

Here we present c-representations as presented by [Kern-Isberner, 2002]. Given a set of conditionals $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$, an OCF κ is *indifferent with respect to* Δ iff there are $\eta_0, \eta_i^- \in \mathbb{Q}$ for $1 \leqslant i \leqslant n$ s.t. for every $\omega \in \Omega$:

$$\kappa(\omega) = \eta_0 + \sum_{\omega \models A_i \overline{B_i}}^{1 \leqslant i \leqslant n} \eta_i \tag{4}$$

We can then ensure that $\kappa \models \Delta$ holds by making sure the following inequality holds for every $1 \leqslant i \leqslant n$:[1]

$$\eta_i > \min_{\omega \models A_i B_i} \left(\sum_{\omega \models A_j \overline{B_j}}^{j \neq i} \eta_j \right) - \min_{\omega \models A_i \overline{B_i}} \left(\sum_{\omega \models A_j \overline{B_j}}^{j \neq i} \eta_j \right) \tag{5}$$

A *c-representation* of a set of conditionals Δ is an OCF κ of the form (4) whose parameters satisfy equation (5). It can be observed that for any c-representation κ of Δ, $\kappa \models \Delta$.

Example 3 (Example 1 ctd.). For the Tweety belief base Δ as in Example 1, with $r_1 = (f|b), r_2 = (b|p), r_3 = (\neg f|p)$, we consider the (minimal) c-representation characterised by $\eta_1 = 1$ and $\eta_2 = \eta_3 = 2$. This results in the following OCF κ (on the basis of equation 4):

ω	$\kappa(\omega)$	ω	$\kappa(\omega)$	ω	$\kappa(\omega)$	ω	$\kappa(\omega)$
pbf	2	$pb\overline{f}$	1	$p\overline{b}f$	4	$p\overline{b}\,\overline{f}$	2
$\overline{p}bf$	0	$\overline{p}b\overline{f}$	1	$\overline{p}\overline{b}f$	0	$\overline{p}\overline{b}\,\overline{f}$	0

[1] To save space, we sometimes denote $A \wedge \neg B$ by $A\overline{B}$.

As an example of an inference, observe that e.g. $\top \mid\!\sim_\kappa \neg p$ and $p \wedge f \mid\!\sim_\kappa b$.

3 Lexicographic Entailment Satisfies Syntax Splitting

In this short but central section, we state one of the main results of the paper, namely that lexicographic entailment satisfies syntax splitting. Syntax splitting is shown by showing that **Ind** and **Rel** are satisfied. **Rel** is quite straightforward. In fact it also holds for e.g. system Z [Kern-Isberner *et al.*, 2020], and is proved in a way that is similar to the proof for system Z in [Kern-Isberner *et al.*, 2020]. For **Ind**, things are a bit more involved. Due to spatial restrictions, the proof could not be included (but is given, like the proofs of all results in this paper, in the appendix), but we illustrate the main ideas using an example.

Example 4 (Based on [Kern-Isberner *et al.*, 2020, Ex. 9]). Consider $\Delta = \{(a|\top), (b|\top)\}$. Then $\Delta = \{(a|\top)\} \bigcup_{\{a\},\{b\}} \{(b|\top)\}$ and $\Delta = \Delta_0$. We have the following κ^Z_Δ respectively lex-values over Ω:

ω	$\kappa^Z_\Delta(\omega)$	$\mathsf{lex}(\omega)$	ω	$\kappa^Z_\Delta(\omega)$	$\mathsf{lex}(\omega)$
ab	0	(0)	$a\overline{b}$	1	(1)
$\overline{a}b$	1	(1)	$\overline{a}\overline{b}$	1	(2)

We see that $\top \mid\!\sim^Z_\Delta a$ (since $\kappa^Z_\Delta(a) < \kappa^Z_\Delta(\overline{a})$) yet $\overline{b} \not\mid\!\sim^Z_\Delta a$ (since $\kappa^Z_\Delta(a\overline{b}) = \kappa^Z_\Delta(\overline{a}\overline{b})$), i.e. inferences in System Z over the signature $\{a\}$ are not independent from the signature $\{b\}$ even though $\{a\}, \{b\}$ splits Δ. The reason is that $a\overline{b}$ and $\overline{a}\overline{b}$ have the same κ^Z_Δ-value, and thus the κ^Z_Δ-value does not preserve the original order between a and $\overline{a}$ when looking at $\overline{b}$-worlds. This is not the case for lexicographic inference since $\mathsf{lex}(a\overline{b}) \prec^{\mathsf{lex}} \mathsf{lex}(\overline{a}\overline{b})$.

More generally, the reason for lexicographic entailment satisfying **Ind** is that, since lexicographic vectors count the number of conditionals violated by a world, the relative order of marginalized world is independent of other subsignatures, i.e. (for $\Delta = \Delta^1 \bigcup_{\Sigma_1,\Sigma_2} \Delta^2$, $\omega^i_1, \omega^i_2 \in \Omega_i$, $\omega^j \in \Omega_j$, $i, j = 1, 2$ and $i \neq j$):

$$\omega^i_1 \preceq^{\mathsf{lex}}_\Delta \omega^i_2 \text{ iff } \omega^i_1\omega^j \preceq^{\mathsf{lex}}_\Delta \omega^i_2\omega^j$$

Altogether, the following central theorem of the paper can be shown:

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

It is important to notice that the satisfaction of syntax splitting by lexicographic entailment does not improve the worst-case complexity of lexicographic inference. However, syntax splitting can be used to develop implementations that are more efficient than current implementations for many concrete queries, as it ensures that we can restrict attention to the part of the conditional belief base relevant to the query.

4 Lexicographic Entailment in C-Representations

In previous work, only c-representations were shown to be an inductive inference method that satisfies syntax splitting (see [Kern-Isberner *et al.*, 2020]). Given the results shown in the previous section, we can thus say that lexicographic inference and c-representations are two inductive inference methods of very high quality. In order to obtain a precise view on these two high-quality inductive inference relations, we investigate the exact relationship between c-representations and lexicographic inference in this section. The comparison between lexicographic inference and c-representations was started in [Strotkamp, 2020], where an algorithm (Algorithm 1) originally proposed in [Bourne and Parsons, 1999] was adapted to an algorithm that generates, given a conditional belief base Δ, an OCF κ that (1) is a c-representation of Δ and (2) preserves lexicographic inference, i.e. for any ω, ω': $\omega \prec^{\mathsf{lex}}_\Delta \omega'$ iff $\kappa(\omega) < \kappa(\omega')$. In more detail, given a belief base Δ (Z-partitioned into $(\Delta_0, \ldots, \Delta_n)$), we use Algorithm 1 to calculate η_i for any rule $r_i \in \Delta$. We can then obtain the rank of any $\omega \in \Omega$ by setting:

Algorithm 1 Algorithm for generating a c-representation equivalent to lexicographic entailment

---

Input: a belief base $\Delta = \{r_i \mid i \in I\}$ Z-partitioned in $(\Delta_0, \ldots, \Delta_n)$.
Output: A set of κ^-_i for $i \in I$.

$\eta^{\mathtt{lex2c}}_{\Delta_0} = 1$
$\eta_i = \eta^{\mathtt{lex2c}}_{\Delta_0}$ for every $r_i \in \Delta_0$
for $1 \leqslant k \leqslant n$ **do**
  $\eta^{\mathtt{lex2c}}_{\Delta_k} = (|\Delta_{k-1}| + 1) \times \eta^{\mathtt{lex2c}}_{\Delta_{k-1}}$
  $\eta_i = \kappa^{\mathtt{lex2c}}_{\Delta_k}$ for all $r_i \in \Delta_k$;
end for

$$\kappa^c_{\mathsf{lex}}(\omega) := \sum^{1 \leqslant i \leqslant n}_{\omega \models A_i\overline{B_i}} \eta_i \tag{6}$$

[Strotkamp, 2020] shows that this algorithm results in a c-representation that captures lexicographic inference:

Proposition 2. (1) κ^c_{lex} (as defined in Equation 6) forms a c-representation of Δ. (2) For any $\omega, \omega' \in \Omega$, $\kappa^c_{\mathsf{lex}}(\omega) \leqslant \kappa^c_{\mathsf{lex}}(\omega')$ iff $\omega \preceq^{\mathsf{lex}} \omega'$.

However, this does not imply that $\preceq^{\mathsf{lex}}$ can be unambiguously equated with (a subclass of) c-representations. Indeed, the algorithm from [Strotkamp, 2020] might give rise to non-convex OCFs, i.e. OFCs with "empty" layers. One might wonder whether it is also possible to construct a convex c-representation equivalent to lexicographic inference. This question is answered negatively:

Proposition 3. There exist conditional belief bases Δ s.t. there is no c-representation $\kappa : \Omega \to \mathbb{N}$ for Δ s.t. (1) $\kappa^{-1}(i) \neq \emptyset$ implies $\kappa^{-1}(i-1) \neq \emptyset$ for any $i \in \mathbb{N}$, (2) $\eta_i \in \mathbb{Q}^+$ for every $r_i \in \Delta$, and (3) $\kappa(\omega) \leqslant \kappa(\omega')$ iff $\omega \preceq^{\mathsf{lex}} \omega'$ for any $\omega, \omega' \in \Omega$.

5 Syntax Splitting and the Drowning Problem

The main result of this paper (Theorem 1) is that lexicographic inference satisfies syntax splitting, which means that it is the first RC-extending inference relation shown to satisfy syntax splitting (since c-representations are incomparable with system Z [Beierle *et al.*, 2018] and system Z itself does not satisfy syntax splitting (see Example 4). To the best of our knowledge, lexicographic inference and c-representations are also the only two inductive inference operators known to avoid the drowning effect.

We start by recalling the drowning problem:

Example 5 (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 together with the additional conditional "birds typically have beaks" ($(e|b)$). 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 blocks or *drowns* inferences about penguins with beaks. It can be verified that lexicographic inference and c-representations do not suffer from the drowning problem.

As lexicographic inference and c-representations are at the same time (1) the only two inductive inference operators that avoid the drowning effect and (2) the only two inductive inference operators that have been shown to satisfy syntax splitting, we investigate connections between these two properties in this section. Indeed, since the drowning effect is, essentially, caused by unrelated information (e) influencing inferences, and the two only inductive inference operations known to satisfy syntax splitting also are the two only known inductive inference relations known to avoid the drowning effect, we might conjecture there is an intimate relationship between syntax splitting and the drowning effect. The answer to this question is negative. We show this by constructing an inductive inference relation, system Z^{ind}, that satisfies syntax splitting and is RC-extending, yet still suffers from the drowning problem.

For this, we first recall the concept of a finest syntax splitting [Parikh, 1999], adapted to our setting of conditional belief bases.

Definition 6. Given a conditional belief base Δ, and $\Sigma_1,\ldots,\Sigma_n$ s.t. $\bigcup_{i=1}^{n}\Sigma_i$ and $\Sigma_i \cap \Sigma_j \neq \emptyset$ for every $i,j = 1,\ldots,n$ s.t. $i \neq j$, $\Sigma_1,\ldots,\Sigma_n$ *splits* Δ iff $\Delta^i \subset (\mathcal{L}(\Sigma_i)|\mathcal{L}(\Sigma_i))$ for every $i = 1,\ldots,n$ and $\Delta = \bigcup_i^n \Delta^i$. We also say $\Sigma_1,\ldots,\Sigma_n$ is a *splitting* of Δ.

Given two partitions $\Sigma_1,\ldots,\Sigma_n$ and $\Sigma'_1,\ldots,\Sigma'_m$ of Σ, we say that $\Sigma_1,\ldots,\Sigma_n$ *is finer than* $\Sigma'_1,\ldots,\Sigma'_m$ iff for every $1 \leqslant i \leqslant n$, there is some $1 \leqslant j \leqslant m$ s.t. $\Sigma_i \subseteq \Sigma_j$.

It can be easily shown, similarly to [Parikh, 1999, Lemma 1], that every conditional belief base admits a unique finest splitting. We define the system Z^{ind}-TPO based on a variation of the lexicographic vectors defined in section 2.5. In particular, instead of counting the number rules in a part Δ_i of the partition $\Delta_1,\ldots,\Delta_n$, we now use $V'(\omega,\Delta')$ to count for how many sub-alphabets of the finest splitting $\Sigma_1,\ldots,\Sigma_m$ of Δ, a rule in Δ' is violated by ω. More formally:

Definition 7. Given a conditional belief base Δ partitioned in $(\Delta_0,\ldots,\Delta_n)$ with finest splitting $\Sigma_1,\ldots,\Sigma_m$, $\Theta \subseteq \Delta$ and $\omega \in \Omega(\Sigma)$, we define:

$$V'(\omega,\Theta) = |\{i \in \{1,\ldots,m\} \mid \exists (B|A) \in \Theta^i \text{ s.t. } \omega \models A\overline{B}\}|$$

The Z^{ind}*-vector of* ω is defined as $Z^{\mathsf{ind}}(\omega) = (V'(\omega,\Delta_0),\ldots,V'(\omega,\Delta_n))$.

We can then define $\preceq^{Z^{\mathsf{ind}}}_{\Delta}$ over Ω as: $\omega \preceq^{Z^{\mathsf{ind}}}_{\Delta} \omega'$ iff $Z^{\mathsf{ind}}(\omega) \preceq^{\mathsf{lex}} Z^{\mathsf{ind}}(\omega')$. We define *System* Z^{ind} as the inductive inference operator $C^{tpo}_{\preceq^{Z^{\mathsf{ind}}}_{\Delta}}$. We denote $C^{tpo}_{\preceq^{Z^{\mathsf{ind}}}}$ by $C^{Z^{\mathsf{ind}}}$ and the appertaining inference relation $C^{Z^{\mathsf{ind}}}(\Delta)$ by $\mathrel{|\!\sim}^{Z^{\mathsf{ind}}}_{\Delta}$.

Example 6. We consider the conditional belief base $\Delta = \{(a|\top),(b|\top),(a|b),(c|\top)\}$. Notice that the Z-partitioning of Δ is $\Delta_0 = \Delta$ whereas its finest spltting is $\Sigma_1 = \{a,b\}, \Sigma_2 = \{c\}$. We have:

ω	$Z^{\mathsf{ind}}(\omega)$	ω	$Z^{\mathsf{ind}}(\omega)$	ω	$Z^{\mathsf{ind}}(\omega)$	ω	$Z^{\mathsf{ind}}(\omega)$
abc	(0)	$ab\overline{c}$	(1)	$a\overline{b}c$	(1)	$a\overline{b}\,\overline{c}$	(2)
$\overline{a}bc$	(1)	$\overline{a}b\overline{c}$	(2)	$\overline{a}\overline{b}c$	(1)	$\overline{a}\overline{b}\,\overline{c}$	(2)

Notice that e.g. $Z^{\mathsf{ind}}(\overline{a}b\overline{c}) = 2$ since $\overline{a}b\overline{c}(a|\top) = 0$, $\overline{a}b\overline{c}(c|\top) = 0$ and $\{a,b\}$ and $\{c\}$ are the finest splitting of Δ. Examples of (non-)inferences include: $a \mathrel{|\!\sim}^{Z^{\mathsf{ind}}}_{\Delta} b \wedge c$ and $\neg(b \wedge c) \mathrel{|\!\not\sim}^{Z^{\mathsf{ind}}}_{\Delta} a$.

We can see that for this belief base, Z^{ind} is different from lexicographic entailment. Indeed, we have:

ω	$\mathsf{lex}(\omega)$	ω	$\mathsf{lex}(\omega)$	ω	$\mathsf{lex}(\omega)$	ω	$\mathsf{lex}(\omega)$
abc	(0)	$ab\overline{c}$	(1)	$a\overline{b}c$	(1)	$a\overline{b}\,\overline{c}$	(2)
$\overline{a}bc$	(2)	$\overline{a}b\overline{c}$	(3)	$\overline{a}\overline{b}c$	(2)	$\overline{a}\overline{b}\,\overline{c}$	(3)

We see that $\neg(b \wedge c) \mathrel{|\!\not\sim}^{Z^{\mathsf{ind}}}_{\Delta} a$ whereas $\neg(b \wedge c) \mathrel{|\!\sim}^{\mathsf{lex}}_{\Delta} a$.

System Z^{ind} satisfies syntax splitting:

Theorem 2. System Z^{ind} satisfies **SynSplit**.

We can furthermore show that System Z^{ind} extends System Z and is itself extended by lexicographic entailment:

Proposition 4. (1) System Z^{ind} is RC-extending; and (2) $A \mathrel{|\!\sim}^{Z^{\mathsf{ind}}}_{\Delta} B$ implies $A \mathrel{|\!\sim}^{\mathsf{lex}}_{\Delta} B$ (for any conditional belief base Δ).

Example 6 shows that $A \mathrel{|\!\sim}^{\mathsf{lex}}_{\Delta} B$ does *not* imply $A \mathrel{|\!\sim}^{Z^{\mathsf{ind}}}_{\Delta} B$.

We now show that system Z^{ind} coincides with system Z if the finest splitting of Δ is the whole signature Σ. Intuitively, if the whole signature is the finest splitting, $V'(\omega,\Delta_i) \leqslant 1$, which means $\preceq^{Z^{\mathsf{ind}}}_{\Delta}$ will coincide with κ^{Z}_{Δ}.

Proposition 5. Given a conditional belief base Δ over the signature Σ with finest splitting Σ, $C^{Z^{\mathsf{ind}}}(\Delta) = C^{Z}(\Delta)$.

Proposition 4 shows that System Z^{ind} is weaker than lexicographic inference. How much weaker? Well, for one, Z^{ind} does still suffer from the drowning problem, as the unique finest splitting of Δ is the (trivial) splitting $\{p,b,f,e\}$, we see that $\mathrel{|\!\sim}^{Z^{\mathsf{ind}}}_{\Delta} = \mathrel{|\!\sim}^{Z}_{\Delta}$ for Δ as in Example 5 (with Proposition 5). Altogether, this shows that the property of **SynSplit** is independent of the drowning effect, as Z^{ind} is an inductive inference relation that satisfies **SynSplit** but suffers from the drowning effect.

Remark 1. The fact that system Z^{ind} satisfies syntax splitting and RC-extending also disproves a potential conjecture about how to characterise lexicographic inference axiomatically, namely by syntax splitting and being RC-extending.

Remark 2. What seems to distinguish System Z ind and lexicographic inference and ensure that the latter does not suffer from the drowning problem is the following property:

Qual If $|\{r \in \Delta \mid \omega(r) = 0\}| \subset |\{r \in \Delta \mid \omega'(r) = 0\}|$ then $\omega \prec_\Delta \omega'$.[2]

Qual says that if ω falsifies strictly fewer conditionals than ω', this should be reflected in the ordering $\prec_\Delta$. Clearly, this is satisfied by lexicographic but not by Z^{ind}- or Z-inference. It is not hard to see that, at least for the Tweety-example, **Qual** ensures the drowning effect is avoided, as any world ω s.t. $\omega \models b \wedge \neg e$ will be ranked strictly higher than any ω' that agrees with ω on all atoms except w s.t. $\omega' \models b \wedge e$, as $\omega' \models p \wedge b \wedge e$ implies $(e|b)(\omega') = 0$ and $(e|b)(\omega)$ implies $\omega \prec_\Delta \omega'$. It remains a question for future work to formalize the drowning effect into a general property of inductive inference relations, and to study the conjectured connection between such a property and **Qual**.

We conclude by noticing that lexicographic inference is neither the least refined (witnessed by Z^{ind}) nor the most refined RC-extending inference relation that satisfies **SynSplit**. The following example illustrates the latter claim:

Example 7. Consider again $\Delta = \{r_1 = (a|\top), r_2 = (b|\top)\}$ from Example 4. $\prec^{\mathsf{lex}}_\Delta$ is detailed in Example 4.

For c-representations of Δ, however, we get the following c-representations:

ω	$\kappa^c_1(\omega)$	$\kappa^c_2(\omega)$	$\kappa^c_3(\omega)$	$\kappa^c_4(\omega)$
ab	0	0	0	0
$a\overline{b}$	2	1	1	2
$\overline{a}b$	1	2	1	2
$\overline{a}\overline{b}$	3	3	2	4

κ^c_1 and κ^c_2 are both strict refinements of $\preceq_\Delta$, and therefore are RC-extending. In view of [Kern-Isberner *et al.*, 2020, Proposition 9], κ^c_1 and κ^c_2 satisfy **SynSplit**.

In this section, we have shown that (1) lexicographic inference is not the only inductive inference relation that is RC-extending and satisfies syntax splitting and (2) that there are (RC-extending) inductive inference relations that still suffer from the drowning effect. This points to several open questions in the field of non-monotonic reasoning. First, how to axiomatically characterize lexicographic inference? Secondly, what is the exact relation between notions of independence and relevance on the one hand, and the drowning effect on the other hand. Intuitively, these concepts seem to be on the core of why the drowning effect is undesirable: in the belief base Δ, we have no reasons to assume that having a beak (e) is related to flying (f), once we know that something is a bird (b).

[2] An alternative version of **Qual** is **Qual'**: If $\{r \in \Delta \mid \omega(r) = 0\} \subset \{r \in \Delta \mid \omega'(r) = 0\}$ then $\omega \prec_\Delta \omega'$. Notice that this version is stronger than our formulation (i.e. **Qual** implies **Qual'**, but not vice-versa).

6 Related Work

Considerations pertaining to relevance and independence in logic are numerous, but the notion of syntax splitting comes from [Kern-Isberner *et al.*, 2020], which was in turn inspired by similar considerations in belief revision [Kern- Isberner and Brewka, 2017b]. Other notions of relevance and syntax splitting in belief revision can be found in [Parikh, 1999; Aravanis *et al.*, 2019] and are compared in depth with syntax splitting in [Kern-Isberner *et al.*, 2020].

To the best of our knowledge, no axiomatic notion of relevance or independence has been studied for lexicographic inference before. Furthermore, System Z^{Ind} has, as far as we know, not been proposed in the literature. This is not surprising, as it was designed to disprove several conjectures on lexicographic inference. There are a couple of non-monotonic logics that are somewhat related to system Z^{Ind} and lexicographic inference which we mention here. Firstly, System W [Komo and Beierle, 2021] bases inferences on a partial order over worlds obtained by comparing, given a Z-partitionend conditional belief base $(\Delta_0, \ldots, \Delta_n$, the vector $(\eta_0(\omega), \ldots, \eta_n(\omega))$, where $\eta_i(\omega)$ contains the conditionals in Δ_i violated by ω. These vectors are compared lexicographically, but using the $\subseteq$-relation instead of the cardinality relation. This results in partial order over worlds, which means, among others, that the resulting inference relation does not satisfy rational monotony. In [Haldimann and Beierle, 2022], it is shown that system W satisfies syntax splitting as well. The exact relationship with lexicographic inference remains a question for future work. Another related logic is System ARS [Ritterskamp and Kern-Isberner, 2008], where so-called tolerance layers are combined lexicographically. However, in [Ritterskamp and Kern-Isberner, 2008], it is shown that lexicographic inference and system ARS are incomparable.

In [Casini *et al.*, 2019], 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.

7 Conclusion

The main results of this paper are the following: (1) lexicographic inference satisfies syntax splitting, thus establishing it as an RC-extending inference relation of high quality; (2) c-representations and lexicographic inference are related, but not identical; (3) syntax splitting does not uniquely characterize lexicographic inference and (4) the drowning problem and syntax splitting are unrelated. The latter negative results point to two open questions, to which we have pointed first suggestions of where to look for answers. Further future work includes taking advantage of syntax splitting in the development of efficient solvers, and the generalization of syntax splitting to first-order and (probabilistic) description logics and richer preferential languages, where lexicographic inference [Casini and Straccia, 2012; Giugno and Lukasiewicz, 2002; Booth *et al.*, 2015] has been applied extensively.

Acknowledgments

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

References

[Aravanis *et al.*, 2019] Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams. Full characterization of parikh's relevance-sensitive axiom for belief revision. *Journal of Artificial Intelligence Research*, 66:765–792, 2019.

[Beierle *et al.*, 2018] Christoph Beierle, Christian Eichhorn, Gabriele Kern-Isberner, and Steven Kutsch. Properties of skeptical c-inference for conditional knowledge bases and its realization as a constraint satisfaction problem. *Annals of Mathematics and Artificial Intelligence*, 83(3):247–275, 2018.

[Booth *et al.*, 2015] Richard Booth, Giovanni Casini, Thomas Andreas Meyer, and Ivan José Varzinczak. On the entailment problem for a logic of typicality. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

[Bourne and Parsons, 1999] Rachel A Bourne and Simon Parsons. Connecting lexicographic with maximum entropy entailment. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 80–91. Springer, 1999.

[Casini and Straccia, 2012] Giovanni Casini and Umberto Straccia. Lexicographic closure for defeasible description logics. In *Proc. of Australasian Ontology Workshop*, volume 969, pages 28–39, 2012.

[Casini *et al.*, 2019] Giovanni Casini, Thomas Meyer, and Ivan Varzinczak. Taking defeasible entailment beyond rational closure. In *European Conference on Logics in Artificial Intelligence*, pages 182–197. Springer, 2019.

[de Finetti, 1974] Bruno de Finetti. Theory of probability (2 vols.). *John Wiley & Sons*, 1974.

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

[Giugno and Lukasiewicz, 2002] Rosalba Giugno and Thomas Lukasiewicz. P-shoq (d): A probabilistic extension of shoq (d) for probabilistic ontologies in the semantic web. In *JELIA*, volume 2, pages 86–97. Springer, 2002.

[Goldszmidt and Pearl, 1996] Moisés Goldszmidt and Judea Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *AI*, 84(1-2):57–112, 1996.

[Haldimann and Beierle, 2022] Jonas Haldimann and Christoph Beierle. Inference with system W satisfies syntax splitting. In *19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022*, 2022. (accepted for publication).

[Kern-Isberner and Brewka, 2017] Gabriele Kern-Isberner and Gerhard Brewka. Strong syntax splitting for iterated belief revision. In C. Sierra, editor, *Proceedings International Joint Conference on Artificial Intelligence, IJCAI 2017*, pages 1131–1137. ijcai.org, 2017.

[Kern-Isberner *et al.*, 2020] Gabriele Kern-Isberner, Christoph Beierle, and Gerhard 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, pages 560–571, 2020.

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

[Komo and Beierle, 2021] Christian Komo and Christoph Beierle. Nonmonotonic reasoning from conditional knowledge bases with system w. *Annals of Mathematics and Artificial Intelligence*, pages 1–38, 2021.

[Kraus *et al.*, 1990] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence*, 44(1-2):167–207, 1990.

[Lehmann and Magidor, 1992] Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artificial intelligence*, 55(1):1–60, 1992.

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

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

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

[Ritterskamp and Kern-Isberner, 2008] Manuela Ritterskamp and Gabriele Kern-Isberner. Preference-based default reasoning. In *FLAIRS Conference*, pages 672–677, 2008.

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

[Strotkamp, 2020] Isabell Philine Strotkamp. Lexikographische inferenzen im vergleich mit c-repräsentationen, 2020. Bachelor's Thesis, Technische Universität Dortmund.