Selection Strategies for Inductive Reasoning From Conditional Belief Bases and for Belief Change Respecting the Principle of Conditional Preservation

Christoph Beierle
FernUniversität in Hagen
Hagen, Germany
christoph.beierle@fernuni-hagen.de

Gabriele Kern-Isberner
TU Dortmund University
Dortmund, Germany
gabriele.kern-isberner@cs.tu-dortmund.de

Abstract

Given a belief base Δ consisting of a set of conditionals, there are many different ways an agent may inductively complete the knowledge represented by Δ to a complete epistemic state; two well-known approaches are given by system P and system Z, and also each ranking model of Δ induces a full inference relation. C-representations are special ranking models that obey the principle of conditional indifference. Inductive reasoning using c-representations can be done with respect to all c-representations, with respect to a subclass of, e.g., minimal c-representations, or with respect to single c-representations. In this paper, we present and investigate selection strategies for determining single c-representations to be used for inductive reasoning from belief bases. We develop axioms for specifying characteristics of selection strategies. We illustrate which desirable properties, like syntax splitting, are ensured by the axioms, and develop constructions for obtaining selection strategies satisfying the axioms. Furthermore, we also present and study the extension of selection strategies to c-revisions that follow the principle of conditional preservation and that have been employed successfully in various belief change settings.

1 Introduction

Conditionals of the form *If A then usually B* are an established standard for expressing plausible, yet uncertain knowledge. Given a belief base Δ consisting of a set of conditionals, there are many different ways an agent may inductively complete the knowledge represented by Δ to a complete epistemic state. Such a completion allows the agent to reason from Δ and to answer all queries of the form whether a proposition A entails a proposition B in the context of Δ. Well-known classical approaches to achieve this are, for instance, given by system P (Lehmann and Magidor 1992), taking all models of Δ into account, or by system Z (Pearl 1990) which selects a specific minimal ranking model for inducing an inference relation. Another approach is to use c-representations which constitute a subclass of all ranking models of Δ respecting the principle of conditional indifference (Kern-Isberner 2004). Inductive reasoning using c-representations can be done with respect to all c-representations, with respect to a subclass of, e.g., minimal c-representations, or with respect to single c-representations (Kern-Isberner 2004; Beierle et al. 2018). A first use of employing a selection strategy for choosing c-representations along with a single postulate closely related to syntax splitting is given in (Kern-Isberner, Beierle, and Brewka 2020), but no further elaboration on selection strategies and no construction of syntax-splitting obeying selection strategies is provided there. In this paper, we present and investigate selection strategies for determining c-representations to be used for inductive reasoning from conditional belief bases in a general form. The principal strength of c-representations is that conditional information in the belief base is processed transparently by so-called conditional impact factors. With strategies, we are able to fully explore this by guiding also the selection of impact factors according to general principles. In this way, inductive reasoning problems can be analysed in more depth by analysing the set of involved conditionals, or the constraint problems that they invoke, respectively, so that similar subproblems of different inductive reasoning tasks can be identified and solved in a coherent way. We develop axioms for specifying characteristics of selection strategies, illustrate which desirable properties are ensured by the axioms, and develop constructions for obtaining selection strategies satisfying the axioms. Furthermore, we also present and discuss the extension of selection strategies for c-representations to c-revisions of ranking functions (Spohn 1988) by sets of conditionals. Here, the use of selection strategies is even more rewarding. They make it possible to identify similar subproblems in two formally different revision tasks involving different prior ranking functions and different new information given by sets of conditionals, and provide coherent solutions. Due to close connections between c-revisions and c-representations, the axioms and properties for c-representations can be lifted to c-revisions.

In summary, the main contributions of this paper are: (1) selection strategies for inductive reasoning and for belief change with respect to sets of conditionals, (2) axioms for specifying characteristics of selection strategies and their interrelationships, and (3) constructions of selection strategies satisfying specific axioms.

2 Background

Let $\mathcal{L}$ be a finitely generated propositional language over an alphabet Σ with atoms $a, b, c, \ldots$, and with formulas $A, B, C, \ldots$. For subsets Θ of Σ, let $\mathcal{L}(\Theta)$ denote the propositional language defined by Θ, with associated set of interpretations $\Omega(\Theta)$. For conciseness of notation, AB stands for $A \wedge B$, and $\overline{A}$ means $\neg A$. Let Ω denote the set of *possible worlds* over $\mathcal{L}$; each $\omega \in \Omega$ will be taken both as a propositional interpretation and as a complete conjunction over $\mathcal{L}$. The set of all models of A is denoted by $Mod(A)$, and $A \models B$ holds iff $Mod(A) \subseteq Mod(B)$, as usual.

By making use of a conditional operator $|$, we introduce the set $(\mathcal{L}|\mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}$ of *conditionals* over $\mathcal{L}$. Conditionals $(B|A)$ are meant to express plausible, defeasible rules "If A then plausibly (usually, possibly, probably, typically etc.) B". For interpreting conditionals, we use *ordinal conditional functions (OCFs)*, (also called *ranking functions*) $\kappa : \Omega \to \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$, first introduced (in a more general form) by (Spohn 1988). They express degrees of plausibility of propositional formulas. Formally, we have $\kappa(A) := \min\{\kappa(\omega) \mid \omega \models A\}$, so that $\kappa(A \vee B) = \min\{\kappa(A), \kappa(B)\}$. A proposition A is believed if $\kappa(\overline{A}) > 0$ (which implies particularly $\kappa(A) = 0$). The *uniform OCF* κ_u is defined by $\kappa_u(\omega) = 0$ for all $\omega \in \Omega$.

A conditional $(B|A)$ is *accepted* in the epistemic state represented by κ, written as $\kappa \models (B|A)$, iff $\kappa(AB) < \kappa(A\overline{B})$, i.e. iff AB is more plausible than $A\overline{B}$. Every OCF κ induces a nonmonotonic inference relation given by:

$$A \mid\!\sim_{\kappa} B \quad \text{iff} \quad A \equiv \bot \quad \text{or} \quad \kappa(AB) < \kappa(A\overline{B}) \tag{1}$$

We say that two conditionals $(B|A)$ and $(B'|A')$ are *conditionally equivalent*, denoted by $(B|A) \equiv_{ce} (B'|A')$, if they have the same verification and the same falsification behaviour, i.e., if $AB \equiv A'B'$ and $A\overline{B} \equiv A'\overline{B'}$ (de Finetti 1937). A belief base Δ is a finite set of conditionals. In order to avoid lengthy case distinctions, througout this paper we will assume, unless stated otherwise, that Δ does not contain any conditional $(B|A)$ that is conradictory ($AB \equiv \bot$) or that is self-fulfilling ($A \models B$) and that Δ is *duplicate free*, i.e., $(B|A), (B'|A') \in \Delta$ with $(B|A) \equiv_{ce} (B'|A')$ implies $(B|A) = (B'|A')$. Furthermore, when dealing with two sets Δ, Δ' of conditionals and enumerating their elements, we will tacitly assume that conditionally equivalent conditionals are given first and in the same order in Δ and in Δ'. Thus, if $\Delta = \{(B_1|A_1), \ldots, (B_m|A_m)\}$ and $\Delta' = \{(B'_1|A'_1), \ldots, (B'_{m'}|A'_{m'})\}$, then there is $k \in \{0, \ldots, \min\{m, m'\}\}$ such that for all $i \in \{1, \ldots, k\}$, $(B_i|A_i) \equiv_{ce} (B'_i|A'_i)$ and there is no conditional in $\{(B_{k+1}|A_{k+1}), \ldots, (B_m|A_m)\}$ that is conditionally equivalent to any conditional in $\{(B'_{k+1}|A'_{k+1}), \ldots, (B'_{m'}|A'_{m'})\}$. We extend $\equiv_{ce}$ to sets of conditionals Δ and Δ' as above and define $\Delta \equiv_{ce} \Delta'$ if $m = m'$ and $(B_i|X_i) \equiv_{ce} (B'_i|X'_i)$ for all $i \in \{1, \ldots, m\}$.

3 Inductive Reasoning from Conditionals

3.1 C-Representations and Selection Strategies

Among the OCF models of Δ, c-representations are special ranking models obtained by assigning individual integer impacts to the conditionals in Δ and generating the world ranks as the sum of impacts of falsified conditionals.

Definition 1 (c-representation (Kern-Isberner 2001)). A *c-representation* of a conditional belief base $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ is an OCF κ constructed from non-negative integer impacts $\eta_j \in \mathbb{N}_0$ assigned to each $(B_j|A_j)$ such that κ accepts Δ and is given by:

$$\kappa(\omega) = \sum_{\substack{1 \leqslant j \leqslant n \\ \omega \models A_j \overline{B}_j}} \eta_j \tag{2}$$

C-representations can conveniently be specified using a constraint satisfaction problem (for detailed explanations, see (Kern-Isberner 2001; Kern-Isberner 2004; Beierle et al. 2018; Beierle et al. 2021)):

Definition 2 ($CR(\Delta)$, cr_i^{Δ}). Let $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$. The *constraint satisfaction problem for c-representations of* Δ, denoted by $CR(\Delta)$, is given by the conjunction of the set of constraints cr_i^{Δ}, for all $i \in \{1, \ldots, n\}$:

$$(cr_i^{\Delta}) \qquad \eta_i > \min_{\omega \models A_i B_i} \sum_{\substack{k \neq i \\ \omega \models A_k \overline{B_k}}} \eta_k - \min_{\omega \models A_i \overline{B}_i} \sum_{\substack{k \neq i \\ \omega \models A_k \overline{B_k}}} \eta_k \tag{3}$$

where the η_i are constraint variables taking values in $\mathbb{N}$.

The condition that the η_i take non-negative values expresses that falsification of conditionals should not make worlds more plausible, and (3) ensures that κ as specified by (2) accepts Δ.

A solution of $CR(\Delta)$ is a vector $\overrightarrow{\eta} = (\eta_1, \ldots, \eta_n)$ of natural numbers. $Sol(CR(\Delta))$ denotes the set of all solutions of $CR(\Delta)$. For $\overrightarrow{\eta} \in Sol(CR(\Delta))$ and κ as in Equation (2), κ is the *OCF induced by* $\overrightarrow{\eta}$ and is denoted by $\kappa_{\overrightarrow{\eta}}$. $CR(\Delta)$ is sound and complete (Kern-Isberner 2001; Beierle et al. 2018): For every $\overrightarrow{\eta} \in Sol(CR(\Delta))$, $\kappa_{\overrightarrow{\eta}}$ is a c-representation with $\kappa_{\overrightarrow{\eta}} \models \Delta$, and for every c-representation κ with $\kappa \models \Delta$, there is $\overrightarrow{\eta} \in Sol(CR(\Delta))$ such that $\kappa = \kappa_{\overrightarrow{\eta}}$.

Every c-representation κ with $\kappa \models \Delta$ expands the beliefs in Δ. A selection function can be used for modelling the different possible choices of which c-representation should be selected.

Definition 3 (selection strategy σ (for c-representations)). A *selection strategy (for c-representations)* is a function σ

$$\sigma : \Delta \mapsto \overrightarrow{\eta}$$

assigning to each conditional belief base Δ an impact vector $\overrightarrow{\eta} \in Sol(CR(\Delta))$.

Any selection strategy σ gives rise to an inductive inference operator for c-representations by assigning an inference relation to every belief base Δ via the function

$$\mathbf{C}_{\sigma}^{\textit{c-rep}} : \Delta \mapsto \mid\!\sim_{\kappa_{\sigma(\Delta)}} \tag{4}$$

where $\mid\!\sim_{\kappa_{\sigma(\Delta)}}$ is obtained via Equation (1) (cf. (Kern-Isberner, Beierle, and Brewka 2020)). In principle, for every Δ, a selection strategy may choose some impact vector independently from the choices for all other belief bases Δ'. However, depending on the relationships between Δ and Δ', there are desirable properties for selection strategies that relate $\sigma(\Delta)$ to $\sigma(\Delta')$. In the following, we specify axioms for selection strategies which aim at ensuring such desirable properties. First, we present a postulate requiring selection strategies to be independent of the particular syntactic representation of Δ.

(SI) A selection strategy σ (for c-representations) is *syntax independent* if for any Δ' obtained from Δ by replacing a conditional $(B|A)$ occurring in Δ by $(B'|A')$ with $(B|A) \equiv_{ce} (B'|A')$, we have $\sigma(\Delta) = \sigma(\Delta')$.

The following observation is a direct consequence of syntax independence.

Proposition 1. *Let σ be a selection strategy satisfying (**SI**) and let Δ, Δ' be sets of conditionals. If $\Delta \equiv_{ce} \Delta'$ then $\sigma(\Delta) = \sigma(\Delta')$ and thus $\mathrel{|\!\sim}_{\kappa_{\sigma(\Delta)}} = \mathrel{|\!\sim}_{\kappa_{\sigma(\Delta')}}$.*

A more general requirement than syntax independence is the requirement that a selection strategy should be dependent only on the solution space of the respective constraint satisfaction problem:

(IP-SOL) A selection strategy σ is *impact preserving with respect to the solution space* if for any two belief bases Δ, Δ' with $Sol(CR(\Delta)) = Sol(CR(\Delta'))$, we have $\sigma(\Delta) = \sigma(\Delta')$.

Note that if $CR(\Delta) = CR(\Delta')$, then a selection strategy fulfilling (**IP-SOL**) will also trivially satisfy $\sigma(\Delta) = \sigma(\Delta')$. In particular, this also implies that the selection strategy σ and the resulting inference relation are syntax independent.

Proposition 2. *If σ satisfies (**IP-SOL**), σ also satisfies (**SI**).*

In the following subsection, we will study further uses of selection strategies and corresponding postulates for them.

3.2 Splittings of Belief Bases

Another desirable property of selection strategies surfaces when considering a belief base Δ that splits into subbases Δ_1, Δ_2 over disjoint sublanguages. Formally, this means $\Delta = \Delta_1 \cup \Delta_2$, $\Delta_i \subset (\mathcal{L}_i|\mathcal{L}_i), \mathcal{L}_i = \mathcal{L}(\Sigma_i)$ for $i = 1, 2$ such that $\Sigma_1 \cap \Sigma_2 = \emptyset$ and $\Sigma_1 \cup \Sigma_2 = \Sigma$; we will write

$$\Delta = \Delta_1 \bigcup_{\Sigma_1,\Sigma_2} \Delta_2$$

in this case. Note that this base case covers the most general case of syntax splitting which can be recursively applied to finer syntax splittings (in the sense of Parikh (Parikh 1999)).

In such a situation, an inference from A to B with A, B being propositional formulas over one of the sublanguages should be determined by the respective subbase, and moreover, the inference should be uneffected by any additional information stemming from the other sublanguage. This view on syntax splitting is proposed and elaborated in (Kern-Isberner, Beierle, and Brewka 2020), covering also semantic counterparts in the frameworks of total preorders (Makinson 1989) and OCFs. For inductive reasoning with c-representations, a selection strategy is used for ensuring syntax splitting of the induced inference operator, but only a single postulate closely related to syntax splitting is given in (Kern-Isberner, Beierle, and Brewka 2020). Writing $(\vec{\eta}^1, \vec{\eta}^2)$ for the composition of impact vectors $\vec{\eta}^1, \vec{\eta}^2$, the postulate (**IP-Split**), called (**IP***c-rep*) in (Kern-Isberner, Beierle, and Brewka 2020), reads as follows:

(IP-Split) A selection strategy σ (for c-representations) is *impact preserving* if $\sigma(\Delta) = (\sigma(\Delta_1), \sigma(\Delta_2))$ for any $\Delta = \Delta_1 \bigcup_{\Sigma_1,\Sigma_2} \Delta_2$.

No further elaboration on selection strategies and no construction of syntax-splitting obeying selection strategies is provided in (Kern-Isberner, Beierle, and Brewka 2020). The following proposition gives a nondeterministic algorithm yielding exactly all selection strategies that are impact preserving with respect to syntax splitting.

Proposition 3. *A selection strategy σ satisfies (**IP-Split**) iff it can be obtained by applying the following steps:*

- *For the empty belief base $\Delta = \emptyset$ choose the empty vector $()$ for $\sigma(\Delta)$.*
- *For any nonempty belief base Δ, if there is a splitting $\Delta = \Delta_1 \bigcup_{\Sigma_1,\Sigma_2} \Delta_2$ with $\Delta_1 \neq \emptyset$ and $\Delta_2 \neq \emptyset$, choose $(\sigma(\Delta_1), \sigma(\Delta_2))$ for $\sigma(\Delta)$.*
- *Otherwise, choose any impact vector $\vec{\eta} \in Sol(CR(\Delta))$ for $\sigma(\Delta)$.*

Proof. $\Leftarrow$ direction: Let σ be obtained from executing the given steps. We show (**IP-Split**) by induction on the number of conditionals in Δ. For any $\Delta = \Delta'_1 \bigcup_{\Sigma'_1,\Sigma'_2} \Delta'_2$, we have to show $\sigma(\Delta) = (\sigma(\Delta'_1), \sigma(\Delta'_2))$.

For $|\Delta| = 0$, we have $\Delta = \Delta_1 = \Delta_2 = \emptyset$ and thus $\sigma(\Delta) = () = ((), ()) = (\sigma(\Delta'_1), \sigma(\Delta'_2))$.

For $|\Delta| \geqslant 1$, a syntax splitting $\Delta = \Delta_1 \bigcup_{\Sigma_1,\Sigma_2} \Delta_2$ must have been taken into account in the construction of σ because there is a syntax splitting for Δ. There are four cases:
(i) If $\Delta'_1 = \Delta_1$, then $\Delta'_2 = \Delta_2$ and we are done.
(ii) If $\Delta'_1 \subsetneq \Delta_1$, then in addition to the syntax splittings $\Delta = \Delta_1 \bigcup_{\Sigma_1,\Sigma_2} \Delta_2$ and $\Delta = \Delta'_1 \bigcup_{\Sigma'_1,\Sigma'_2} \Delta'_2$ we also have:

$$\Delta_1 = \Delta'_1 \bigcup_{\Sigma'_1, \Sigma_1 \setminus \Sigma'_1} \Delta_1 \setminus \Delta'_1 \tag{5}$$

$$\Delta'_2 = \Delta_1 \setminus \Delta'_1 \bigcup_{\Sigma_1 \setminus \Sigma'_1, \Sigma_2} \Delta_2 \tag{6}$$

Using the induction hypothesis on Δ_1 and on Δ'_2, we get:

$$\begin{aligned}\sigma(\Delta) = (\sigma(\Delta_1), \sigma(\Delta_2)) &= ((\sigma(\Delta'_1), \sigma(\Delta_1 \setminus \Delta'_1)), \sigma(\Delta_2)) \\ &= (\sigma(\Delta'_1), (\sigma(\Delta_1 \setminus \Delta'_1), \sigma(\Delta_2))) = (\sigma(\Delta'_1), \sigma(\Delta'_2))\end{aligned}$$

(iii) If $\Delta'_1 \subsetneq \Delta_2$, the proof is as in case (ii) after swapping the roles of Δ'_1 and Δ'_2.
(iv) In the remaining case, we have $\Delta_i \cap \Delta'_j \neq \emptyset$ for $i, j \in \{1, 2\}$ with $i \neq j$ and the additional syntax splittings:

$$\Delta_1 = \Delta_1 \setminus \Delta'_1 \bigcup_{\Sigma_1 \setminus \Sigma'_1, \Sigma_1 \setminus \Sigma'_2} \Delta_1 \setminus \Delta'_2 \tag{7}$$

$$\Delta_2 = \Delta_2 \setminus \Delta'_1 \bigcup_{\Sigma_2 \setminus \Sigma'_1, \Sigma_2 \setminus \Sigma'_2} \Delta_2 \setminus \Delta'_2 \tag{8}$$

$$\Delta'_2 = \Delta_2 \setminus \Delta'_1 \bigcup_{\Sigma_2 \setminus \Sigma'_1, \Sigma_1 \setminus \Sigma'_1} \Delta_1 \setminus \Delta'_1 \tag{9}$$

$$\Delta'_1 = \Delta_1 \setminus \Delta'_2 \bigcup_{\Sigma_1 \setminus \Sigma'_2, \Sigma_2 \setminus \Sigma'_2} \Delta_2 \setminus \Delta'_2 \tag{10}$$

Using again the induction hypothesis, the proof is similar to case (ii).

$\Rightarrow$ direction: Let σ be a selection strategy satisfying (**IP-Split**). We can easily construct σ using the given steps; in particular, if there are several syntax splittings, any choice in the second step yields the same result determined by σ. $\square$

The significance of selection strategies for syntax splitting is given by the fact that (**IP-Split**) precisely characterizes the inductive operators based on single c-representations that satisfy syntax splitting (Kern-Isberner, Beierle, and Brewka 2020, Proposition 9).

3.3 Projections of Belief Bases

If $\vec{\eta}$ is an impact vector with impacts corresponding to the conditionals in Δ, then for $\Delta' \subseteq \Delta$, the subvector of $\vec{\eta}$ containing only the impacts related to the conditionals in Δ' is called the *projection* of $\vec{\eta}$ to Δ' and is denoted by $\vec{\eta}_{\Delta'}$. Hence, if σ is a selection strategy, $\sigma(\Delta)_{\Delta'}$ is the projection of $\sigma(\Delta)$ to Δ'.

The following definition extends the notion of projection to constraint satisfaction problems for c-representations.

Definition 4 (CSP projection $CR(\Delta)_{\Delta'}$)**.** Let $\Delta = \{(B_1|A_1), \dots, (B_n|A_n)\}$ be a set of conditionals, and let $\Delta' \subseteq \Delta$. The *projection of* $CR(\Delta)$ *to* Δ', denoted by $CR(\Delta)_{\Delta'}$, is the constraint satisfaction problem given by the set of constraints $\{cr_i^{\Delta} \mid (B_i|A_i) \in \Delta'\}$.

W.l.o.g. let us assume that $\Delta = \{(B_1|A_1), \dots, (B_n|A_n)\}$ and $\Delta' = \{(B_1|A_1), \dots, (B_k|A_k)\}$ where $k \leqslant n$. Note the difference between

$$CR(\Delta') = \{cr_1^{\Delta'}, \dots, cr_k^{\Delta'}\} \qquad (11)$$

and

$$CR(\Delta)_{\Delta'} = \{cr_1^{\Delta}, \dots, cr_k^{\Delta}\}. \qquad (12)$$

While $CR(\Delta')$ in (11) is a CSP over the constraint variables $\eta_1, \dots, \eta_k$, $CR(\Delta)_{\Delta'}$ in (12) is a CSP over the constraint variables $\eta_1, \dots, \eta_k, \eta_{k+1}, \dots, \eta_m$. Both CSP have $|\Delta'|$-many constraints, but in contrast to cr_i^{Δ}, any of the constraint variables $\eta_{k+1}, \dots, \eta_m$ in the sum in the minimizations terms given in Equation (2) are removed in $cr_i^{\Delta'}$.

Using projections of constraint satisfaction problems for c-revisions, we can generalize the idea of preserving impacts, as expressed by (**IP-SOL**) and (**SI**) to equivalent subproblems as specified in the next axiom.

(IP-ESP) A selection strategy σ is *impact preserving with respect to equivalent subproblems* if for any two belief bases Δ, Δ' with subbases $\Delta_1 \subseteq \Delta$ and $\Delta'_1 \subseteq \Delta'$ and with $\Delta_1 \equiv_{ce} \Delta'_1$ such that $CR(\Delta)_{\Delta_1} = CR(\Delta')_{\Delta'_1}$, we have $\sigma(\Delta)_{\Delta_1} = \sigma(\Delta')_{\Delta'_1}$.

Example 1. We consider the two disjoint signatures $\Sigma_1 = \{p, b, f\}$ and $\Sigma_2 = \{s, h, g, c\}$ where the atoms have the meanings $p = \mathit{penguin}, b = \mathit{bird}, f = \mathit{fly}$ and $s = \mathit{sunny}, h = \mathit{hiking}, g = \mathit{gym}, c = \mathit{cinema}$, and the two conditional belief bases $\Delta_1 = \{r_1^1 = (f|b), r_2^1 = (b|p), r_3^1 = (\overline{f}|p)\}$ and $\Delta_2 = \{r_1^2 = (h|s), r_2^2 = (g|\overline{s})$, expressing that birds usually fly, penguins are birds, but usually, penguins cannot fly, respectively, if the weather is sunny, people usually go on a hike, and if it is not, they usually go to the gym. If we set $\Sigma = \Sigma_1 \cup \Sigma_2$ and $\Delta = \Delta_1 \cup \Delta_2$, then we have $\Delta = \Delta_1 \bigcup_{\Sigma_1, \Sigma_2} \Delta_2$. Following Proposition 3, we choose impact factors for building up a c-representation for each Δ_i first, and combine them in a joint vector afterwards. In both cases, we pursue the general strategy of selecting (pareto-)minimal impact factors and obtain $\sigma(\Delta_1) = (1, 2, 2)$ resp. $\sigma(\Delta_2) = (1, 1)$ (all numbers in the respective ordering of the rules). According to (**IP-Split**), the selection strategy would choose $\sigma(\Delta) = (1, 2, 2, 1, 1)$, and together with (2), we obtain a c-representation κ_Δ of Δ immediately.

A consequence is that if σ satisfies (**IP-ESP**), then it is syntax independent.

Proposition 4. *If σ satisfies (**IP-ESP**), σ also satisfies (**SI**).*

Another consequence of (**IP-ESP**) is that self-fulfilling conditionals do not influence the induced inference relation.

Proposition 5. *If σ satisfies (**IP-ESP**), then $\mathbf{C}_\sigma^{c\text{-}rep}(\Delta) = \mathbf{C}_\sigma^{c\text{-}rep}(\Delta \cup (B|A))$ if $A \models B$.*

Proof. Let σ be a selection strategy, $\Delta_0 = \{(B_1|A_1), \dots, (B_{n-1}|A_{n-1}\}$ a set of conditionals, $\Delta_1 = \{(B_n|A_n)\}$ such that $A_n \models B_n$, and let $\Delta = \Delta_0 \cup \Delta_1$. Thus, due to (4), it suffices to show $\kappa_{\sigma(\Delta)} = \kappa_{\sigma(\Delta \cup \{(B_n|A_n)\})}$.

Consider $CR(\Delta)$ and $CR(\Delta_0)$. Because the condition $\omega \models A_n\overline{B_n}$ used in the construction of the sum expressions in $cr_{\kappa,i}^{\Delta}$ is never satisfied, we get $cr_i^{\Delta} = cr_i^{\Delta_0}$ for all $i \in \{1, \dots, n-1\}$ and hence $CR(\Delta)_{\Delta_0} = CR(\Delta_0)_{\Delta_0} = CR(\Delta_0)$. Because σ satisfies (**IP-ESP**), we obtain $\sigma(\Delta)_{\Delta_0} = \sigma(\Delta_0)$. Moreover, $A_n \models B_n$ implies that cr_n^{Δ} is just $\eta_n > 0$. Therefore, σ may choose any natural number $k > 0$ as impact for $(B_n|A_n)$ in $CR(\Delta)$. For every $k > 0$, both $\vec{\eta} = \sigma(\kappa, \Delta_0)$ and $\vec{\eta}' = (\vec{\eta}, k) = \sigma(\kappa, \Delta)$ induce the same ranking function, i.e., $\kappa_{\vec{\eta}}(\omega) = \kappa_{\vec{\eta}'}(\omega)$ according to (13) for all ω because no world falsifies $(B_n|A_n)$, implying $\kappa_{\sigma(\Delta)} = \kappa_{\sigma(\Delta \cup \{(B_n|A_n)\})}$. $\square$

In the following section, we will show how the axioms and their implied properties evolve in a more general setting.

4 Selection Strategies for Belief Revision

C-representations can be generalized to c-revisions by taking an a priori OCF into account. C-revisions provide a highly general framework for revising ranking functions by sets of conditionals while respecting the so-called *principle of conditional preservation* (Kern-Isberner 2004; Kern-Isberner 2018). Each c-revision is an iterated revision in the sense of Darwiche and Pearl (Darwiche and Pearl 1997) because the DP-postulates are implied by the principle of conditional preservation. The following definition provides the basic version of c-revisions.

Definition 5 (C-revisions for OCFs)**.** Let κ be an OCF and $\Delta = \{(B_1|A_1), \dots, (B_n|A_n)\}$ a set of conditionals. Then a *c-revision of* κ *by* Δ is an OCF κ^* constructed from nonnegative integers η_i assigned to each $(B_i|A_i)$ and an integer κ_0

such that κ^* accepts Δ and is given by:

$$\kappa^*(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} \eta_i \tag{13}$$

Note that κ_0 in (13) is just a normalization factor ensuring $\kappa^*(\omega) = 0$ for at least one ω. The η_i can be considered as impact factors of the single conditionals $(B_i|A_i)$ for falsifying the conditionals in Δ which have to be chosen so as to ensure success of the revision. Similar as for c-representations, being c-revisions of κ with $\kappa(\omega) = 0$ for all ω (cf. Definition 1), the possible values for the η_i can be specified by a constraint satisfaction problem.

Definition 6 ($CR(\kappa, \Delta)$, $cr^{\Delta}_{\kappa,i}$)**.** Let κ be an OCF and $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ a set of conditionals. The *constraint satisfaction problem for c-revisions of* κ *by* Δ, denoted by $CR(\kappa, \Delta)$, is given by the set of constraints $cr^{\Delta}_{\kappa,i}$, for $i \in \{1, \ldots, n\}$:

$$(cr^{\Delta}_{\kappa,i}) \qquad \begin{aligned} \eta_i > &\min_{\omega \models A_i B_i} \Big\{ \kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \eta_j \Big\} \\ &- \min_{\omega \models A_i \overline{B}_i} \Big\{ \kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \eta_j \Big\} \end{aligned} \tag{14}$$

where the η_i are constraint variables taking values in $\mathbb{N}$.

With $Sol(CR(\kappa, \Delta))$ we denote the set of all solutions $\vec{\eta} = (\eta_1, \ldots, \eta_n)$ of $CR(\kappa, \Delta)$. For any $\vec{\eta} \in \mathbb{N}^n$, the function κ^* as defined by Equation (13) with

$$\kappa_0 = -\min_{\omega \in \Omega} \{ \kappa(\omega) + \sum_{1 \leqslant i \leqslant m,\, \omega \models A_i \overline{B}_i} \eta_i \} \tag{15}$$

is denoted by $\kappa^*_{\vec{\eta}}$.

Proposition 6 (Soundness and completeness of $CR(\kappa, \Delta)$)**.** *Let* κ *be an OCF and* $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ *be a set of conditionals.*

- *If* $\vec{\eta} \in Sol(CR(\kappa, \Delta))$ *then* $\kappa^*_{\vec{\eta}}$ *is a c-revision of* κ *by* Δ *with* $\kappa^*_{\vec{\eta}} \models \Delta$.
- *If* κ^* *is a c-revision of* κ *by* Δ *then there is a vector* $\vec{\eta} \in Sol(CR(\kappa, \Delta))$ *such that* $\kappa^* = \kappa^*_{\vec{\eta}}$.

The proof of Proposition 6 is a direct consequence of the propositions presented in (Kern-Isberner 2004). The constraints given by (14) ensure that a c-revision κ^* of κ by Δ as given by (13) accepts Δ, and κ_0 given by (15) is a normalization factor ensuring that $\kappa^*(\omega) = 0$ for at least one world ω.

Analogously to c-representations, we introduce selection strategies for c-revisions determining a c-revision with respect to the general schema provided by (13) and (14).

Definition 7 (Selection strategy σ, strategic c-revision $*_\sigma$)**.** A *selection strategy (for c-revisions)* is a function

$$\sigma : (\kappa, \Delta) \mapsto \vec{\eta}$$

assigning to each pair of an OCF κ and a (consistent) set of conditionals Δ an impact vector $\vec{\eta} \in Sol(CR(\kappa, \Delta))$. If $\sigma(\kappa, \Delta) = \vec{\eta}$, the *c-revision of* κ *by* Δ *determined by* σ is $\kappa^*_{\vec{\eta}}$, denoted by $\kappa *_\sigma \Delta$, and $*_\sigma$ is a *strategic c-revision.*

Selection strategies for c-revisions generalize selection strategies for c-representations.

Proposition 7. *Every selection strategy* σ *for c-representations uniquely corresponds to a selection strategy* σ_{rev} *for c-revisons via the equation* $\sigma(\Delta) = \sigma_{rev}(\kappa_u, \Delta)$ *for every set of conditionals* Δ.

Note that the c-revision operator $*_\sigma$ determined by a selection strategy σ selects a single c-revision for each κ and each (consistent) Δ. By identifying (plausible) propositional statements A with conditionals having a tautological antecedent, i.e., $A \equiv (A|\top)$, this concept of selection strategy-based c-revision operator $*_\sigma$ can also be applied in a straightforward way to propositional iterated revision and multiple iterated revision, where κ is revised by A resp. a set $\{A_1, \ldots, A_m\}$ of propositions (cf. (Kern-Isberner and Huvermann 2017)).

Also the concept of projections for c-representations developed in Section 3.3 carries over to c-revisions.

Definition 8 (CSP projection $CR(\kappa, \Delta)_{\Delta'}$)**.** Let κ be an OCF, let $\Delta = \{(B_1|A_1), \ldots, (B_m|A_m)\}$ be a set of conditionals, and let $\Delta' \subseteq \Delta$. The *projection of* $CR(\kappa, \Delta)$ *to* Δ', denoted by $CR(\kappa, \Delta)_{\Delta'}$, is the constraint satisfaction problem given by the set of constraints $\{cr^{\Delta}_{\kappa,i} \mid (B_i|A_i) \in \Delta'\}$.

Transforming the axioms for selection strategies for c-representations to c-revisions yields the following axioms.

(SI*) A selection strategy σ is *syntax independent* if for any Δ' obtained from Δ by replacing a conditional $(B|A)$ occurring in Δ by a conditional $(B'|A')$ with $(B|A) \equiv (B'|A')$, we have $\sigma(\kappa, \Delta) = \sigma(\kappa, \Delta')$.

(IP-SOL*) A selection strategy σ is *impact preserving with respect to the solution space* if for any two constraint satisfaction problems $CR(\kappa, \Delta)$, $CR(\kappa', \Delta')$ with $Sol(CR(\kappa, \Delta)) = Sol(CR(\kappa', \Delta'))$, we have $\sigma(\kappa, \Delta) = \sigma(\kappa', \Delta')$.

(IP-ESP*) A selection strategy σ is *impact preserving with respect to equivalent subproblems* if for any two revision problems $(\kappa, \Delta), (\kappa', \Delta')$ with $\Delta_1 \subseteq \Delta, \Delta'_1 \subseteq \Delta'$ and $\Delta_1 \equiv_e \Delta'_1$ such that $CR(\kappa, \Delta)_{\Delta_1} = CR(\kappa', \Delta')_{\Delta'_1}$, we have $\sigma(\kappa, \Delta)_{\Delta_1} = \sigma(\kappa', \Delta')_{\Delta'_1}$.

The following example uses the *marginal* $\kappa|_\Theta$ of an OCF κ on $\Theta \subseteq \Sigma$ given by $\kappa|_\Theta(\omega^\Theta) = \kappa(\omega^\Theta)$ for any $\omega^\Theta \in \Omega(\Theta)$.

Example 2. We continue Example 1. Recently, pollsters have found out that people also like going to the cinema if the weather is not sunny. So, we would like to revise our κ_Δ from Example 1 by $(c|\overline{s})$. When setting up $CR(\kappa_\Delta, \{(c|\overline{s})\})$, we observe that the conditionals from Δ_1, or more generally, the Σ_1-parts of the possible worlds are irrelevant for the constraint problem, and that effectively, we have $CR(\kappa_\Delta, \{(c|\overline{s})\}) = CR((\kappa_\Delta)_{|\Sigma_2}, \{(c|\overline{s})\})$, where $(\kappa_\Delta)_{|\Sigma_2}$ is the marginalisation of κ_Δ on Σ_2. Following (**IP-ESP***), the selection strategy σ chooses $\sigma(\kappa_\Delta, \{(c|\overline{s})\}) = \sigma((\kappa_\Delta)_{|\Sigma_2}, \{(c|\overline{s})\})$. In this way, the impact factors can be found by solving a significantly smaller constraint problem, and the revised OCF $\kappa_\Delta * \{(c|\overline{s})\}$ can be computed via (13) and (15).

The axioms for selection strategies for c-revisions ensure the following properties, generalizing the corresponding observations for c-representations.

Proposition 8. *Let $*_\sigma$ be a strategic c-revision satisfying (**SI**$^\sigma$) and let Δ, Δ' be sets of conditionals. If $\Delta \equiv_e \Delta'$ then $\sigma(\kappa, \Delta) = \sigma(\kappa, \Delta')$ and thus $\kappa *_\sigma \Delta = \kappa *_\sigma \Delta'$ for any κ.*

Proposition 9. *If σ satisfies (**IP-SOL***), σ satisfies (**SI***).*

Proposition 10. *If σ satisfies (**IP-ESP***), σ satisfies (**SI***).*

Proof. Let κ be an OCF, σ a selection strategy, $\Delta_0 = \{(B_1|A_1), \ldots, (B_{m-1}|A_{m-1}\}$ a set of conditionals, and let $(B_m|A_m)$ and $(B'_m|A'_m)$ be two conditionals with $(B_m|A_m) \equiv (B'_m|A'_m)$. For $\Delta = \Delta_0 \cup \Delta_1$ and $\Delta' = \Delta_0 \cup \Delta'_1$ with $\Delta_1 = \{(B_m|A_m)\}$ and $\Delta'_1 = \{(B'_m|A'_m)\}$ we have to show $\sigma(\kappa, \Delta) = \sigma(\kappa, \Delta')$. Let us compare $cr^{\Delta}_{\kappa,i}$ and $cr^{\Delta'}_{\kappa,i}$. According to Equation (14), the only difference in the construction of the two constraints is that whenever $\omega \models A_m B_m$ (or $\omega \models A_m \overline{B_m}$, respectively) is used as a condition in the minimization or in the sum expressions in $cr^{\Delta}_{\kappa,i}$, then $\omega \models A'_m B'_m$ (or $\omega \models A'_m \overline{B'_m}$, respectively) is used in $cr^{\Delta'}_{\kappa,i}$. Thus, because $A_m B_m \equiv A'_m B'_m$ and $A_m \overline{B_m} \equiv A'_m \overline{B'_m}$, we get $cr^{\Delta}_{\kappa,i} = cr^{\Delta'}_{\kappa,i}$ and hence $CR(\kappa, \Delta) = CR(\kappa, \Delta')$. Because $CR(\kappa, \Delta)_{\Delta_0} = CR(\kappa, \Delta')_{\Delta_0}$ and $CR(\kappa, \Delta)_{\Delta_1} = CR(\kappa, \Delta')_{\Delta'_1}$, satisfaction of (**IP-ESP***) by σ implies $\sigma(\kappa, \Delta) = \sigma(\kappa, \Delta')$. $\square$

Another consequence of (**IP-ESP***) is that it ensures that self-fulfilling conditionals do not influence the revision result as expressed by the following postulate (cf. Prop. 5).

(**SF***) If $A \models B$ then $\kappa * (\Delta \cup \{(B|A)\}) = \kappa * \Delta$.

Proposition 11. *If σ satisfies (**IP-ESP***), $*_\sigma$ satisfies (**SF***).*

Because a propositional tautology can be modelled by the self-filfilling conditional $(\top|\top)$, a direct consequence of Proposition 11 is that (**IP-ESP***) ensures *tautological vacuity*, i.e., adding $(\top|\top)$ to a set of conditionals for revison does not have any influence on the result of the revsion.

5 Conclusions and Future Work

We presented selection strategies for inductive reasoning and for belief revision with sets of conditionals. Using axioms for specifying characteristics of such selection strategies, we elaborated various interrelationships among the axioms, and provided constructive approaches for obtaining selection strategies satisfying specific axioms.

C-revisions have been generalized to model other forms of belief change. Regarding contractions, using selection strategies satisfying most sensible axioms similar to the ones presented in this paper invalidate all corresponding counterexamples given in (Haldimann, Kern-Isberner, and Beierle 2020) and show how the required syntax splitting properties can be ensured (Haldimann, Beierle, and Kern-Isberner 2021). Also when modelling Hansson's framework of descriptor revision (Hansson 2014) by employing the principle of conditional preservation as it is done in (Sauerwald et al. 2020), using selection strategies will be beneficial.

References

[Beierle et al. 2018] Beierle, C.; Eichhorn, C.; Kern-Isberner, G.; and Kutsch, S. 2018. Properties of skeptical c-inference for conditional knowledge bases and its realization as a constraint satisfaction problem. *Ann. Math. Artif. Intell.* 83(3-4):247–275.

[Beierle et al. 2021] Beierle, C.; Eichhorn, C.; Kern-Isberner, G.; and Kutsch, S. 2021. Properties and interrelationships of skeptical, weakly skeptical, and credulous inference induced by classes of minimal models. *Artificial Intelligence* 297. (in press, online 2 Mar 2021).

[Darwiche and Pearl 1997] Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artif. Intell.* 89:1–29.

[de Finetti 1937] de Finetti, B. 1937. La prévision, ses lois logiques et ses sources subjectives. *Ann. Inst. H. Poincaré* 7(1):1–68. Engl. transl. *Theory of Probability*, J. Wiley & Sons, 1974.

[Haldimann, Beierle, and Kern-Isberner 2021] Haldimann, J.; Beierle, C.; and Kern-Isberner, G. 2021. Syntax splitting for iterated contractions, ignorations, and revisions on ranking functions using selection strategies. In *JELIA-2021*, LNAI. Springer.

[Haldimann, Kern-Isberner, and Beierle 2020] Haldimann, J.; Kern-Isberner, G.; and Beierle, C. 2020. Syntax splitting for iterated contractions. In *KR-2020*, 465–475.

[Hansson 2014] Hansson, S. O. 2014. Descriptor revision. *Studia Logica* 102(5):955–980.

[Kern-Isberner and Huvermann 2017] Kern-Isberner, G., and Huvermann, D. 2017. What kind of independence do we need for multiple iterated belief change? *J. Applied Logic* 22:91–119.

[Kern-Isberner, Beierle, and Brewka 2020] Kern-Isberner, G.; Beierle, C.; and Brewka, G. 2020. Syntax splitting = relevance + independence: New postulates for nonmonotonic reasoning from conditional belief bases. In *KR-2020*, 560–571.

[Kern-Isberner 2001] Kern-Isberner, G. 2001. *Conditionals in nonmonotonic reasoning and belief revision*. Springer, Lecture Notes in Artificial Intelligence LNAI 2087.

[Kern-Isberner 2004] Kern-Isberner, G. 2004. A thorough axiomatization of a principle of conditional preservation in belief revision. *Ann. Math. Artif. Intell.* 40(1-2):127–164.

[Kern-Isberner 2018] Kern-Isberner, G. 2018. Axiomatizing a qualitative principle of conditional preservation for iterated belief change. In *KR-2018*, 248–256. AAAI Press.

[Lehmann and Magidor 1992] Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.

[Makinson 1989] Makinson, D. 1989. General theory of cumulative inference. In Reinfrank, M., et al., eds., *Non-monotonic Reasoning*, LNAI, Vol. 346. Springer. 1–18.

[Parikh 1999] Parikh, R. 1999. Beliefs, belief revision, and splitting languages. *Logic, language, and computation* 2:266–278.

[Pearl 1990] Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proc. TARK'1990*, 121–135. Morgan Kaufmann.

[Sauerwald et al. 2020] Sauerwald, K.; Haldimann, J.; von Berg, M.; and Beierle, C. 2020. Descriptor revision for conditionals: Literal descriptors and conditional preservation. In *KI-2020*, volume 12325 of *LNCS*, 204–218. Springer.

[Spohn 1988] Spohn, W. 1988. Ordinal conditional functions: a dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics, II*. Kluwer Academic Publishers. 105–134.