Conditional indifference and conditional preservation

Gabriele Kern-Isberner
FernUniversitaet Hagen, Department of Computer Science
P.O. Box 940, D-58084 Hagen, Germany
gabriele.kern-isberner@fernuni-hagen.de

Abstract

The idea of preserving conditional beliefs emerged recently as a new paradigm apt to guide the revision of epistemic states. Conditionals are substantially different from propositional beliefs and need specific treatment. In this paper, we present a new approach to conditionals, capturing particularly well their dynamic part as revision policies. We thoroughly axiomatize a *principle of conditional preservation* as an indifference property with respect to *conditional structures* of worlds. This principle is developed in a semi-quantitative setting, so as to reveal its fundamental meaning for belief revision in quantitative as well as in qualitative frameworks. In fact, it is shown to cover other proposed approaches to conditional preservation.

Introduction

Within the last years, the propositional limitations of classical belief revision have been overcome piece by piece. For instance, Boutilier (Boutilier 1994) investigated belief revision within a modal framework, and Williams (Williams 1994) proposed *transmutation schemas for knowledge systems*. In general, *epistemic states* have moved into the center of interest as representations of belief states of some individual or intelligent agent at a given time. Besides propositional or first-order facts, reflecting certain knowledge, they may contain assumptions, preferences among beliefs, and, as a crucial ingredient, conditional knowledge. They may be represented in different ways, e.g. by gradings of plausibility or disbelief, by making use of epistemic entrenchment, or as a probability distribution.

Epistemic states provide an excellent framework to study *iterative revisions* which are important to build fully dynamic systems[1]. While propositional AGM theory only observes the results of revisions, considering epistemic states under change allows one to focus on the mechanisms underlying that change, taking conditional beliefs as *revision policies* explicitly into account (cf. (Boutilier & Goldszmidt 1993; Rott 1991; Darwiche & Pearl 1997)). The connection between epistemic states, Ψ, (iterative) revision operators, $*$, and conditionals, $(B|A)$, is established by the *Ramsey test*

$$\Psi \models (B|A) \quad \text{iff} \quad \Psi * A \models B. \tag{1}$$

Hence revising epistemic states does not only mean to deal with propositional beliefs – it also requires studying how conditional beliefs are changed.

Darwiche and Pearl (Darwiche & Pearl 1997) rephrased the AGM postulates for epistemic states. Applying the *minimal change paradigm* of propositional belief revision in that extended framework, as Goldszmidt and Boutilier did in (Boutilier & Goldszmidt 1993), however, may produce unintuitive results (Darwiche & Pearl 1997). So, Darwiche and Pearl (Darwiche & Pearl 1997) advanced four postulates as a cautious approach to describe *principles of conditional preservation* when revising epistemic states by *propositional* beliefs. In (Kern-Isberner 1999c; 1999b), we extended their approach in considering revisions of epistemic states by *conditionals*. We proposed a set of axioms outlining conditional revisions which are in accordance with some fundamental postulates of revisions (like, for instance, *success*) and with propositional AGM theory, and which preserve conditional beliefs in observing conditional interactions. These interactions were specified by two newly introduced relations between conditionals, that of *subconditionality* and that of *perpendicularity* (see (Kern-Isberner 1999b)).

Earlier, in (Kern-Isberner 1998), we defined a *principle of conditional preservation* in quite a different, namely probabilistic, framework. This principle there was based on the algebraic notion of conditional structures, made use of group theoretical means and postulated the numerical values of the given probability distribution to follow the *conditional structures* of worlds.

In this paper, we will bring together both approaches to conditional preservation, the qualitative one and the quantitative one, in the semi-quantitative framework of ordinal conditional functions. We will rephrase the probabilistic principle of conditional preservation of (Kern-Isberner 1998) for ordinal conditional functions, and we will show, that this quantitative principle of conditional preservation implies the corresponding

[1] An interesting approach to iterated revisions of *belief sets* was proposed quite recently by Lehmann, Magidor and Schlechta in (Lehmann, Magidor, & Schlechta 1999).

qualitative postulates of (Kern-Isberner 1999b). Actually, the quantitative version is much stronger than the qualitative one, dealing with *sets of conditionals* instead of only one revising conditional, and handling interactions of conditionals of arbitrary complexity. Although numerical in nature, the principle is based on a symbolic representation of conditional influences on worlds, called *conditional structures*. Numbers (rankings or probabilities) are only considered as manifestations of complex conditional interactions. The representation of conditional structures by the aid of group theory provides a rich methodological framework to study conditionals in belief revision and nonmonotonic reasoning.

In the following section, we fix some notations and describe the relationship between conditionals and epistemic states. Then we introduce the crucial notion of conditional structures which the property of conditional indifference is based upon. By applying the concept of indifference to revision functions, we obtain a precise formalization of the principle of conditional preservation, and we characterize ordinal conditional functions observing this principle. Further on, we compare our approach to Goldszmidt, Morris & Pearl's system-Z and system-Z^* in several examples. Finally, we bring together qualitative and quantitative approaches to the principle of conditional preservation, proving the formalization given here to be a most fundamental one. A summary and an outlook conclude this paper.

Conditionals and epistemic states

We consider a propositional language $\mathcal{L}$ over a finite alphabet $a, b, c \ldots$. Let Ω denote the set of possible worlds for $\mathcal{L}$, i.e. Ω is a complete set of interpretations of $\mathcal{L}$. Throughout this paper, we will write $\overline{A}$ instead of $\neg A$, and AB instead of $A \wedge B$, for formulas $A, B \in \mathcal{L}$.

Conditionals $(B|A)$ represent statements of the form "*If A then B*", expressing a relationship between two (propositional) formulas A, the *antecedent* or *premise*, and B, the *consequent*. $(\mathcal{L} \mid \mathcal{L})$ denotes the set of all conditionals $(B|A)$ with $A, B \in \mathcal{L}$. A conditional $(B|\top)$ with a tautological antecedent is taken to correspond to its (propositional) consequent, B. $(D|C)$ is called a *subconditional* of $(B|A)$, written as

$$(D|C) \sqsubseteq (B|A), \tag{2}$$

iff $CD \models AB$ and $C\overline{D} \models A\overline{B}$. Typically, subconditionals arise by strengthening the antecedent of a conditional, e.g. $(b|ac)$ is a subconditional of $(b|a)$, $(b|ac) \sqsubseteq (b|a)$.

Each possible world $\omega \in \Omega$ either confirms $(B|A)$, in case that $\omega \models AB$, or refutes it, if $\omega \models A\overline{B}$, or does not even satisfy its premise, $\omega \not\models A$, and so is of no relevance for it. Although conditionals are evaluated with respect to worlds, they cannot really be accepted (as entities) in single, isolated worlds. To validate conditionals, we need richer epistemic structures than plain propositional interpretations, at least to compare different worlds with regard to their relevance for a conditional (see, for example, (Nute 1980; Boutilier 1994; Darwiche & Pearl 1997)). Epistemic states as representations of cognitive states of intelligent agents provide an adequate framework for conditionals.

An epistemic notion that turned out to be of great importance both for conditionals and epistemic states, in particular in the context of belief revision, is that of *plausibility*: conditionals are supposed to represent plausible conclusions, and plausibility relations on formulas or worlds, respectively, guide AGM-revisions of belief sets and of epistemic states (Nute 1980; Katsuno & Mendelzon 1991; Darwiche & Pearl 1997).

As Spohn (Spohn 1988) emphasized, however, it is not enough to consider the qualitative ordering of propositions according to their plausibility – also relative distances between degrees of plausibility should be taken into account. So he introduced ordinal conditional functions κ (*OCF's, ranking functions*) (Spohn 1988) from worlds to ordinals such that some worlds are mapped to the minimal element 0. Here, we will simply assume that OCF's are functions $\kappa : \Omega \rightarrow \mathbb{N} \cup \{0, \infty\}$ from the set of worlds to the natural numbers, extended by 0 and ∞. They specify non-negative integers as degrees of plausibility – or, more precisely, as degrees of *disbelief* – for worlds. The smaller $\kappa(\omega)$ is, the more plausible the world ω appears, and what is believed (for certain) in the epistemic state represented by κ is described precisely by the set $Mod(\kappa) := \{\omega \in \Omega \mid \kappa(\omega) = 0\}$. For propositional formulas $A, B \in \mathcal{L}$, we set $\kappa(A) = \min\{\kappa(\omega) \mid \omega \models A\}$, so that $\kappa(A \vee B) = \min\{\kappa(A), \kappa(B)\}$. In particular, $0 = \min\{\kappa(A), \kappa(\overline{A})\}$, so that at least one of A or $\overline{A}$ is considered mostly plausible. A proposition A is believed iff $\kappa(\overline{A}) > 0$, which is denoted by $\kappa \models A$. A conditional $(B|A) \in (\mathcal{L} \mid \mathcal{L})$ may be assigned a degree of plausibility via $\kappa(B|A) = \kappa(AB) - \kappa(A)$. Each OCF κ induces a (propositional) AGM-revision operator $*$ by setting $Mod(\kappa * A) = \min_\kappa(Mod(A))$ (see (Darwiche & Pearl 1997)). The Ramsey test (1) then reads $\kappa \models (B|A)$ iff $\kappa * A \models B$. This is in accordance with the plausibility relation imposed by κ, as the following lemma shows:

Lemma 1 *Let $(B|A)$ be a conditional in $(\mathcal{L} \mid \mathcal{L})$, let κ be an ordinal conditional function. Then $\kappa \models (B|A)$ (by applying the Ramsey test) iff $\kappa(AB) < \kappa(A\overline{B})$.*

So κ accepts a conditional (via the Ramsey test) iff AB is more plausible than $A\overline{B}$. The proof of this lemma is immediate.

Conditional structures

By observing the behavior of worlds with respect to it, each conditional $(B|A)$ can be considered as a generalized (namely three-valued) indicator function on worlds:

$$(B|A)(\omega) = \begin{cases} 1 & : \ \omega \models AB \\ 0 & : \ \omega \models A\overline{B} \\ u & : \ \omega \models \overline{A} \end{cases} \tag{3}$$

where u stands for *undefined* (DeFinetti 1974; Calabrese 1991)). Intuitively, incorporating a conditional as a plausible conclusion in an epistemic state means to make – at least some – worlds confirming the conditional more plausible than the worlds refuting it. In this sense, conditionals to be learned have effects on possible worlds (more exactly, on their degrees of plausibility), shifting them appropriately to establish the intended plausible relationship. (3) then provides a classification of worlds for achieving this: On confirming worlds $\omega \models AB$, i.e. $(B|A)(\omega) = 1$, $(B|A)$ possibly has a positive effect, while on refuting worlds $\omega \models A\overline{B}$, $(B|A)$ possibly has a negative effect; the effects on worlds ω with $(B|A)(\omega) = u$ is unclear. Which worlds will actually be shifted depends on the chosen revision procedure – for the conditional, all worlds in either of the partitioning sets are indistinguishable.

When we consider (finite) sets of conditionals $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$, we have to modify the representation (3) appropriately to identify the effect of each conditional in $\mathcal{R}$ on worlds in Ω. This leads to introducing the functions $\sigma_i = \sigma_{(B_i|A_i)}$ below (see (4)) which generalize (3) by replacing the numbers 0 and 1 by abstract symbols. Moreover, we will make use of a group structure to represent the joint impact of conditionals on worlds.

To each conditional $(B_i|A_i)$ in $\mathcal{R}$ we associate two symbols $\mathbf{a}_i^+, \mathbf{a}_i^-$. Let

$$\mathcal{F}_\mathcal{R} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$$

be the free abelian group with generators $\mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^-$, i.e. $\mathcal{F}_\mathcal{R}$ consists of all elements of the form $(\mathbf{a}_1^+)^{r_1}(\mathbf{a}_1^-)^{s_1} \ldots (\mathbf{a}_n^+)^{r_n}(\mathbf{a}_n^-)^{s_n}$ with integers $r_i, s_i \in \mathbb{Z}$ (the ring of integers). Each element of $\mathcal{F}_\mathcal{R}$ can be identified by its exponents, so that $\mathcal{F}_\mathcal{R}$ is isomorphic to $\mathbb{Z}^{2n}$ (Lyndon & Schupp 1977). The commutativity of $\mathcal{F}_\mathcal{R}$ corresponds to the fact that the conditionals in $\mathcal{R}$ shall be effective simultaneously, without assuming any order of application. So our way of dealing with conditionals is a symmetric, homogeneous one – we do not need (user-defined) priorities among conditionals. Note that, although we will speak of *multiplication* and *products* in $\mathcal{F}_\mathcal{R}$, the generators of $\mathcal{F}_\mathcal{R}$ are merely juxtaposed, like words.

For each $i, 1 \leqslant i \leqslant n$, we define a function $\sigma_i : \Omega \to \mathcal{F}_\mathcal{R}$ by setting

$$\sigma_i(\omega) = \begin{cases} \mathbf{a}_i^+ & \text{if} \quad (B_i|A_i)(\omega) = 1 \\ \mathbf{a}_i^- & \text{if} \quad (B_i|A_i)(\omega) = 0 \\ 1 & \text{if} \quad (B_i|A_i)(\omega) = u \end{cases} \qquad (4)$$

$\sigma_i(\omega)$ represents the manner in which the conditional $(B_i|A_i)$ applies to the possible world ω. The neutral element 1 of $\mathcal{F}_\mathcal{R}$ corresponds to the non-applicability of $(B_i|A_i)$ in case that the antecedent A_i is not satisfied. The function $\sigma_\mathcal{R} : \Omega \to \mathcal{F}_\mathcal{R}$,

$$\sigma_\mathcal{R}(\omega) = \prod_{1 \leqslant i \leqslant n} \sigma_i(\omega) = \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \mathbf{a}_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} \mathbf{a}_i^-$$

describes the all-over effect of $\mathcal{R}$ on ω. $\sigma_\mathcal{R}(\omega)$ is called *(a representation of) the conditional structure of* ω *with respect to* $\mathcal{R}$. For each world ω, $\sigma_\mathcal{R}(\omega)$ contains at most one of each $\mathbf{a}_i^+$ or $\mathbf{a}_i^-$, but never both of them because each conditional applies to ω in a well-defined way. The next lemma (which is easy to prove) shows that this property characterizes conditional structure functions:

Lemma 2 *Let* $\sigma : \Omega \to \mathcal{F}$ *be a map from the set of worlds* Ω *to the free abelian group* $\mathcal{F} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$ *generated by* $\mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^-$, *such that* $\sigma(\omega)$ *contains at most one of each* $\mathbf{a}_i^+$ *or* $\mathbf{a}_i^-$, *for each world* $\omega \in \Omega$. *Then there is a set of conditionals* $\mathcal{R}$ *with* $card(\mathcal{R}) \leqslant n$ *such that* $\sigma = \sigma_\mathcal{R}$.

Example 3 Let $\mathcal{R} = \{(c|a), (c|b)\}$, where a, b, c are atoms, and let $\mathcal{F}_\mathcal{R} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \mathbf{a}_2^-, \mathbf{a}_2^- \rangle$. We associate $\mathbf{a}_1^\pm$ with the first conditional, $(c|a)$, and $\mathbf{a}_2^\pm$ with the second one, $(c|b)$. The following table shows the values of the function $\sigma_\mathcal{R}$ on worlds $\omega \in \Omega$:

ω	$\sigma_\mathcal{R}(\omega)$	ω	$\sigma_\mathcal{R}(\omega)$
abc	$\mathbf{a}_1^+\mathbf{a}_2^+$	$\overline{a}bc$	$\mathbf{a}_2^+$
$ab\overline{c}$	$\mathbf{a}_1^-\mathbf{a}_2^-$	$\overline{a}b\overline{c}$	$\mathbf{a}_2^-$
$a\overline{b}c$	$\mathbf{a}_1^+$	$\overline{a}\overline{b}c$	1
$a\overline{b}\overline{c}$	$\mathbf{a}_1^-$	$\overline{a}\overline{b}\overline{c}$	1

abc confirms both conditionals, so its conditional structure is represented by $\mathbf{a}_1^+\mathbf{a}_2^+$. This corresponds to the product (in $\mathcal{F}_\mathcal{R}$) of the conditional structures of the worlds $\overline{a}bc$ and $a\overline{b}c$. Two worlds, namely $\overline{a}\overline{b}c$ and $\overline{a}\overline{b}\overline{c}$, are not affected at all by the conditionals in $\mathcal{R}$. ■

The logical structure of antecedents and consequents of the conditionals in $\mathcal{R}$ does not really matter, nor do logical relationships between the conditionals. All that we need is a conditional's partitioning property on the set of worlds (cf. (3) and (4)). $\sigma_\mathcal{R}$ labels each world appropriately and allows us to compare different worlds with respect to the impact the conditionals in $\mathcal{R}$ exert on them. The following example illustrates that also multiple copies of worlds may be necessary to relate conditional structures:

Example 4 Consider the set $\mathcal{R} = \{(d|a), (d|b), (d|c)\}$ of conditionals using the atoms a, b, c, d. Let $\mathbf{a}_1^\pm, \mathbf{a}_2^\pm, \mathbf{a}_3^\pm$ be the group generators associated with $(d|a), (d|b)$, $(d|c)$, respectively. Then we have

$$\begin{aligned} \sigma_\mathcal{R}(ab\overline{c}d)\sigma_\mathcal{R}(a\overline{b}cd)\sigma_\mathcal{R}(\overline{a}bcd) &= (\mathbf{a}_1^+\mathbf{a}_2^+)(\mathbf{a}_1^+\mathbf{a}_3^+)(\mathbf{a}_2^+\mathbf{a}_3^+) \\ = (\mathbf{a}_1^+)^2(\mathbf{a}_2^+)^2(\mathbf{a}_3^+)^2 &= (\mathbf{a}_1^+\mathbf{a}_2^+\mathbf{a}_3^+)^2 \\ &= \sigma_\mathcal{R}(abcd)^2. \end{aligned}$$

Here two copies of $abcd$, or of its structure, respectively, are necessary to match the product of the conditional structures of $ab\overline{c}d, a\overline{b}cd$ and $\overline{a}bcd$. ■

To compare worlds adequately with respect to their conditional structures, we take the worlds $\omega \in \Omega$ as formal generators of the free abelian group

$$\widehat{\Omega} := \langle \omega \mid \omega \in \Omega \rangle$$

$\widehat{\Omega}$ consists of all products $\widehat{\omega} = {\omega_1}^{r_1} \ldots {\omega_m}^{r_m}$, with $\omega_1, \ldots, \omega_m \in \Omega$, and $r_1, \ldots r_m$ integers. Introducing such a "multiplication between worlds" is nothing but a technical means to comply with the multiplicative structure the effects of conditionals impose on worlds. As in $\mathcal{F}_\mathcal{R}$, multiplication in $\widehat{\Omega}$ actually means juxtaposition. In (Kern-Isberner 1998), where we first developed these ideas, we considered multi-sets of worlds (corresponding to elements in $\widehat{\Omega}$ with only positive exponents) and calculated the conditional structure of such a multi-set as the *conditional weight* it is carrying. Making use of arbitrary elements of $\widehat{\Omega}$ as group elements, however, provides a much more convenient and elegant framework to deal with conditional structures. We will usually write $\frac{\omega_1}{\omega_2}$ instead of $\omega_1 \omega_2^{-1}$.

Now $\sigma_\mathcal{R}$ may be extended to $\widehat{\Omega}$ in a straightforward manner by setting

$$\sigma_\mathcal{R}(\widehat{\omega}) = \sigma_\mathcal{R}(\omega_1)^{r_1} \ldots \sigma_\mathcal{R}(\omega_m)^{r_m},$$

yielding a *homomorphism of groups* $\sigma_\mathcal{R} : \widehat{\Omega} \to \mathcal{F}_\mathcal{R}$. For $\widehat{\omega} = {\omega_1}^{r_1} \ldots {\omega_m}^{r_m} \in \widehat{\Omega}$, we obtain

$$\sigma_\mathcal{R}({\omega_1}^{r_1} \ldots {\omega_m}^{r_m}) = \prod_{1 \leqslant i \leqslant n} (\mathbf{a}_i^+)^{\sum_{k : \sigma_i(\omega_k) = \mathbf{a}_i^+} r_k} \prod_{1 \leqslant i \leqslant n} (\mathbf{a}_i^-)^{\sum_{k : \sigma_i(\omega_k) = \mathbf{a}_i^-} r_k},$$

as a representation of its conditional structure. The exponent of $\mathbf{a}_i^+$ in $\sigma_\mathcal{R}(\widehat{\omega})$ indicates the number of worlds in $\widehat{\omega}$ which confirm the conditional $(B_i|A_i)$, each world being counted with its multiplicity, and in the same way, the exponent of $\mathbf{a}_i^-$ indicates the number of worlds that are in conflict with $(B_i|A_i)$.

By investigating suitable elements of $\widehat{\Omega}$, it is possible to isolate the (positive or negative) net impacts of conditionals in $\mathcal{R}$, as the following example illustrates:

Example 5 (continued) In Example 3 above, we have

$$\sigma_\mathcal{R}\left(\frac{abc}{\overline{a}bc}\right) = \frac{\mathbf{a}_1^+ \mathbf{a}_2^+}{\mathbf{a}_2^+} = \mathbf{a}_1^+$$

So $\frac{abc}{\overline{a}bc}$ reveals the positive net impact of the conditional $(c|a)$ within $\mathcal{R}$, symbolized by $\mathbf{a}_1^+$.

Similarly, in Example 4, the element $\frac{ab\overline{c}\overline{d} \cdot \overline{a}\overline{b}c\overline{d}}{a\overline{b}c\overline{d}}$ isolates the negative net impact of the second conditional, $(d|b)$:

$$\sigma_\mathcal{R}\left(\frac{ab\overline{c}\overline{d} \cdot \overline{a}\overline{b}c\overline{d}}{a\overline{b}c\overline{d}}\right) = \frac{\mathbf{a}_1^- \mathbf{a}_2^- \cdot \mathbf{a}_3^-}{\mathbf{a}_1^- \mathbf{a}_3^-} = \mathbf{a}_2^- . \qquad \blacksquare$$

The following example is taken from (Goldszmidt & Pearl 1996, p. 68f):

Example 6 Consider the set $\mathcal{R}$ consisting of the following conditionals:

$r_1 : (f|b)$ *Birds fly.*
$r_2 : (b|p)$ *Penguins are birds.*
$r_3 : (\overline{f}|p)$ *Penguins do not fly.*
$r_4 : (w|b)$ *Birds have wings.*
$r_5 : (a|f)$ *Animals that fly are airborne.*

In Table 1, we list the conditional structures of all

ω	$\sigma_\mathcal{R}(\omega)$	ω	$\sigma_\mathcal{R}(\omega)$
$pbfwa$	$\mathbf{a}_1^+ \mathbf{a}_2^+ \mathbf{a}_3^- \mathbf{a}_4^+ \mathbf{a}_5^+$	$\overline{p}bfwa$	$\mathbf{a}_1^+ \mathbf{a}_4^+ \mathbf{a}_5^+$
$pbfw\overline{a}$	$\mathbf{a}_1^+ \mathbf{a}_2^+ \mathbf{a}_3^- \mathbf{a}_4^+ \mathbf{a}_5^-$	$\overline{p}bfw\overline{a}$	$\mathbf{a}_1^+ \mathbf{a}_4^+ \mathbf{a}_5^-$
$pbf\overline{w}a$	$\mathbf{a}_1^+ \mathbf{a}_2^+ \mathbf{a}_3^- \mathbf{a}_4^- \mathbf{a}_5^+$	$\overline{p}bf\overline{w}a$	$\mathbf{a}_1^+ \mathbf{a}_4^- \mathbf{a}_5^+$
$pbf\overline{w}\,\overline{a}$	$\mathbf{a}_1^+ \mathbf{a}_2^+ \mathbf{a}_3^- \mathbf{a}_4^- \mathbf{a}_5^-$	$\overline{p}bf\overline{w}\,\overline{a}$	$\mathbf{a}_1^+ \mathbf{a}_4^- \mathbf{a}_5^-$
$pb\overline{f}wa$	$\mathbf{a}_1^- \mathbf{a}_2^+ \mathbf{a}_3^+ \mathbf{a}_4^+$	$\overline{p}b\overline{f}wa$	$\mathbf{a}_1^- \mathbf{a}_4^+$
$pb\overline{f}w\overline{a}$	$\mathbf{a}_1^- \mathbf{a}_2^+ \mathbf{a}_3^+ \mathbf{a}_4^+$	$\overline{p}b\overline{f}w\overline{a}$	$\mathbf{a}_1^- \mathbf{a}_4^+$
$pb\overline{f}\overline{w}a$	$\mathbf{a}_1^- \mathbf{a}_2^+ \mathbf{a}_3^+ \mathbf{a}_4^-$	$\overline{p}b\overline{f}\overline{w}a$	$\mathbf{a}_1^- \mathbf{a}_4^-$
$pb\overline{f}\overline{w}\,\overline{a}$	$\mathbf{a}_1^- \mathbf{a}_2^+ \mathbf{a}_3^+ \mathbf{a}_4^-$	$\overline{p}b\overline{f}\overline{w}\,\overline{a}$	$\mathbf{a}_1^- \mathbf{a}_4^-$
$p\overline{b}fwa$	$\mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_5^+$	$\overline{p}\overline{b}fwa$	$\mathbf{a}_5^+$
$p\overline{b}fw\overline{a}$	$\mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_5^-$	$\overline{p}\overline{b}fw\overline{a}$	$\mathbf{a}_5^-$
$p\overline{b}f\overline{w}a$	$\mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_5^+$	$\overline{p}\overline{b}f\overline{w}a$	$\mathbf{a}_5^+$
$p\overline{b}f\overline{w}\,\overline{a}$	$\mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_5^-$	$\overline{p}\overline{b}f\overline{w}\,\overline{a}$	$\mathbf{a}_5^-$
$p\overline{b}\,\overline{f}wa$	$\mathbf{a}_2^- \mathbf{a}_3^+$	$\overline{p}\overline{b}\,\overline{f}wa$	1
$p\overline{b}\,\overline{f}w\overline{a}$	$\mathbf{a}_2^- \mathbf{a}_3^+$	$\overline{p}\overline{b}\,\overline{f}w\overline{a}$	1
$p\overline{b}\,\overline{f}\overline{w}a$	$\mathbf{a}_2^- \mathbf{a}_3^+$	$\overline{p}\overline{b}\,\overline{f}\overline{w}a$	1
$p\overline{b}\,\overline{f}\overline{w}\,\overline{a}$	$\mathbf{a}_2^- \mathbf{a}_3^+$	$\overline{p}\overline{b}\,\overline{f}\overline{w}\,\overline{a}$	1

Table 1: Conditional structures for Example 6

possible worlds; this table will be helpful in the sequel. $\blacksquare$

Having the same conditional structure defines an equivalence relation $\equiv_\mathcal{R}$ on $\widehat{\Omega}$:

$$\widehat{\omega}_1 \equiv_\mathcal{R} \widehat{\omega}_2 \quad \text{iff} \quad \sigma_\mathcal{R}(\widehat{\omega}_1) = \sigma_\mathcal{R}(\widehat{\omega}_2). \tag{5}$$

Those elements of $\widehat{\Omega}$ that are balanced with respect to the effects of conditionals in $\mathcal{R}$ are contained in the *kernel of* $\sigma_\mathcal{R}$, *ker* $\sigma_\mathcal{R} = \{\widehat{\omega} \in \widehat{\Omega} \mid \sigma_\mathcal{R}(\widehat{\omega}) = 1\}$. *ker* $\sigma_\mathcal{R}$ does not depend on the chosen representation of conditional structures by symbols in $\mathcal{F}_\mathcal{R}$ and thus, it is an invariant of $\mathcal{R}$ (Kern-Isberner 1999d).

Often, besides the conditionals explicitly given in $\mathcal{R}$, implicit normalizing constraints have to be taken into account, like, e.g. $\kappa(\top) = 0$ for ordinal conditional functions. This can be achieved by focusing on equivalence with repect to $\sigma_\top$. Since $\sigma_\top$ simply counts the generators occurring in $\widehat{\omega}$, two elements $\widehat{\omega}_1 = \omega_1^{r_1} \ldots \omega_m^{r_m}$, $\widehat{\omega}_2 = \nu_1^{s_1} \ldots \nu_p^{s_p} \in \widehat{\Omega}$ are $\sigma_\top$-equivalent, $\widehat{\omega}_1 \equiv_\top \widehat{\omega}_2$, iff $\sum_{1 \leqslant j \leqslant m} r_j = \sum_{1 \leqslant k \leqslant p} s_k$. This means, $\widehat{\omega}_1 \equiv_\top \widehat{\omega}_2$ iff they both are a (cancelled) product of the same number of generators, each generator being counted with its corresponding exponent.

Conditional indifference

To study conditional interactions, we now focus on the behavior of OCF's $\kappa : \Omega \to \mathbb{N} \cup \{0, \infty\}$ with respect to the multiplication in $\widehat{\Omega}$. Each such function may be extended to a homomorphism, $\kappa : \widehat{\Omega}_+ \to (\mathbb{Z}, +)$, by setting

$$\kappa({\omega_1}^{r_1} \cdot \ldots \cdot {\omega_m}^{r_m}) = r_1 \kappa(\omega_1) + \ldots + r_m \kappa(\omega_m),$$

where $\widehat{\Omega}_+$ is the subgroup of $\widehat{\Omega}$ generated by the set $\Omega_+ := \{\omega \in \Omega \mid \kappa(\omega) \neq \infty\}$. This allows us to analyze numerical relationships holding between different $\kappa(\omega)$. Thereby, it will be possible to elaborate the conditionals whose structures κ follows, that means, to determine sets of conditionals $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ with respect to which κ is indifferent:

Definition 7 Suppose $\kappa : \Omega \to \mathbb{N} \cup \{0, \infty\}$ is an OCF, and $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ is a set of conditionals such that $\kappa(A) \neq \infty$ for all $(B|A) \in \mathcal{R}$. κ is *indifferent with respect to* $\mathcal{R}$ iff the following two conditions hold:

(i) If $\kappa(\omega) = \infty$ then there is $(B|A) \in \mathcal{R}$ such that $\sigma_{(B|A)}(\omega) \neq 1$ and $\kappa(\omega') = \infty$ for all ω' with $\sigma_{(B|A)}(\omega') = \sigma_{(B|A)}(\omega)$.

(ii) $\kappa(\widehat{\omega}_1) = \kappa(\widehat{\omega}_2)$ whenever $\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \sigma_{\mathcal{R}}(\widehat{\omega}_2)$ for $\widehat{\omega}_1 \equiv_{\top} \widehat{\omega}_2 \in \widehat{\Omega}_+$.

If κ is indifferent with respect to $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$, then it does not distinguish between different elements $\widehat{\omega}_1, \widehat{\omega}_2$ with the same conditional structure with respect to $\mathcal{R}$. Normalizing constraints are taken into account by observing $\equiv_{\top}$-equivalence. Conversely, any deviation $\kappa(\widehat{\omega}) \neq 0$ can be explained by the conditionals in $\mathcal{R}$ acting on $\widehat{\omega}$ in a non-balanced way. Condition (i) in Definition 7 is necessary to deal with worlds $\omega \notin \Omega_+$. Conditional indifference, as defined in Definition 7, captures interactions of conditionals of arbitrary depth by making use of the homomorphism induced by κ. It also respects, however, indifference on the superficial level of the function κ itself:

Lemma 8 *If the ordinal conditional function κ is indifferent with respect to $\mathcal{R}$, then $\sigma_{\mathcal{R}}(\omega_1) = \sigma_{\mathcal{R}}(\omega_2)$ implies $\kappa(\omega_1) = \kappa(\omega_2)$ for all worlds $\omega_1, \omega_2 \in \Omega$.*

The next theorem gives a simple criteria to check conditional indifference with ordinal conditional functions. Moreover, it provides an intelligible schema to construct conditional indifferent functions.

Theorem 9 *An OCF κ is indifferent with respect to a set $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ of conditionals iff $\kappa(A_i) \neq \infty$ for all $i, 1 \leqslant i \leqslant n$, and there are rational numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}$, $1 \leqslant i \leqslant n$, such that for all $\omega \in \Omega$,*

$$\kappa(\omega) = \kappa_0 + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^- \tag{6}$$

Sketch of proof. According to Lemma 8, the equivalence relation (5) provides a rough classification of the worlds in Ω with respect to the conditionals in $\mathcal{R}$. Obtaining a representation of the form (6) then amounts to checking the solvability of a linear equational system. The proof of this theorem is very similar to the proof of the analogous theorem for probabilistic representation of knowledge given in (Kern-Isberner 1998). □

The principle of conditional preservation

Minimality of change is a crucial paradigm for belief revision, and a "principle of conditional preservation" is to realize this idea of minimality when conditionals are involved in change. Minimizing absolutely the changes in conditional beliefs, as in (Boutilier & Goldszmidt 1993), is an important proposal to this aim, but it does not always lead to intuitive results (Darwiche & Pearl 1997). The idea we will develop here rather aims at *preserving the conditional structure of knowledge* within an epistemic state which we assume to be represented by an OCF κ.

We just explained what it means for an OCF κ to follow the structure imposed by $\mathcal{R}$ on the set of worlds by introducing the notion of conditional indifference (cf. Definition 7). Pursuing this approach further in the framework of belief revision, a revision of κ by simultaneously incorporating the conditionals in $\mathcal{R}$, $\kappa^* = \kappa * \mathcal{R}$, can be said to preserve the conditional structure of κ with respect to $\mathcal{R}$ if the *relative change function* $\kappa^* - \kappa$ is indifferent with respect to $\mathcal{R}$[2]. Taking into regard the worlds ω with $\kappa(\omega) = \infty$ appropriately, this gives rise to the following definitions:

Definition 10 Let κ be an OCF, and let $\mathcal{R}$ be a finite set of conditionals. Let $\kappa^* = \kappa * \mathcal{R}$ denote the result of revising κ by $\mathcal{R}$. Presuppose further[3] that $\kappa^*(A) \neq \infty$ for all $(B|A) \in \mathcal{R}$.

1. κ^* is called *κ-consistent* iff $\kappa(\omega) = \infty$ implies $\kappa^*(\omega) = \infty$.
2. κ^* is *indifferent with respect to $\mathcal{R}$ and κ* iff κ^* is κ-consistent and the following two conditions hold:
   (i) If $\kappa^*(\omega) = \infty$ then $\kappa(\omega) = \infty$, or there is $(B|A) \in \mathcal{R}$ such that $\sigma_{(B|A)}(\omega) \neq 1$ and $\kappa^*(\omega') = \infty$ for all ω' with $\sigma_{(B|A)}(\omega') = \sigma_{(B|A)}(\omega)$.
   (ii) $(\kappa^* - \kappa)(\widehat{\omega}_1) = (\kappa^* - \kappa)(\widehat{\omega}_2)$ whenever $\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \sigma_{\mathcal{R}}(\widehat{\omega}_2)$ and $\widehat{\omega}_1 \equiv_{\top} \widehat{\omega}_2$ for $\widehat{\omega}_1, \widehat{\omega}_2 \in \widehat{\Omega}^*_+$, where $\widehat{\Omega}^*_+ = \langle \omega \in \Omega \mid \kappa^*(\omega) \neq \infty \rangle$.

The principle of conditional preservation is now realized as an indifference property:

Definition 11 Let κ be an OCF, and let $\mathcal{R}$ be a finite set of conditionals. A revision $\kappa^* = \kappa * \mathcal{R}$ satisfies the *principle of conditional preservation* iff κ^* is indifferent with respect to $\mathcal{R}$ and κ.

So $\kappa * \mathcal{R}$ satisfies the principle of conditional preservation if any change in plausibility is clearly and unambigously induced by $\mathcal{R}$. The next theorem characterizes

[2] Why just $\kappa^* - \kappa$? First, it is more accurate than, e.g., $\max\{0, \kappa^* - \kappa\}$, in the sense of taking differences in degrees of plausibility seriously. Second, it makes use of the conditional "$-$", considering revision as a generalized conditional operation; for more details, see (Kern-Isberner 1999d).

[3] Note that success $\kappa^* \models \mathcal{R}$ is not compulsory for conditional indifference and conditional preservation. We only presuppose $\kappa^*(A) \neq \infty$ for all $(B|A) \in \mathcal{R}$ in order to exclude pathological cases.

revisions of ordinal conditional functions that satisfy the principle of conditional preservation. The theorem is obvious by observing Theorem 9.

Theorem 12 *Let κ, κ^* be OCF's, and let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ be a (finite) set of conditionals in $(\mathcal{L} \mid \mathcal{L})$. A revision $\kappa^* = \kappa * \mathcal{R}$ satisfies the principle of conditional preservation iff $\kappa^*(A_i) \neq \infty$ for all $i, 1 \leqslant i \leqslant n$, and there are numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leqslant i \leqslant n$, such that for all $\omega \in \Omega$,*

$$\kappa^*(\omega) = \kappa(\omega) + \kappa_0 + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^- \qquad (7)$$

Comparing Theorems 9 and 12 with one another, we see that an OCF κ is indifferent with respect to a finite set of conditionals $\mathcal{R}$ iff it can be taken as a revision $\kappa_0 * \mathcal{R}$ satisfying the principle of conditional preservation, where $\kappa_0(\omega) = 0$ for all $\omega \in \Omega$ is the *uniform ordinal conditional function.*

Up to now, we have not yet taken the *success condition* $\kappa^* \models \mathcal{R}$ into regard, postulating that the revised OCF in fact represents the conditionals in $\mathcal{R}$.

Definition 13 Let κ, κ^* be OCF's, and let $\mathcal{R}$ be a set of conditionals. $\kappa^* = \kappa * \mathcal{R}$ is called a *c-revision* iff $\kappa^* \models \mathcal{R}$ and κ^* satisfies the principle of conditional preservation. κ is called a *c-representation* of $\mathcal{R}$, iff $\kappa \models \mathcal{R}$ and κ is indifferent with respect to $\mathcal{R}$.

Theorems 9 and 12 provide simple schemes to construct c-revisions and c-representations. The numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leqslant i \leqslant n$, then have to be chosen appropriately to ensure that $\kappa^{(*)}$ is an ordinal conditional function, and such that $\kappa^{(*)}(AB) < \kappa^{(*)}(A\overline{B})$ for all conditionals $(B|A) \in \mathcal{R}$ (cf. Lemma 1). In the special case that κ is a representation of $\mathcal{R}$, we obtain the following corollary by some easy calculations:

Corollary 14 *Let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ be a (finite) set of conditionals in $(\mathcal{L} \mid \mathcal{L})$, and let κ be an OCF.*

κ is a c-representation of $\mathcal{R}$ iff $\kappa(A_i) \neq \infty$ for all $i, 1 \leqslant i \leqslant n$, and there are numbers $\kappa_0, \kappa_i^+, \kappa_i^- \in \mathbb{Q}, 1 \leqslant i \leqslant n$, such that for all $\omega \in \Omega$,

$$\kappa(\omega) = \kappa_0 + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^-, \qquad (8)$$

and

$$\kappa_i^- - \kappa_i^+ > \min_{\omega \models A_i B_i} \left(\sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \kappa_j^- \right) - \min_{\omega \models A_i \overline{B_i}} \left(\sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \kappa_j^- \right) \qquad (9)$$

Proof. κ is a c-representation of R iff $\kappa \models \mathcal{R}$ and κ is indifferent with respect to $\mathcal{R}$. From Theorem 9, we obtain representation (8). Due to Lemma 1, $\kappa \models \mathcal{R}$ iff $\kappa(A_i B_i) < \kappa(A_i \overline{B}_i)$, i.e. iff

$$\min_{\omega \models A_i B_i} \kappa_0 + \sum_{\substack{1 \leqslant j \leqslant n \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{1 \leqslant j \leqslant n \\ \omega \models A_j \overline{B_j}}} \kappa_j^-$$
$$< \min_{\omega \models A_i \overline{B}_i} \kappa_0 + \sum_{\substack{1 \leqslant j \leqslant n \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{1 \leqslant j \leqslant n \\ \omega \models A_j \overline{B_j}}} \kappa_j^-,$$

which is equivalent to

$$\min_{\omega \models A_i B_i} \kappa_i^+ + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \kappa_j^-$$
$$< \min_{\omega \models A_i \overline{B}_i} \kappa_i^- + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \kappa_j^-.$$

This shows (9). □

The difference $\kappa_i^- - \kappa_i^+$, or the right hand side of (9), respectively, measures the effort needed to establish the i-th conditional. To calculate suitable constants κ_i^+, κ_i^-, we apply the following heuristics: To establish conditional beliefs, one can make confirming worlds more plausible (if required, which amounts to choose $\kappa_i^+ \leqslant 0$), or refuting worlds less plausible (if required, which means $\kappa_i^- \geqslant 0$). The normalizing constant κ_0 then has to be chosen appropriately to ensure that actually an OCF is obtained.

We prefer the second alternative, presupposing

$$\kappa_i^- \geqslant 0, \text{ and } \kappa_i^+ = 0 \text{ for } 1 \leqslant i \leqslant n \qquad (10)$$

Then (9) reduces to

$$\kappa_i^- > \min_{\omega \models A_i B_i} \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \kappa_j^- - \min_{\omega \models A_i \overline{B}_i} \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B_j}}} \kappa_j^- \qquad (11)$$

for $1 \leqslant i \leqslant n$. If there are worlds ω with neutral conditional structure, $\sigma_{\mathcal{R}}(\omega) = 1$, we may set $\kappa_0 = 0$. So, we obtain a c-representation of $\mathcal{R}$ via

$$\kappa(\omega) = \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^-, \qquad (12)$$

where the κ_i^- have to satisfy (11).

Example 15 We will use Corollary 14 and the heuristics (10) to obtain a c-representation (12) of the conditionals $\mathcal{R} = \{r_1, \ldots, r_5\}$ of Example 6. To calculate constants κ_i^- according to (11), Table 1 proves to be helpful. We only have to focus on the $\mathbf{a}^-$-labels of worlds, and we obtain

$$\kappa_5^-, \kappa_4^-, \kappa_1^- > 0,$$
$$\kappa_3^- > \min\{\kappa_1^-, \kappa_2^-\}, \ \kappa_2^- > \min\{\kappa_1^-, \kappa_3^-\}$$

So we set

$$\kappa_5^- = \kappa_4^- = \kappa_1^- = 1, \ \kappa_2^- = \kappa_3^- = 2. \qquad (13)$$

Since there are also worlds ω with $\sigma_{\mathcal{R}}(\omega) = 1$ (cf. Table 1), we set $\kappa_0 = 0$. So we obtain a c-representation of $r_1, \ldots, r_5$ by

$$\kappa(\omega) = \sum_{\substack{1 \leqslant i \leqslant 5 \\ \omega \models A_i \overline{B}_i}} \kappa_i^-, \tag{14}$$

where the A_i, B_i's are the antecedents and consequents of the rules r_i, and the κ_i^- are defined as in (13) for $1 \leqslant i \leqslant 5$ (see also Table 2 in Example 16 below). ■

A comparison with system-Z and system-Z^*

A well-known method to represent a (finite) set $\mathcal{R} = \{r_i = (B_i|A_i) \mid 1 \leqslant i \leqslant n\}$ of conditionals by an OCF is to apply the *system-Z* of Goldszmidt and Pearl (Goldszmidt & Pearl 1992; 1996). The corresponding ranking function κ^z is given by

$$\kappa^z(\omega) = \begin{cases} 0, \text{ if } \omega \text{ does not falsify any } r_i, \\ 1 + \max\limits_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} Z(r_i), \text{otherwise} \end{cases} \tag{15}$$

where Z is an ordering on $\mathcal{R}$ observing the (logical) interactions of the conditionals (for a detailed description of Z, see, for instance, (Goldszmidt & Pearl 1996)). κ^z assigns to each world ω the lowest possible rank admissible with respect to the constraints in $\mathcal{R}$. Comparing (15) with (6), we see that in general, κ^z is *not* a c-representation of $\mathcal{R}$, since in its definition (15), *maximum* is used instead of *summation* (see Example 16 below). The numbers $Z(r_i)$, however, may well serve to define appropriate constants κ_i^- in (6). Setting $\kappa_0 = \kappa_i^+ = 0$, and $\kappa_i^- = Z(r_i) + 1$ for $1 \leqslant i \leqslant n$, we obtain from Z a c-representation κ_c^z of $\mathcal{R}$ via

$$\kappa_c^z(\omega) = \begin{cases} 0, \text{ if } \omega \text{ does not falsify any } r_i, \\ \sum\limits_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} (Z(r_i) + 1), \quad \text{otherwise.} \end{cases} \tag{16}$$

An even more sophisticated representation is obtained by combining the system-Z approach with the principle of maximum entropy (*ME-principle*), yielding system-Z^* (Goldszmidt, Morris, & Pearl 1993). The corresponding Z^*-rankings of the conditionals in $\mathcal{R}$ have to satisfy the following equation (see equation (16) in (Goldszmidt, Morris, & Pearl 1993, p. 225))

$$Z^*(r_i) + \min_{\omega \models A_i \overline{B}_i} \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} Z^*(r_j) \tag{17}$$

$$= 1 + \min_{\omega \models A_i B_i} \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} Z^*(r_j),$$

and κ^* is then calculated by

$$\kappa^*(\omega) = \sum_{\omega \models A_i \overline{B}_i} Z^*(r_i) \tag{18}$$

(see equation (18) in (Goldszmidt, Morris, & Pearl 1993, p. 225)). For so-called *minimal-core sets* – these are sets $\mathcal{R}$ allowing each conditional to be separable from the other rules by restricting conditional interactions –, a procedure is given to calculate Z^*-rankings in (Goldszmidt, Morris, & Pearl 1993).

Like our method, system-Z^* makes use of summation instead of maximization, as in system-Z. And equations (17) determining the Z^*-rankings look similar to our inequality constraints (9). More exactly, if we follow the heuristics (10) and set $Z^*(r_i) = \kappa_i^-$, then system-Z^* turns out to be a special instance of our more general scheme in Corollary 14. In particular, system-Z^* yields a c-representation.

This similarity is not accidental – the ME-principle not only provides a powerful base for system-Z^*, but also influenced the idea of conditional indifference presented in this paper. In (Kern-Isberner 1998), we characterized the ME-principle by four axioms, one of which was the *postulate of conditional preservation*. Conditional preservation for probability functions there was realized in full analogy to that for OCF's defined here. So both c-representations and ME-distributions comply with a fundamental principle for representing conditionals, and it is this principle of conditional preservation (or principle of conditional indifference, respectively) that is responsible for a peculiar thoroughness and accuracy when incorporating conditionals. Since we realized this principle completely in a semi-quantitative setting, we did not have to refer to probabilities and to ME-distributions, and we were able to formalize the acceptance conditions, (9), in a purely qualitative manner.

We will illustrate our method by various examples which are taken from (Goldszmidt, Morris, & Pearl 1993) and (Goldszmidt & Pearl 1996) to allow a direct comparison with system-Z and system-Z^*.

Example 16 Consider once again the conditionals $r_1, \ldots, r_5$ from Example 6. Here we have $Z(r_1) = Z(r_4) = Z(r_5) = 0$ and $Z(r_2) = Z(r_3) = 1$ (for the details, see (Goldszmidt & Pearl 1996, p. 69)). By setting $\kappa_i^- = Z(r_i) + 1$ for $1 \leqslant i \leqslant 5$, we obtain the same constants, (13), as in Example 15. Furthermore, by applying the procedure *Z-rank* in (Goldszmidt, Morris, & Pearl 1993), we calculate $Z^*(r_1) = Z^*(r_4) = Z^*(r_5) = 1$ and $Z^*(r_2) = Z^*(r_3) = 2$, therefore also $Z^*(r_i) = \kappa_i^-$, $1 \leqslant i \leqslant 5$. So, in this example, κ_c^z from (16) actually is the OCF from (14) and coincides with κ^*. For instance,

$$\begin{aligned} \kappa_c^z(\overline{p}\overline{b}\,\overline{f}wa) &= 0, \\ \kappa_c^z(pb\overline{f}wa) &= \kappa_1^- = 1, \\ \kappa_c^z(p\overline{b}fw\overline{a}) &= \kappa_2^- + \kappa_3^- + \kappa_5^- = 5. \end{aligned}$$

In Table 2, we list the ranks of all possible worlds, first computed by system-Z, according to (15), and then computed as a c-representation, κ_c^z, of $\mathcal{R}$, according to (16).

This table reveals clearly that κ^z is not a c-representation of $\mathcal{R}$: Associating symbols $\mathbf{a}_i^+, \mathbf{a}_i^-$ with the conditionals r_i in $\mathcal{R}$, $1 \leqslant i \leqslant 5$, respectively, we ob-

ω	$\kappa^z(\omega)$	$\kappa_c^z(\omega)$	ω	$\kappa^z(\omega)$	$\kappa_c^z(\omega)$
$pbfwa$	2	2	$\overline{p}bfwa$	0	0
$pbfw\overline{a}$	2	3	$\overline{p}bfw\overline{a}$	1	1
$pbf\overline{w}a$	2	3	$\overline{p}bf\overline{w}a$	1	1
$pbf\overline{w}\,\overline{a}$	2	4	$\overline{p}bf\overline{w}\,\overline{a}$	1	2
$pb\overline{f}wa$	1	1	$\overline{p}b\overline{f}wa$	1	1
$pb\overline{f}w\overline{a}$	1	1	$\overline{p}b\overline{f}w\overline{a}$	1	1
$pb\overline{f}\overline{w}a$	1	2	$\overline{p}b\overline{f}\overline{w}a$	1	2
$pb\overline{f}\overline{w}\,\overline{a}$	1	2	$\overline{p}b\overline{f}\overline{w}\,\overline{a}$	1	2
$p\overline{b}fwa$	2	4	$\overline{p}\overline{b}fwa$	0	0
$p\overline{b}fw\overline{a}$	2	5	$\overline{p}\overline{b}fw\overline{a}$	1	1
$p\overline{b}f\overline{w}a$	2	4	$\overline{p}\overline{b}f\overline{w}a$	0	0
$p\overline{b}f\overline{w}\,\overline{a}$	2	5	$\overline{p}\overline{b}f\overline{w}\,\overline{a}$	1	1
$p\overline{b}\,\overline{f}wa$	2	2	$\overline{p}\overline{b}\,\overline{f}wa$	0	0
$p\overline{b}\,\overline{f}w\overline{a}$	2	2	$\overline{p}\overline{b}\,\overline{f}w\overline{a}$	0	0
$p\overline{b}\,\overline{f}\overline{w}a$	2	2	$\overline{p}\overline{b}\,\overline{f}\overline{w}a$	0	0
$p\overline{b}\,\overline{f}\overline{w}\,\overline{a}$	2	2	$\overline{p}\overline{b}\,\overline{f}\overline{w}\,\overline{a}$	0	0

Table 2: Rankings for Example 16

tain

$$\sigma_{\mathcal{R}}\left(\frac{pbfwa \cdot \overline{p}bfw\overline{a}}{pbfw\overline{a} \cdot \overline{p}bfwa}\right) = \frac{\mathbf{a}_1^+\mathbf{a}_2^+\mathbf{a}_3^-\mathbf{a}_4^+\mathbf{a}_5^+ \cdot \mathbf{a}_1^+\mathbf{a}_4^+\mathbf{a}_5^-}{\mathbf{a}_1^+\mathbf{a}_2^+\mathbf{a}_3^-\mathbf{a}_4^+\mathbf{a}_5^- \cdot \mathbf{a}_1^+\mathbf{a}_4^+\mathbf{a}_5^+} = 1,$$

but $\kappa^z\left(\frac{pbfwa \cdot \overline{p}bfw\overline{a}}{pbfw\overline{a} \cdot \overline{p}bfwa}\right) = \kappa^z(pbfwa) + \kappa^z(\overline{p}bfw\overline{a}) - \kappa^z(pbfw\overline{a}) - \kappa^z(\overline{p}bfwa) = 2 + 1 - 2 = 1 \neq 0$.

What is the actual benefit of this formal principle of conditional preservation? Comparing $\kappa^z(\omega)$ to $\kappa_c^z(\omega)$, we see that κ_c^z is more fine-grained. Therefore, it represents more conditionals. Consider, e.g., the conditional $(w|pb\overline{f}a)$ – does a non-flying, but airborne penguin possess wings or not? κ^z does not know, we have $\kappa^z(pb\overline{f}aw) = \kappa^z(pb\overline{f}a\overline{w}) = 1$. On the other hand, κ_c^z accepts this conditional: $\kappa_c^z(pb\overline{f}aw) = 1 < 2 = \kappa_c^z(pb\overline{f}a\overline{w})$. To show that this is more than pure speculation, consider the conditional structures of $pb\overline{f}aw$ and $pb\overline{f}a\overline{w}$: $\sigma_{\mathcal{R}}(pb\overline{f}aw) = \mathbf{a}_1^-\mathbf{a}_2^+\mathbf{a}_3^+\mathbf{a}_4^+$, $\sigma_{\mathcal{R}}(pb\overline{f}a\overline{w}) = \mathbf{a}_1^-\mathbf{a}_2^+\mathbf{a}_3^+\mathbf{a}_4^-$. Thus, except for r_4, both worlds behave exactly the same with respect to the conditionals in $\mathcal{R}$, but, due to r_4, our penguin is supposed to have wings since it is a bird.

Similar arguments apply when considering the conditionals $(a|pbf), (a|pbfw), (w|pbf), (w|pbfa)$: κ^z is totally indifferent when confronted with flying superpenguins - it assigns the same degree of plausibility, 2, to any of the involved worlds. So, it accepts neither of these conditionals, whereas κ_c^z accepts all of them.

Finally, let us consider the conditional $(a|bfw)$. It is accepted by κ^z, as well as by κ_c^z, as may easily be checked. But what happens when the variable p is taken into account? κ^z establishes $(a|\overline{p}bfw)$, but is undecided with respect to $(a|pbfw)$. By contrast, κ_c^z not only accepts both of these conditionals, but also establishes them with equal strength:

$$\kappa_c^z\left(\frac{pbfwa}{pbfw\overline{a}}\right) = \kappa_c^z\left(\frac{\overline{p}bfwa}{\overline{p}bfw\overline{a}}\right) = -1. \tag{19}$$

This is a simple consequence of the principle of conditional preservation in this case, since $\sigma_{\mathcal{R}}\left(\frac{pbfwa \cdot \overline{p}bfw\overline{a}}{pbfw\overline{a} \cdot \overline{p}bfwa}\right) = 1$ (see above). (19) is justified because each of the involved quotients has the same conditional structure, i.e. shows the same behavior, with respect to $\mathcal{R}$. By considering arbitrary group elements in *ker* $\sigma_{\mathcal{R}}$, even very complicated interrelationships between degrees of strength associated with conditionals can be observed. ■

Therefore, using System-Z and the max-operator means to establish conditionals only on a superficial level, whereas obeying the principle of conditional preservation ensures that conditional knowledge is propagated thoroughly and deeply in plausibility structures of epistemic states. This well-behavedness with respect to subconditionals is also observed in (Goldszmidt, Morris, & Pearl 1993). It is illustrated by the next example, too:

Example 17 Consider the conditionals

$r_1 : (f|s)$ *Swedes are fair-haired.*
$r_2 : (t|s)$ *Swedes are tall.*

We apply Corollary 14 and heuristics (10) and calculate $\kappa_1^-, \kappa_2^- \geqslant 0$. So we set $\kappa_1^- = \kappa_2^- = 1$, and we obtain a c-representation, κ, of the form (12). Then not only the subconditional $(f|st)$ of $(f|s)$ is accepted, but also the subconditional $(f|s\overline{t})$, because $\kappa(sf\overline{t}) = 1 < 2 = \kappa(s\overline{f}\,\overline{t})$.

Goldszmidt, Morris and Pearl (Goldszmidt, Morris, & Pearl 1993) compared this situation to the one where instead of r_1, r_2, merely the conditional $(ft|s)$ is learned. Here, only the first subconditional, $(f|st)$, is accepted, but not the second one, $(f|s\overline{t})$. This becomes intelligible by considering conditional structures: $sf\overline{t}$ and $s\overline{f}\,\overline{t}$ both show the same behavior with respect to $(ft|s)$ (namely, they refute it), while sft and $s\overline{f}t$ show different behaviors.

Observing conditional structures emphasizes once again that in general, the joint integration of conditionals cannot be achieved by learning only one conditional – conditionals resist to propositional treatment. In our framework, each conditional constitutes an independent piece of knowledge. ■

Finally, let us consider an example that cannot be dealt with by system-Z^* in a straightforward manner because the involved set of conditionals is not a minimal-core set:

Example 18 Let $\mathcal{R}$ consist of the rules

$$r_1 : (b|a), \quad r_2 : (c|b), \quad r_3 : (c|a).$$

ω	$\sigma_{\mathcal{R}}(\omega)$	ω	$\sigma_{\mathcal{R}}(\omega)$
abc	$\mathbf{a}_1^+\mathbf{a}_2^+\mathbf{a}_3^+$	$\overline{a}bc$	$\mathbf{a}_2^+$
$ab\overline{c}$	$\mathbf{a}_1^+\mathbf{a}_2^-\mathbf{a}_3^-$	$\overline{a}b\overline{c}$	$\mathbf{a}_2^-$
$a\overline{b}c$	$\mathbf{a}_1^-\mathbf{a}_3^+$	$\overline{a}\overline{b}c$	1
$a\overline{b}\overline{c}$	$\mathbf{a}_1^-\mathbf{a}_3^-$	$\overline{a}\overline{b}\overline{c}$	1

Table 3: Conditional structures for Example 18

We list the conditional structures in Table 3 to make argumentation easier. $\mathcal{R}$ is not a minimal-core set in the sense of (Goldszmidt, Morris, & Pearl 1993) because r_3 is only refuted by worlds that also refute either r_1 or r_2. So equation (17) is not solvable to yield Z^*-rankings.

For our approach, however, dealing with $\mathcal{R}$ is no problem: Using (10) and (11), we obtain

$$\kappa_1^-, \kappa_2^- > 0, \quad \kappa_3^- > 0 - \min\{\kappa_1^-, \kappa_2^-\}$$

We set $\kappa_1^- = \kappa_2^- = 1$ and $\kappa_3^- = 0$, and we obtain by (12) an appropriate ranking function, κ (see Table 4). Actually, r_3 seems to be redundant since $\kappa_3^+ = \kappa_3^- = 0$,

ω	$\kappa(\omega)$	$\kappa_1(\omega)$	ω	$\kappa(\omega)$	$\kappa_1(\omega)$
abc	0	0	$\overline{a}bc$	0	0
$ab\overline{c}$	1	2	$\overline{a}b\overline{c}$	1	1
$a\overline{b}c$	1	1	$\overline{a}\overline{b}c$	0	0
$a\overline{b}\overline{c}$	1	2	$\overline{a}\overline{b}\overline{c}$	0	0

Table 4: Rankings for Example 18

and it is in fact already established by r_1 and r_2. But what about the subconditionals $(c|ab)$ and $(c|a\overline{b})$ of r_3? While the first one is accepted, due to the impact of r_2, the second subconditional is not, since r_3 is taken to be redundant (see the conditional structures in Table 3).

To ensure the thorough propagation of conditional knowledge to subconditionals, we have to postulate $\kappa_i^- > 0$ in (10), in order to protect the influence of each conditional against numerical cancellations. With $\kappa_1^- = \kappa_2^- = \kappa_3^- = 1$, we obtain the OCF κ_1 in Table 4, which accepts both subconditionals of r_3, as desired. ■

Unifying qualitative and quantitative approaches

We defined the principle of conditional preservation as an indifference property of the revised ranking function (cf. Definition 11) – conditional preservation means to maintain numerical relationships. Therefore, it appears here as an essentially quantitative notion. Nevertheless, it is applied to ranking functions which are used in a qualitative, or at least, semi-quantitative setting. Moreover, note that conditional indifference is based on the notion of conditional structures which are represented in a completely symbolic way. So, what is the connection between the principle of conditional preservation introduced here, and the approaches to conditional preservation in a purely qualitative framework, as were proposed in (Darwiche & Pearl 1997) and in (Kern-Isberner 1999b)?

In (Kern-Isberner 1999b), we advanced a set of postulates apt to guide the revisions of epistemic states, Ψ, by conditional beliefs, $(B|A)$. There, the idea of conditional preservation was put in formal terms by making use of two relations, *subconditionality*, $\sqsubseteq$ (see (2)), and *perpendicularity*, $\underline{\|}$, on the set of conditionals:

$$(D|C) \,\underline{\|}\, (B|A) \text{ iff either } C \models AB, \text{ or } A\overline{B}, \text{ or } \overline{A}.$$

If $(D|C) \,\underline{\|}\, (B|A)$, then for all worlds which $(D|C)$ may be applied to, $(B|A)$ has the same effect and thus, it yields no further partitioning – $(B|A)$ is *irrelevant* for $(D|C)$.

Conditional preservation then was described in (Kern-Isberner 1999b) by the following three postulates:

(CR5) If $(D|C) \,\underline{\|}\, (B|A)$ then $\Psi \models (D|C)$ iff $\Psi * (B|A) \models (D|C)$.

(CR6) If $(D|C) \sqsubseteq (B|A)$ and $\Psi \models (D|C)$ then $\Psi * (B|A) \models (D|C)$.

(CR7) If $(D|C) \sqsubseteq (\overline{B}|A)$ and $\Psi * (B|A) \models (D|C)$ then $\Psi \models (D|C)$.

These axioms cover the postulates of Darwiche and Pearl in (Darwiche & Pearl 1997) (see (Kern-Isberner 1999b; 1999c)) and support their intuitive ideas with more formal arguments. Applying Definition 11 to the case $\mathcal{R} = \{(B|A)\}$, the principle of conditional preservation reduces to postulate

$$(\kappa^* - \kappa)(\omega_1) = (\kappa^* - \kappa)(\omega_2) \quad \text{if} \quad (B|A)(\omega_1) = (B|A)(\omega_2)$$

for the revised function $\kappa^* = \kappa * \mathcal{R}$ (cf. also (Kern-Isberner 1999a)). It can then be shown that such a revision also satisfies the postulates (CR5)-(CR7) stated above (Kern-Isberner 1999a).

This means that the principle of conditional preservation, phrased in its full (numerical) complexity in this paper, also covers the approaches to conditional preservation proposed in a qualitative framework.

To guarantee the thorough propagation of conditional knowledge to subconditionals, we may supplement (CR5)-(CR7) with another postulate:

(CR8) If $(D|C) \sqsubseteq (B|A)$ and $\Psi \not\models (\overline{D}|C)$, then $\Psi * (B|A) \models (D|C)$.

(CR8) clearly exceeds the paradigm of conditional preservation, in favor of imposing conditional structure as long as there are no conflicts.

Axioms (CR5)-(CR8) only deal, however, with revising an epistemic state by one single conditional. A topic of our ongoing research is to generalize them so as to apply for incorporating sets of conditionals, too.

Summary and outlook

In this paper, we presented an approach to realize the idea of conditional preservation in knowledge representation and belief revision in a very comprehensive way. We introduced a formal notion of conditional structure of worlds, and then we phrased a *principle of conditional preservation* for (revisions of) ordinal conditional functions as indifference with respect to these structures. We presented simple schemes to construct ordinal conditional functions observing this principle, and we compared our approach to system-Z and system-Z^*. Finally, we showed that the principle formalized here also covers other, qualitative approaches to conditional preservation.

In (Kern-Isberner 1999d), we developed these ideas in an even more general setting, with so-called *conditional valuation functions* representing epistemic states. Ordinal conditional functions are special instances of these functions, as well as probability distributions and possibility distributions. So the principle of conditional preservation, as was presented here, can be applied to nearly all major (semi-)quantitative representation forms for epistemic states. Actually, this paper continues and extends work begun in (Kern-Isberner 1998).

The algebraic notion of *conditional structures* of worlds played a crucial part for conditional preservation here. Its representation via group theory does not only provide an elegant methodological framework for handling conditionals. In (Kern-Isberner 1999d), we present an approach to solve the "inverse representation problem" – which conditionals are most appropriate to represent an epistemic state? – by making extensive use of group theoretical means to elaborate conditional structures.

Acknowledgements

I am grateful to three anonymous referees whose comments helped me to improve the presentation of my results.

References

Boutilier, C., and Goldszmidt, M. 1993. Revision by conditional beliefs. In *Proceedings 11th National Conference on Artificial Intelligence (AAAI'93)*, 649–654.

Boutilier, C. 1994. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence* 68:33–85.

Calabrese, P. 1991. Deduction and inference using conditional logic and probability. In Goodman, I.; Gupta, M.; Nguyen, H.; and Rogers, G., eds., *Conditional Logic in Expert Systems*. Elsevier, North Holland. 71–100.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial Intelligence* 89:1–29.

DeFinetti, B. 1974. *Theory of Probability*, volume 1,2. New York: John Wiley and Sons.

Goldszmidt, M., and Pearl, J. 1992. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proceedings Third International Conference on Principles of Knowledge Representation and Reasoning*, 661–672.

Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence* 84:57–112.

Goldszmidt, M.; Morris, P.; and Pearl, J. 1993. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3):220–232.

Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52:263–294.

Kern-Isberner, G. 1998. Characterizing the principle of minimum cross-entropy within a conditional-logical framework. *Artificial Intelligence* 98:169–208.

Kern-Isberner, G. 1999a. Following conditional structures of knowledge. In *KI-99: Advances in Artificial Intelligence, Proceedings of the 23rd Annual German Conference on Artificial Intelligence*, 125–136. Springer Lecture Notes in Artificial Intelligence LNAI 1701.

Kern-Isberner, G. 1999b. Postulates for conditional belief revision. In *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI-99*, 186–191. Morgan Kaufmann.

Kern-Isberner, G. 1999c. Revising by conditional beliefs. In *Proceedings Fourth Dutch-German Workshop on Nonmonotonic Reasoning Rechniques And Their Applications, DGNMR-99.* University of Amsterdam: Institute for Logic, Language and Computation.

Kern-Isberner, G. 1999d. A unifying framework for symbolic and numerical approaches to nonmonotonic reasoning and belief revision. (manuscript).

Lehmann, D.; Magidor, M.; and Schlechta, K. 1999. Distance semantics for belief revision. manuscript.

Lyndon, R., and Schupp, P. 1977. *Combinatorial group theory*. Berlin Heidelberg New York: Springer.

Nute, D. 1980. *Topics in Conditional Logic*. Dordrecht, Holland: D. Reidel Publishing Company.

Rott, H. 1991. A nonmonotonic conditional logic for belief revision, part I: Semantics and logic of simple conditionals. In Fuhrmann, A., and Morreau, M., eds., *The logic of theory change*. Berlin, Heidelberg, New York: Springer. 135–181.

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.

Williams, M.-A. 1994. Transmutations of knowledge systems. In *Proceedings Fourth International Confer-*

ence on Principles of Knowledge Representation and Reasoning, KR-94, 619–629. Morgan Kaufman.