The principle of conditional preservation in belief revision

Gabriele Kern-Isberner

FernUniversität Hagen, Department of Computer Science,
P.O. Box 940, D-58084 Hagen, Germany

email: gabriele.kern-isberner@fernuni-hagen.de

Abstract. Although the crucial role of if-then-conditionals for the dynamics of knowledge has been known for several decades, they do not seem to fit well in the framework of classical belief revision theory. In particular, the propositional paradigm of minimal change guiding the AGM-postulates of belief revision proved to be inadequate for preserving conditional beliefs under revision. In this paper, we present a thorough axiomatization of a principle of conditional preservation in a very general framework, considering the revision of epistemic states by sets of conditionals. This axiomatization is based on a non-standard approach to conditionals, which focuses on their dynamic aspects, and uses the newly introduced notion of conditional valuation functions as representations of epistemic states. In this way, probabilistic revision as well as possibilistic revision and the revision of ranking functions can all be dealt with within one framework. Moreover, we show that our approach can also be applied in a merely qualitative environment, extending AGM-style revision to properly handling conditional beliefs.

1 Introduction

Knowledge is subject to change, either due to changes in the real world, or by obtaining new findings about the domain under consideration. New information may simply extend the prior knowledge base, or be in conflict with it, in which case its incorporation makes complex revision processes necessary. In any case, the modification of knowledge bases brought about by learning new information may drastically alter the response behavior of knowledge systems to queries; e.g. answers that were meaningful in the context of the prior knowledge base, might become irrelevant or even false in the light of new information.

Belief revision, the theory of dynamics of knowledge, has been mainly concerned with propositional beliefs for a long time. The most basic approach here is the *AGM-theory* presented in the seminal paper [1] as a set of postulates outlining appropriate revision mechanisms in a propositional logical environment. If-then-conditionals, on the other hand, seem to play an ambivalent role in belief revision: Although their dynamic power as *revision policies* has been appreciated (see e.g. [36, 5]), Gärdenfors' *triviality result* [16] describes an obvious incompatibility between conditionals and *classical* AGM-approaches. This incompatibility, however, can be resolved by leaving the narrow framework of classical logic – first, conditional beliefs must be understood as fundamentally different from propositional beliefs (cf. [31]) and hence be treated differently, and second, instead of focusing on belief sets (i.e. deductively closed propositional theories) containing all certain beliefs, one should consider belief states or epistemic states, respectively, as representations of cognitive states of intelligent agents. Although the close connections between belief revision and conditionals, on the one side, and between belief revision and epistemic orderings, on the other side, has been apparent for many years (cf. [36, 19]), it was only quite recently that first approaches extending the AGM-theory to that broader framework have been brought forth: Darwiche and Pearl [8] reformulated the AGM-postulates for revising epistemic states by propositional beliefs and, moreover, they formulated four new postulates dealing explicitly with conditional beliefs. Instead of following the *minimal change paradigm* which guides propositional AGM-revision, Darwiche and Pearl's postulates vaguely outline how to preserve conditional beliefs under propositional revision. In [22], we then presented a complete set of axioms for revising epistemic states by conditional beliefs, extending propositional AGM-revision and covering the postulates of Darwiche and Pearl.

Instead of only regarding the results of belief change, as in AGM-theory, studying belief revision in the framework of epistemic states and conditionals means to observe the very process of belief dynamics. Perhaps the most important consequence of this is that, in overcoming classical borders and peculiarities, it opens up the view to a most general framework which unifies belief revision, nonmonotonic reasoning and inductive representation of complex, conditional knowledge. To be more precise, belief revision and nonmonotonic reasoning can be linked via conditionals in epistemic states in the following way: A nonmonotonically implies B, based on the knowledge given by the epistemic state Ψ, $(A \mid\!\sim_{\Psi} B)$, iff the conditional $(B|A)$ is accepted in Ψ ($\Psi \models (B|A)$) iff revising Ψ by A yields belief in B ($\Psi * A \models B$). Note that here background knowledge represented by Ψ can be taken explicitly into account, in contrast to the purely propositional view in [33]. Furthermore, inductive knowledge representation can be understood as revising a uniform belief state, expressing complete ignorance, by the (conditional) knowledge to be represented.

For these reasons, revising epistemic states by conditional beliefs should not be considered as an artifact, but rather be understood as one of the most fundamental and powerful processes in formal knowledge management. A thorough axiomatization of an appropriate *principle of conditional preservation* in the sense of [8] and [22] will be able to serve as an important guideline for handling those revisions. To this end, we even go one step further: We leave the purely qualitative framework and enter into semi-quantitative (i.e. ordinal) and quantitative (i.e. probabilistic) environments, finding once again that more complex surroundings provide clearer, unifying views. In this paper, *conditional valuation functions* are introduced as quite general representation of epistemic states. Ordinal conditional functions, possibility distributions and probability functions are special instances of conditional valuation functions. We then formalize a most general principle of conditional preservation, dealing with the revision of conditional valuation functions by sets of (quantified) conditional beliefs. This principle is inspired by properties of optimal information-theoretical methods, and hence can be regarded as a most appropriate paradigm to deal with conditional information. As ordinal epistemic states (such as ordinal conditional functions and possibility distributions) also allow a purely qualitative view, we investigate the consequences of this quantitative principle of conditional preservation in a qualitative setting. We show that our quantitative principle of conditional preservation implies the validity of the axioms for conditional belief revision of [22] and hence also provides a high-level formalization of Darwiche and Pearl's ideas [8].

The principle of conditional preservation to be axiomatized in this paper is based on a non-standard theory of conditionals which captures the dynamic effects of establishing conditional relationships within epistemic states. *Conditional structures* make interactions of conditionals transparent and computable – a crucial problem when revising by sets of conditionals. This framework we consider conditionals in not only concerns belief revision, nonmonotonic reasoning and inductive knowledge representation, but also helps unifying qualitative and quantitative approaches. We clearly differentiate between numerical and structural aspects of conditionals, by first building up a formal, algebraic frame for conditionals and then linking this frame to numerical values. It is just this idea of separating structures from numbers that provides a solid basic theory of conditionals with applications in (apparently) very different domains.

This paper is organized as follows: Section 2 contains some formal preliminaries, and here we briefly explain the different types of epistemic states we are going to consider. In Section 3, conditional valuation functions are introduced as basic representations of (semi-)quantitative epistemic states. In Section 4, we present a new, dynamic view on conditionals; in particular, we define the notions of *subconditional* and of *perpendicular conditionals*, which are crucial for formalizing a qualitative principle of conditional preservation in Section 5. Section 6 prepares the axiomatization of the quantitative principle of conditional preservation in Section 7 by explaining *conditional structures* and *conditional indifference*. Finally, Section 8 shows that both principles are compatible. Section 9 concludes with highlighting the main results of this paper.

All proofs are omitted for the sake of brevity, but can be found in [27].

2 Conditionals, epistemic states and belief revision

We start with a finitely generated propositional language $\mathcal{L}$, with atoms $a, b, c, \ldots$, and with formulas $A, B, C, \ldots$. For conciseness of notation, we will omit the logical *and*-connector, writing AB instead of $A \wedge B$, and barring formulas will indicate negation, i.e. $\overline{A}$ means $\neg A$. Let Ω denote the set of possible worlds over $\mathcal{L}$; Ω will be taken here simply as the set of all propositional interpretations over $\mathcal{L}$. $\omega \models A$ means that the propositional formula $A \in \mathcal{L}$ holds in the possible world $\omega \in \Omega$.

By introducing a new binary operator $|$, we obtain the set $(\mathcal{L} \mid \mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}$ of conditionals over $\mathcal{L}$. $(B|A)$ formalizes "*if A then B*" and establishes a plausible, probable, possible etc connection between the *antecedent* A and the *consequent* B. Here, conditionals are supposed not to be nested, that is, antecedent and consequent of a conditional will be propositional formulas.

Conditionals are usually considered within richer structures such as *epistemic states*. Besides certain (logical) knowledge, epistemic states also allow the representation of preferences, beliefs, assumptions etc of an intelligent agent. In a purely qualitative setting, preferences are assumed to be given by an epistemic pre-ordering on $\mathcal{L}$ (reflexive and transitive, but not symmetrical, and mostly induced by pre-orderings on worlds). In a (semi-)quantitative setting, also degrees of plausibility, probability, possibility, necessity etc can be expressed. Here, most widely used representations of epistemic states are

- *probability functions* (or *probability distributions*) $P : \Omega \to [0,1]$ with $\sum_{\omega \in \Omega} P(\omega) = 1$. The probability of a formula $A \in \mathcal{L}$ is given by $P(A) = \sum_{\omega \models A} P(\omega)$. Note that, since $\mathcal{L}$ is finitely generated, Ω is finite, too, and we only need additivity instead of σ-additivity. Conditionals are interpreted via conditional probability, so we have $P(B|A) = \frac{P(AB)}{P(A)}$ for $P(A) > 0$, and $P \models (B|A)\,[x]$ iff $P(B|A) = x$ $(x \in [0,1])$.
- *ordinal conditional functions, OCFs,* (also called *ranking functions*) $\kappa : \Omega \to \mathbb{N} \cup \{\infty\}$ with $\kappa^{-1}(0) \neq \emptyset$, expressing degrees of plausibility of propositional formulas A by specifying degrees of disbeliefs of their negations $\overline{A}$ (cf. [38]). More formally, we have $\kappa(A) := \min\{\kappa(\omega) \mid \omega \models A\}$, so that $\kappa(A \vee B) = \min\{\kappa(A), \kappa(B)\}$. Hence, at least one of $\kappa(A), \kappa(\overline{A})$ must be 0. A proposition A is believed if $\kappa(\overline{A}) > 0$ (which implies particularly $\kappa(A) = 0$). Degrees of plausibility can also be assigned to conditionals by setting $\kappa(B|A) = \kappa(AB) - \kappa(A)$. A conditional $(B|A)$ is accepted in the epistemic state represented by κ, or κ *satisfies* $(B|A)$, written as $\kappa \models (B|A)$, iff $\kappa(AB) < \kappa(A\overline{B})$, i.e. iff AB is more plausible than $A\overline{B}$. We can also specify a numerical degree of plausibility of a

conditional by defining $\kappa \models (B|A)\,[n]$ iff $\kappa(AB) + n < \kappa(A\overline{B})$ $(n \in \mathbb{N})$. OCF's are the qualitative counterpart of probability distributions. Their plausibility degrees may be taken as order-of-magnitude abstractions of probabilities (cf. [17,18]).

– *possibility distributions* $\pi : \Omega \to [0,1]$ with $\max_{\omega \in \Omega} \pi(\omega) = 1$. Each possibility distribution induces a *possibility measure* on $\mathcal{L}$ via $\pi(A) := \max_{\omega \models A} \pi(\omega)$. Since the correspondence between possibility distributions and possibility measures is straightforward and one-to-one, we will not distinguish between them. A *necessity measure* N_π can also be based on π by setting $N_\pi(A) := 1 - \pi(\overline{A})$. Possibility measures and necessity measures are dual, so it is sufficient to know only one of them. Furthermore, a possibility degree can also be assigned to a conditional $(B|A)$ by setting $\pi(B|A) = \frac{\pi(AB)}{\pi(A)}$, in full analogy to Bayesian conditioning in probability theory. Note that we also make use of the product operation in $[0,1]$. That means, that our approach is not only based upon comparing numbers, but also takes relations between numbers into account. These numerical relationships encode important information about the (relative) strength of conditionals which proves to be particularly crucial for representation and revision tasks. This amounts to carrying over Spohn's argumentation in [38] to a possibilistic framework (see also [23] and [3]).
A conditional $(B|A)$ *is accepted in* π, $\pi \models (B|A)$, iff $\pi(AB) > \pi(A\overline{B})$ (which is equivalent to $\pi(\overline{B}|A) < 1$ and $N_\pi(B|A) = 1 - \pi(\overline{B}|A) > 0$) (cf. [12]). So, in accordance with intuition, a conditional $(B|A)$ is accepted in the epistemic state modeled by a possibility distribution, if its confirmation (AB) is considered to be more possible (or plausible) than its refutation $(A\overline{B})$. This definition can be generalized by saying that π *accepts* $(B|A)$ *with degree* $x \in (0,1]$, $\pi \models (B|A)\,[x]$, iff $N_\pi(B|A) \geqslant x$ iff $\pi(A\overline{B}) \leqslant (1-x)\pi(AB)$.
Possibility distributions are similar to ordinal conditional functions (cf. [2]), but realize degrees of possibility (or plausibility) in a non-discrete, compact domain. They can be taken as fuzzy representations of epistemic states (cf. [30,11]), and are closely related to belief functions (cf. [4]).

With each epistemic state Ψ (either qualitative or (semi-)quantitative) one can associate the set $Bel(\Psi) = \{A \in \mathcal{L} \mid \Psi \models A\}$ of those beliefs the agent accepts as most plausible. $Bel(\Psi)$ is supposed to consist of formulas (or to be a formula, respectively) of $\mathcal{L}$ and hence is subject to classical belief revision theory which investigates the changing of propositional beliefs when new information becomes evident. Here, most important work has been done by Alchourron, Gärdenfors and Makinson in presenting in [1] a catalogue of postulates (the so-called *AGM-postulates*) which a well-behaved revision operator $*$ should obey. The revision of epistemic states, however, cannot be reduced to propositional revision, for two reasons: First, two *different* epistemic states Ψ_1, Ψ_2 may have *equivalent* belief sets $Bel(\Psi_1) \equiv Bel(\Psi_2)$. Thus an epistemic state is not described uniquely by its belief set, and revising Ψ_1 and Ψ_2 by new (propositional) information A may result in different revised belief sets $Bel(\Psi_1 * A) \not\equiv Bel(\Psi_2 * A)$. Second, epistemic states may represent different kinds of beliefs, and beliefs on different levels of acceptance. So "information" in the context of epistemic state must be understood as a much more complex concept than provided by the propositional framework. Incorporating new information in an epistemic state means, for instance, to change degrees of plausibility, or to establish a new conditional relationship. Nevertheless, the revision of Ψ by $A \in \mathcal{L}$ also yields a revised belief set $Bel(\Psi * A) \subseteq \mathcal{L}$, and of course, this revision should obey the standards of the AGM theory. So, Darwiche and Pearl have reformulated the AGM-postulates for belief revision so as to comply with the framework of epistemic states (cf. [8]):

Suppose Ψ, Ψ_1, Ψ_2 to be epistemic states and $A, A_1, A_2, B \in \mathcal{L}$;

(R*1) A is believed in $\Psi * A$: $Bel(\Psi * A) \models A$.
(R*2) If $Bel(\Psi) \wedge A$ is satisfiable, then $Bel(\Psi * A) \equiv Bel(\Psi) \wedge A$.
(R*3) If A is satisfiable, then $Bel(\Psi * A)$ is also satisfiable.
(R*4) If $\Psi_1 = \Psi_2$ and $A_1 \equiv A_2$, then $Bel(\Psi_1 * A_1) \equiv Bel(\Psi_2 * A_2)$.
(R*5) $Bel(\Psi * A) \wedge B$ implies $Bel(\Psi * (A \wedge B))$.
(R*6) If $Bel(\Psi * A) \wedge B$ is satisfiable then $Bel(\Psi * (A \wedge B))$ implies $Bel(\Psi * A) \wedge B$.

Epistemic states, conditionals and revision are related by the so-called *Ramsey test*, according to which a conditional $(B|A)$ is accepted in an epistemic state Ψ, iff revising Ψ by A yields belief in B:

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

The Ramsey test can be generalized for quantified beliefs and conditionals in a straightforward way:

$$\Psi \models (B|A)\,[x] \quad \text{iff}\, \Psi * A \models B\,[x].$$

In this paper, we will consider quantified as well as unquantified (structural) conditionals, where the quantifications are taken from the proper domain ($[0,1]$ or $\mathbb{N} \cup \{\infty\}$, respectively). If $\mathcal{R}^* = \{(B_1|A_1)\,[x_1], \ldots, (B_n|A_n)\,[x_n]\}$ is a set of quantified conditionals, then $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ will be its structural counterpart and vice versa.

3 Conditional valuation functions

What is common to probability functions, ordinal conditional functions, and possibility measures is, that they make use of *two different operations* to handle both purely propositional information and conditionals adequately. Therefore, we will introduce the abstract notion of a conditional valuation function to reveal more clearly and uniformly the way in which (conditional) knowledge may be represented and treated within epistemic states. As an adequate structure, we assume an algebra $\mathcal{A} = (\mathcal{A}, \oplus, \odot, 0^{\mathcal{A}}, 1^{\mathcal{A}})$ of real numbers to be equipped with two operations, $\oplus$ and $\odot$, such that

- $(\mathcal{A}, \oplus)$ is an associative and commutative structure with neutral element $0^{\mathcal{A}}$;
- $(\mathcal{A} - \{0^{\mathcal{A}}\}, \odot)$ is a commutative group with neutral element $1^{\mathcal{A}}$;
- the rule of distributivity holds, i.e. $x \odot (y \oplus z) = (x \odot y) \oplus (x \odot z)$ for $x, y, z \in \mathcal{A}$;
- $\mathcal{A}$ is totally ordered by $\leqslant_{\mathcal{A}}$ which is compatible with $\oplus$ and $\odot$ in that $x \leqslant_{\mathcal{A}} y$ implies both $x \oplus z \leqslant_{\mathcal{A}} y \oplus z$ and $x \odot z \leqslant_{\mathcal{A}} y \odot z$ for all $x, y, z \in \mathcal{A}$.

So $\mathcal{A}$ is close to be an ordered field, except that the elements of $\mathcal{A}$ need not be invertible with respect to $\oplus$.

Definition 3.1 (conditional valuation function). *A* conditional valuation function *is a function $V : \mathcal{L} \to \mathcal{A}$ from the set $\mathcal{L}$ of formulas into the algebra $\mathcal{A}$ satisfying the following conditions:*

1. *$V(\bot) = 0^{\mathcal{A}}, V(\top) = 1^{\mathcal{A}}$, and for exclusive formulas A, B (i.e. $AB \equiv \bot$), we have $V(A \vee B) = V(A) \oplus V(B)$;*
2. *for each conditional $(B|A) \in (\mathcal{L} \mid \mathcal{L})$ with $V(A) \neq 0^{\mathcal{A}}$, $V(B|A) = V(AB) \odot V(A)^{-1}$, where $V(A)^{-1}$ is the $\odot$-inverse element of $V(A)$ in $\mathcal{A}$.*

Conditional valuation functions assign degrees of certainty, plausibility, possibility etc to propositional formulas and to conditionals. Making use of two operations, they provide a framework for considering and treating conditional knowledge as fundamentally different from propositional knowledge, a point that is stressed by various authors and that seems to be indispensable for representing epistemic states adequately (cf. [8]). There is, however, a close relationship between propositions (facts) and conditionals – facts may be considered as conditionals of a degenerate form by identifying A with $(A|\top)$: Indeed, we have $V(A|\top) = V(A) \odot (1^V)^{-1} = V(A)$. Therefore, conditionals should be regarded as extending propositional knowledge by a new dimension.

For each conditional valuation function V, we have $V(A) = \sum_{\omega \models A}^{\oplus} V(\omega)$, so V is determined uniquely by its values on interpretations or on possible worlds, respectively, and we will also write $V : \Omega \to \mathcal{A}$. For all $A \in \mathcal{A}$, we have $0^{\mathcal{A}} \leqslant_{\mathcal{A}} V(A) \leqslant_{\mathcal{A}} 1^{\mathcal{A}}$. It is easy to see that any conditional valuation function $V : \mathcal{L} \to \mathcal{A}$ is a *plausibility measure*, in the sense of Friedman and Halpern, ([15, 14]), that is, it fulfills $V(\bot) \leqslant_{\mathcal{A}} V(A)$ for all $A \in \mathcal{L}$, and $A \models B$ implies $V(A) \leqslant_{\mathcal{A}} V(B)$.

A notion which is well-known from probability theory may be generalized for conditional valuation functions: A conditional valuation function V is said to be *uniform* if $V(\omega) = V(\omega')$ for all worlds ω, ω', i.e. if it assigns the same degree of plausibility to each world.

The following examples show that the newly introduced notion of a conditional valuation function indeed covers probability functions, ordinal conditional functions and possibility distributions:

Example 3.1. Each *probability function* P may be taken as a conditional valuation function $P : \Omega \to (\mathbb{R}^+, +, \cdot, 0, 1)$, where $\mathbb{R}^+$ denotes the set of all non-negative real numbers. Conversely, each conditional valuation function $V : \Omega \to (\mathbb{R}^+, +, \cdot, 0, 1)$ is a probability function.

Similarly, each *ordinal conditional function* κ is a conditional valuation function $\kappa : \Omega \to (\mathbb{Z} \cup \{\infty\}, \min, +, \infty, 0)$, where $\mathbb{Z}$ denotes the set of all integers, and any *possibility measure* π can be regarded as a conditional valuation function $\pi : \Omega \to (\mathbb{R}^+, \max, \cdot, 0, 1)$.

Conditional valuation functions not only provide an abstract means to quantify epistemological attitudes. Their extended ranges allow us to calculate and compare arbitrary proportions of values attached to single worlds. This will prove quite useful to handle complex conditional interrelationships.

By means of a conditional valuation function $V : \mathcal{L} \to \mathcal{A}$, we are able to validate propositional as well as conditional beliefs. We may say, for instance, that proposition A is *believed in* V, $V \models A$, iff $V(A) = 1^{\mathcal{A}}$, or that the conditional $(B|A)$ is *valid* or *accepted in* V, $V \models (B|A)$, iff $V(A) \neq 0^{\mathcal{A}}$ and $V(A\overline{B}) <_{\mathcal{A}} V(AB)$, i.e. iff AB is more plausible (probable, possible etc.) than $A\overline{B}$. In this way, conditional valuation functions are apt to represent epistemic states.

Note that there is a difference between taking a proposition A for granted or to be true, which would be properly expressed by $V(A) = 1^{\mathcal{A}}$, and considering A to be plausible, which amounts to stating $V(A) >_{\mathcal{A}} V(\overline{A})$. It is only from the second point of view, that propositions, A, can be consistently identified with degenerate conditionals, $(A|\top)$. Since belief revision is mostly concerned with revising plausible beliefs by new plausible beliefs, conditionals offer a most adequate framework to study revision methods in, and conditional valuation functions allow us to distinguish between true facts and plausible propositions.

4 A dynamic view on conditionals

As it is well-known, a conditional $(B|A)$ is an object of a three-valued nature, partitioning the set of worlds Ω in three parts: those worlds satisfying AB and thus *verifying* the conditional, those worlds satisfying $A\overline{B}$, thus *falsifying* the conditional, and those worlds not fulfilling the premise A and so which the conditional may not be applied to at all. The following representation of $(B|A)$ as a *generalized indicator function* goes back to de Finetti [9]:

$$(B|A)(\omega) = \begin{cases} 1 & : \quad \omega \models AB \\ 0 & : \quad \omega \models A\overline{B} \\ u & : \quad \omega \models \overline{A} \end{cases} \tag{2}$$

where u stands for *unknown* or *indeterminate*. Two conditionals are considered *equivalent* iff the corresponding indicator functions are identical, i.e. $(B|A) \equiv (D|C)$ iff $A \equiv C$ and $AB \equiv CD$ (see e.g. [6]). Usually, equation (2) is applied in a static way, namely, to check if possible worlds verify, or falsify a conditional, or are simply neutral with respect to it. In the context of inductive knowledge representation or belief revision, however, when conditionals are to be learned, it also provides a dynamic view on how to incorporate conditional dependencies adequately in a belief state (which might be the uniform one): The conditional $(B|A)$ distinguishes clearly between verifying, falsifying and neutral worlds, but it does not distinguish between worlds within one and the same of these partitioning sets. So, in order to establish $(B|A)$, if demanded with a suitable degree of certainty, the plausibilities or probabilities of worlds have to be shifted uniformly, depending on to which of the partitioning sets the worlds belong. In this sense, conditionals have effects on possible worlds, taking an active role (like *agents*) in the revision (or representation) process.

To make things more precise, we define the *verifying set* $(B|A)^+ := Mod(AB)$, and the *falsifying set* $(B|A)^- := Mod(A\overline{B})$ of a conditional $(B|A)$. $Mod(\overline{A})$ is called the *neutral set* of $(B|A)$. Each of these sets may be empty. If $(B|A)^+ = \emptyset$, $(B|A)$ is called *contradictory*, if $(B|A)^- = \emptyset$, $(B|A)$ is called *tautological*, and if $Mod(\overline{A}) = \emptyset$, i.e. A is tautological, $(B|A)$ is called a *fact*. Verifying and falsifying set clearly identify a conditional up to equivalence. Note that, although $(B|A)$ and $(\overline{B}|A)$ induce the same partitioning on Ω, their verifying and falsifying sets are different, in that $(B|A)^+ = (\overline{B}|A)^-$ and $(B|A)^- = (\overline{B}|A)^+$.

Example 4.1. $(\overline{A}|A)$ is a contradictory conditional, $(A|A)$ is tautological and $(A|\top)$ is a fact.

As usual, propositional formulas $A \in \mathcal{L}$ may be identified with factual conditionals $(A|\top)$. Hence, the results to be presented can be related to the theory of propositional revision, as will be done in Section 5. It should be emphasized, however, that in our framework, $(A|\top)$ should be understood as "A is plausible" or "A is believed", whereas A actually means "A is true". Hence a clear distinction between propositions as *logical* statements and propositions as *epistemic* statements is possible, and is indeed respected in our framework (see [26]).

Next, we introduce the notion of a subconditional:

Definition 4.1 (subconditional, $\sqsubseteq$). *A conditional $(D|C)$ is called a* subconditional *of $(B|A)$, $(D|C) \sqsubseteq (B|A)$, iff $(D|C)^+ \subseteq (B|A)^+$ and $(D|C)^- \subseteq (B|A)^-$.*

The $\sqsubseteq$-relation may be expressed using the standard ordering $\leqslant$ between propositional formulas: $A \leqslant B$ iff $A \models B$, i.e. iff $Mod(A) \subseteq Mod(B)$:

Lemma 4.1. *Let $(B|A), (D|C) \in (\mathcal{L} \mid \mathcal{L})$. Then $(D|C)$ is a subconditional of $(B|A)$, $(D|C) \sqsubseteq (B|A)$, iff $CD \leqslant AB$ and $C\overline{D} \leqslant A\overline{B}$; in particular, if $(D|C) \sqsubseteq (B|A)$ then $C \leqslant A$.*

Thus $(D|C) \sqsubseteq (B|A)$ if the effect of the former conditional on worlds is in line with the latter one, but $(D|C)$ possibly applies to fewer worlds. Furthermore, the well-known equivalence relation for conditionals can also be taken as to be induced by $\sqsubseteq$:

Lemma 4.2. *Two conditionals $(B|A)$ and $(D|C)$ are equivalent, $(B|A) \equiv (D|C)$ iff $(B|A) \sqsubseteq (D|C)$ and $(D|C) \sqsubseteq (B|A)$.*

We will now introduce another relation between conditionals that is quite opposite to the subconditional relation and so describes another extreme of possible conditional interaction:

Definition 4.2 (perpendicular conditionals, $\perp\!\!\!\perp$). *Suppose $(B|A), (D|C) \in (\mathcal{L} \mid \mathcal{L})$ are two conditionals. $(D|C)$ is called* perpendicular *to $(B|A)$, $(D|C) \perp\!\!\!\perp (B|A)$, iff either $Mod(C) \subseteq (B|A)^+$, or $Mod(C) \subseteq (B|A)^-$, or $Mod(C) \subseteq Mod(\overline{A})$, i.e. iff either $C \leqslant AB$, or $C \leqslant A\overline{B}$, or $C \leqslant \overline{A}$.*

The perpendicularity relation symbolizes a kind of *irrelevance* of one conditional for another one. We have $(D|C) \perp\!\!\!\perp (B|A)$ if $Mod(C)$, i.e. the range of application of the conditional $(D|C)$, is completely contained in exactly one of the sets $(B|A)^+, (B|A)^-$ or $Mod(\overline{A})$. So for all worlds which $(D|C)$ may be applied to, $(B|A)$ has the same effect and yields no further partitioning. Note, that $\perp\!\!\!\perp$ is not a symmetric relation; $(D|C) \perp\!\!\!\perp (B|A)$ rather expresses that $(D|C)$ is *not affected* by $(B|A)$, or, that $(B|A)$ is *irrelevant* for $(D|C)$.

Example 4.2. Suppose a, b, c are atoms of the language $\mathcal{L}$. Subconditionals of $(b|a)$ are typically obtained by strengthening the antecedent: $(b|ac)$ and $(b|a\overline{c})$ are both subconditionals of $(b|a)$, $(b|ac), (b|a\overline{c}) \sqsubseteq (b|a)$. As an example for perpendicularity, consider the conditionals $(c|ab), (c|a\overline{b})$ and $(c|\overline{a})$ which are all perpendicular to $(b|a)$: $(c|ab), (c|a\overline{b}), (c|\overline{a}) \perp\!\!\!\perp (b|a)$.

It should be remarked that neither $\sqsubseteq$ nor $\perp\!\!\!\perp$ provide new insights for (flat) propositions, when identifying propositions with factual conditionals. It is easily seen that $(B|\top) \sqsubseteq (A|\top)$ if and only if A and B are logically equivalent, and $(B|\top) \perp\!\!\!\perp (A|\top)$ can only hold if A is tautological or contradictory. Both relations need the richer epistemic framework of conditionals to show their usefulness. For a more thorough discussion of the relations $\sqsubseteq$ and $\perp\!\!\!\perp$, see [27].

5 A principle of conditional preservation in a qualitative framework

In [8], Darwiche and Pearl discussed the problem of preserving conditional beliefs under (propositional) belief revision in an AGM-environment. They emphasized that conditional beliefs are different in nature from propositional beliefs, and that the *minimal change paradigm* which is crucial for the AGM-theory [1] should not be blindly applied when considering conditionals. They reformulated the AGM-postulates in the richer framework of epistemic states (cf. Section 2) and extended this approach by phrasing four new postulates explicitly dealing with the acceptance of conditionals in epistemic states, in the following denoted as *DP-postulates*:

DP-postulates for conditional preservation:

(C1) If $C \models B$ then $\Psi \models (D \mid C)$ iff $\Psi * B \models (D \mid C)$.
(C2) If $C \models \overline{B}$ then $\Psi \models (D \mid C)$ iff $\Psi * B \models (D \mid C)$.
(C3) If $\Psi \models (B \mid A)$ then $\Psi * B \models (B \mid A)$.
(C4) If $\Psi * B \models (\overline{B} \mid A)$ then $\Psi \models (\overline{B} \mid A)$.

The DP-postulates were supported by plausible arguments and many examples (for a further discussion, see the original paper [8]). They are crucial for handling iterated revisions via the Ramsey test (1). For instance, by applying (1), (C2) can be reformulated to guide iterated revisions, as follows: If $C \models \overline{B}$ then $\Psi * C \models D$ iff $\Psi * B * C \models D$. The DP-postulates are not indisputable. An objection often made is the following: Let $C = \overline{p}$ and $B = pq$ (p, q atoms), such that $C \models \overline{B}$. Then (C2) yields $\Psi * \overline{p} \models D$ iff $\Psi * pq * \overline{p} \models D$, which implies $Bel(\Psi * \overline{p}) = Bel(\Psi * pq * \overline{p})$ – the information conveyed by learning (p and) q has apparently been extinguished when $\overline{p}$ becomes evident. As atoms are assumed to be independent, this seems to be counterintuitive. Actually, this example does not really cast doubt on the DP-postulates, rather it proves the inappropriateness of a strictly propositional framework for belief revision. In such a framework, it is impossible to distinguish between revising by p *and* q, on the one hand, and $p \wedge q \equiv pq$, on the other hand, since sets of formulas are identified with the conjunction of the corresponding formulas. pq, however, suggests an intensional connection between p and q, whereas $\{p, q\}$ does not. Furthermore, (C2) does not demand the equivalence of the involved epistemic states $\Psi * \overline{p}$ and $\Psi * pq * \overline{p}$, but only the identity of the corresponding belief sets (cf. Section 2). Again, this distinction gets lost when focusing on propositional beliefs.

In [22], we considered conditionals under revision in an even broader framework, setting up *postulates for revising epistemic states by conditional beliefs*:

Postulates for conditional revision:

Suppose Ψ is an epistemic state and $(B|A), (D|C)$ are conditionals. Let $\Psi * (B|A)$ denote the result of revising Ψ by a non-contradictory conditional $(B|A)$.

(CR0) $\Psi * (B|A)$ is an epistemic state.
(CR1) $\Psi * (B|A) \models (B|A)$ (*success*).
(CR2) $\Psi * (B|A) = \Psi$ iff $\Psi \models (B|A)$ (*stability*).
(CR3) $\Psi * B := \Psi * (B|\top)$ induces a propositional AGM-revision operator.
(CR4) $\Psi * (B|A) = \Psi * (D|C)$ whenever $(B|A) \equiv (D|C)$.
(CR5) If $(D|C) \perp\!\!\!\perp (B|A)$ 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 $\Psi \models (D|C)$.

The postulates (CR0)-(CR2) and (CR4) realize basic ideas of AGM-revision in this more general framework, and (CR3) links conditional belief revision to propositional AGM-revision. (CR5)-(CR7) are the proper axioms to formalize a qualitative principle of conditional preservation. They realize the idea of preserving conditional beliefs by use of the two relations $\sqsubseteq$ and $\perp\!\!\!\perp$, which reflect possible interactions between conditionals. In detail, (CR5) claims that revising by a conditional should preserve all conditionals to which that conditional is irrelevant, in the sense described by the relation $\perp\!\!\!\perp$. The rationale behind this postulate is the following: The validity of a conditional $(B|A)$ in an epistemic state Ψ depends on the relation between (some) worlds in $Mod(AB)$ and (some) worlds in $Mod(A\overline{B})$. So incorporating $(B|A)$ into Ψ may require a shift between $Mod(AB)$ on one side and $Mod(A\overline{B})$ on the other side, but should leave intact any relations between worlds within $Mod(AB)$, $Mod(A\overline{B})$, or $Mod(\overline{A})$. These relations may be captured by conditionals $(D|C)$ not affected by $(B|A)$, that is, by conditionals $(D|C) \perp\!\!\!\perp (B|A)$.

(CR6) states that conditional revision should bring about no change for conditionals that are already in line with the revising conditional, and (CR7) guarantees that no conditional change contrary to the revising conditional is caused by conditional revision.

In particular, by considering a propositional formula as a degenerated conditional with tautological antecedent, each conditional revision operator induces a propositional revision operator, as described by (CR3). For this propositional revision operator, the postulates (CR0)-(CR6) above are trivially fulfilled within an AGM-framework. Postulate (CR7) then reads

(CR7)$^{\mathbf{prop}}$ If $\Psi * A \models \overline{A}$, then $\Psi \models \overline{A}$

An AGM-revision operator, obeying the postulate of success and yielding a consistent belief state, as long as the revising proposition A is not inconsistent, would never fulfill the precondition $\Psi * A \models \overline{A}$. Hence (CR7) is vacuous in an AGM-framework. If we only presuppose that $*$ satisfies the AGM-postulate of success, then $\Psi * A \models \overline{A}$ implies the inconsistency of $\Psi * A$, although A is assumed to be non-contradictory. A reasonable explanation for this would be that Ψ itself is inconsistent, in which case it would entail anything, particularly $\Psi \models \overline{A}$ would be fulfilled. The handling of an inconsistent prior belief state is one of the crucial differences between *revision* and *update*, as characterized in [20] by the so-called *KM-postulates*. An AGM-revision demands $\Psi * A$ to be consistent, regardless if the prior state Ψ is inconsistent or not, whereas update does not remedy the inconsistence of a prior state, even if the new information is consistent. So (CR7) would be trivially fulfilled for KM-updates. If we also give up the postulate of success, then (CR7) describes a reasonable behavior of a revision process in an extreme case: A revision should not establish the negation of the revising proposition if this negated proposition is not already implied by the prior belief state.

The following theorem shows that the postulates (CR0)-(CR7) cover the DP-postulates (C1)-(C4):

Theorem 5.1. *Suppose $*$ is a conditional revision operator obeying the postulates* (CR0)-(CR7). *Then for the induced propositional revision operator, postulates* (C1)-(C4) *are satisfied, too.*

Therefore, the idea of conditional preservation inherent to the postulates (C1)-(C4) of Darwiche and Pearl ([8]) is indeed captured by our postulates. While (CR0) - (CR4) only serve as basic, unspecific postulates, the last three postulates (CR5)-(CR7) can be taken as properly axiomatizing a *principle of conditional preservation in a qualitative framework.* Moreover, our framework provides further, formal justifications for the DP-postulates by making interactions of conditionals more precise.

6 Conditional structures and conditional indifference

The notion of *conditional structures* has been presented and exemplified in several papers (see, e.g., [25, 24, 28]). Since they are basic to the results to be obtained in this paper, we will summarize the main ideas and definitions here. The concept of *conditional indifference* has also been a major topic in [25]; in the present paper, however, it is developed in the general framework of conditional valuation functions.

In Section 4, we presented a dynamic approach to conditionals, focusing on the effects of only one conditional in the revision process. When considering sets $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ of conditionals, the effects each of these conditionals exerts on worlds must be clearly identified. To this end, we replace the numbers 0 and 1 in (2) by formal symbols, one pair of symbols $\mathbf{a}_i^+, \mathbf{a}_i^-$ for each conditional $(B_i|A_i)$ in $\mathcal{R}$, giving rise to the functions $\sigma_i = \sigma_{(B_i|A_i)}$, $1 \leqslant i \leqslant n$:

$$\sigma_i(\omega) = \begin{cases} \mathbf{a}_i^+ & \text{if} \quad (B_i|A_i)(\omega) = 1 \\ \mathbf{a}_i^- & \text{if} \quad (B_i|A_i)(\omega) = 0 \\ 1 & \text{if} \quad (B_i|A_i)(\omega) = u \end{cases} \tag{3}$$

$\sigma_i(\omega)$ represents the effects each conditional $(B_i|A_i)$ has on possible worlds ω. To make these conditional effects computable, we make use of a group structure, introducing the *free abelian group* $\mathcal{F}_\mathcal{R} = \langle \mathbf{a}_1^+, \mathbf{a}_1^-, \ldots, \mathbf{a}_n^+, \mathbf{a}_n^- \rangle$ 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 $\mathbb{Z}^{2n}$ (cf. [32, 13]). The commutativity of $\mathcal{F}_\mathcal{R}$ corresponds to a simultaneous application of the conditionals in $\mathcal{R}$, without assuming any order of application. Note that the neutral element 1 of $\mathcal{F}_\mathcal{R}$ is assigned to possible worlds in the neutral sets of the conditionals. The function $\sigma_\mathcal{R} = \prod_{1 \leqslant i \leqslant n} \sigma_i : \Omega \to \mathcal{F}_\mathcal{R}$, given by

$$\sigma_\mathcal{R}(\omega) = \prod_{1 \leqslant i \leqslant n} \sigma_i(\omega) = \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \mathbf{a}_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} \mathbf{a}_i^- \tag{4}$$

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 following simple example illustrates the notion of conditional structures and shows how to calculate in this framework:

Example 6.1. 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	$\sigma_\mathcal{R}(\omega)$	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

We find that $\dfrac{\sigma_\mathcal{R}(abc)\sigma_\mathcal{R}(\overline{a}\overline{b}c)}{\sigma_\mathcal{R}(a\overline{b}c)\sigma_\mathcal{R}(\overline{a}bc)} = \dfrac{\mathbf{a}_1^+\mathbf{a}_2^+ \cdot 1}{\mathbf{a}_1^+ \cdot \mathbf{a}_2^+} = 1$, which may be interpreted by stating that the sets of worlds $\{abc, \overline{a}\overline{b}c\}$ and $\{a\overline{b}c, \overline{a}bc\}$ are balanced with respect to the effects of the conditionals in $\mathcal{R}$.

To comply with the group theoretical structure of $\mathcal{F}_\mathcal{R}$, we also impose a multiplication on Ω, introducing the free abelian group $\widehat{\Omega} := \langle \omega \mid \omega \in \Omega \rangle$ generated by all $\omega \in \Omega$, and consisting of all products $\widehat{\omega} = \omega_1{}^{r_1} \ldots \omega_m{}^{r_m}$ with $\omega_1, \ldots, \omega_m \in \Omega$, and integers $r_1, \ldots r_m$. Now $\sigma_\mathcal{R}$ may be extended to $\widehat{\Omega}$ in a straightforward manner by setting

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

yielding a *homomorphism of groups* $\sigma_\mathcal{R} : \widehat{\Omega} \to \mathcal{F}_\mathcal{R}$. We will often use fractional representations for the elements of $\widehat{\Omega}$, that is, for instance, we will write $\dfrac{\omega_1}{\omega_2}$ instead of $\omega_1\omega_2^{-1}$. Having the same conditional structure defines an equivalence relation $\equiv_\mathcal{R}$ on $\widehat{\Omega}$: $\widehat{\omega}_1 \equiv_\mathcal{R} \widehat{\omega}_2$ iff $\sigma_\mathcal{R}(\widehat{\omega}_1) = \sigma_\mathcal{R}(\widehat{\omega}_2)$ The equivalence classes are in one-to-one correspondence to the elements of the quotient group $\widehat{\Omega}/_{ker\ \sigma_\mathcal{R}} = \{\widehat{\omega} \cdot (ker\ \sigma_\mathcal{R}) \mid \widehat{\omega} \in \widehat{\Omega}\}$, where $ker\ \sigma_\mathcal{R} := \{\widehat{\omega} \in \widehat{\Omega} \mid \sigma_\mathcal{R}(\widehat{\omega}) = 1\}$ denotes the kernel of the homomorphism $\sigma_\mathcal{R}$, since $\widehat{\omega}_1 \equiv_\mathcal{R} \widehat{\omega}_2$ iff $\widehat{\omega}_1\widehat{\omega}_2^{-1} \in ker\ \sigma_\mathcal{R}$. Therefore, the kernel plays an important role in identifying conditional structures. It contains exactly all group elements $\widehat{\omega} \in \widehat{\Omega}$ with a balanced conditional structure, that means, where all effects of conditionals in $\mathcal{R}$ on worlds occurring in $\widehat{\omega}$ are completely cancelled. For instance, in Example 6.1 above, the element $\dfrac{abc \cdot \overline{a}\overline{b}c}{a\overline{b}c \cdot \overline{a}bc}$ is an element of the kernel of $\sigma_\mathcal{R}$.

Besides the conditional information in $\mathcal{R}$ (or $\mathcal{R}^*$, if one is concerned with quantified conditionals), one usually has to take normalization constraints such as $P(\top) = 1$ for probability distributions P, or $\kappa(\top) = 0$ for ordinal conditional functions κ, or $\pi(\top) = 1$ for possibility distributions π, into regard. This is done by focusing on the subgroup

$$\widehat{\Omega}_0 = ker\,\sigma_{(\top|\top)} = \{\widehat{\omega} = {\omega_1}^{r_1} \cdot \ldots \cdot {\omega_m}^{r_m} \in \widehat{\Omega} \mid \sum_{j=1}^{m} r_j = 0\}$$

of $\widehat{\Omega}$. 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 *equivalent modulo* $\widehat{\Omega}_0$, $\widehat{\omega}_1 \equiv_\top \widehat{\omega}_2$, iff $\widehat{\omega}_1\widehat{\Omega}_0 = \widehat{\omega}_2\widehat{\Omega}_0$, iff $\sum_{1 \leqslant j \leqslant m} r_j = \sum_{1 \leqslant k \leqslant p} s_k$. This means that $\widehat{\omega}_1 \equiv_\top \widehat{\omega}_2$ iff they both are a (cancelled) product of the same number of generators ω, each generator being counted with its corresponding exponent. Let $ker_0\,\sigma_{\mathcal{R}} := ker\,\sigma_{\mathcal{R}} \cap \widehat{\Omega}_0$ be the part of $ker\,\sigma_{\mathcal{R}}$ which is included in $\widehat{\Omega}_0$.

Finally, we will show how to describe the relations $\sqsubseteq$ and $\perp\!\!\!\perp$ between conditionals, introduced in Definitions 4.1 and 4.2, respectively, by considering the kernels of the corresponding σ-homomorphisms. As a convenient notation, for each proposition $A \in \mathcal{L}$, we define $\widehat{A} := \{\widehat{\omega} = \omega_1^{r_1} \ldots \omega_m^{r_m} \in \widehat{\Omega} \mid \omega_i \models A$ for all $i, 1 \leqslant i \leqslant m\}$.

Proposition 6.1. *Let $(B|A), (D|C) \in (\mathcal{L} \mid \mathcal{L})$ be conditionals.*

1. *$(D|C)$ is either a subconditional of $(B|A)$ or of $(\overline{B}|A)$ iff $C \leqslant A$ and $ker\,\sigma_{(D|C)} \cap \widehat{C} \subseteq ker\,\sigma_{(B|A)} \cap \widehat{C}$.*
2. *$(D|C) \perp\!\!\!\perp (B|A)$ iff $\widehat{C} \cap \widehat{\Omega}_0 \subseteq ker\,\sigma_{(B|A)}$.*

To study conditional interactions, we now focus on the behavior of conditional valuation functions $V : \mathcal{L} \to \mathcal{A}$ with respect to the "multiplication" $\odot$ in $\mathcal{A}$ (see Definition 3.1). Each such function may be extended to a homomorphism $V : \widehat{\Omega}_+ \to (\mathcal{A}, \odot)$ by setting $V({\omega_1}^{r_1} \cdot \ldots \cdot {\omega_m}^{r_m}) = V(\omega_1)^{r_1} \odot \ldots \odot V(\omega_m)^{r_m}$, where $\widehat{\Omega}_+$ is the subgroup of $\widehat{\Omega}$ generated by the set $\Omega_+ := \{\omega \in \Omega \mid V(\omega) \neq 0^{\mathcal{A}}\}$. This allows us to analyze numerical relationships holding between different $V(\omega)$. Thereby, it will be possible to elaborate the conditionals whose structures V follows, that means, to determine sets of conditionals $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ with respect to which V is *indifferent*:

Definition 6.1 (indifference wrt $\mathcal{R}$). *Suppose $V : \mathcal{L} \to \mathcal{A}$ is a conditional valuation function and $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ is a set of conditionals such that $V(A) \neq 0^{\mathcal{A}}$ for all $(B|A) \in \mathcal{R}$.*

V is indifferent with respect to $\mathcal{R}$ *iff the following two conditions hold:*

(i) If $V(\omega) = 0^{\mathcal{A}}$ then there is $(B|A) \in \mathcal{R}$ such that $V(\omega') = 0^{\mathcal{A}}$ for all ω' with $\sigma_{(B|A)}(\omega') = \sigma_{(B|A)}(\omega)$.

(ii) $V(\widehat{\omega}_1) = V(\widehat{\omega}_2)$ 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 V is indifferent with respect to $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$, then it does not distinguish between different elements $\widehat{\omega}_1, \widehat{\omega}_2$ which are equivalent modulo $\widehat{\Omega}_0$ and have the same conditional structure with respect to $\mathcal{R}$. Conversely, for each $\widehat{\omega} \in \widehat{\Omega}_0$, any deviation $V(\widehat{\omega}) \neq 1^{\mathcal{A}}$ can be explained by the conditionals in $\mathcal{R}$ acting on $\widehat{\omega}$ in a non-balanced way. Condition (i) in Definition 6.1 is necessary to deal with worlds $\omega \notin \Omega_+$. It says that $0^{\mathcal{A}}$-values in an indifferent valuation function V are established only in according with the partitionings induced by the conditionals in $\mathcal{R}$.

The following proposition rephrases conditional indifference by establishing a relationship between the kernels of $\sigma_{\mathcal{R}}$ and V:

Proposition 6.2. *Let $\mathcal{R} \subseteq (\mathcal{L} \mid \mathcal{L})$ be a set of conditionals, and let $V : \mathcal{L} \to \mathcal{A}$ be a conditional valuation function with $V(A) \neq 0^{\mathcal{A}}$ for all $(B|A) \in \mathcal{R}$.*

V is indifferent with respect to $\mathcal{R}$ iff condition (i) of Definition 6.1 holds, and $ker_0\,\sigma_{\mathcal{R}} \cap \widehat{\Omega}_+ \subseteq ker_0\,V$.

We will close this section by characterizing probability functions, ordinal conditional functions and possibility distributions with indifference properties:

Theorem 6.1. *Let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subset (\mathcal{L} \mid \mathcal{L})$ be a (finite) set of conditionals.*

1. *A probability function P is indifferent with respect to $\mathcal{R}$ iff $P(A_i) \neq 0$ for all $i, 1 \leqslant i \leqslant n$, and there are non-negative real numbers $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+, \alpha_0 > 0$ such that, for all $\omega \in \Omega$,*
$$P(\omega) = \alpha_0 \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{5}$$
2. *An ordinal conditional function κ is indifferent with respect to $\mathcal{R}$ 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 $\omega \in \Omega$,*
$$\kappa(\omega) = \kappa_0 + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^- \tag{6}$$
3. *A possibility distribution π is indifferent with respect to $\mathcal{R}$ iff there are non-negative real numbers $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+, \alpha_0 > 0$, such that for all $\omega \in \Omega$,*
$$\pi(\omega) = \alpha_0 \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{7}$$

Note that conditional indifference is a structural notion, without making any reference to degrees of certainty which may be assigned to the conditionals in $\mathcal{R}$. Theorem 6.1, however, also provides simple schemata how to obtain indifferent probabilistic, OCF and possibilistic representations of quantified conditionals: One has to simply set up functions of the corresponding type according to (5), (6) or (7), respectively, and to determine the constants $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^-$ or $\kappa_0, \kappa_i^+, \kappa_i^-$, respectively, appropriately so as to ensure that all necessary numerical relationships are established. Conditional valuation functions which represent a set $\mathcal{R}^{(*)}$ of (quantified) conditionals and are indifferent to it, are called *c-representations of* $\mathcal{R}^{(*)}$ (for further details and examples, cf. [21, 25, 29]; see also Section 7).

Theorem 6.1 also shows, that most important and well-behaved inductive representation methods realize conditional indifference: Namely, the principle of maximum entropy in probabilistics [34], system-Z^* in the OCF-framework [17], and the LCD-functions of Benferhat, Saffiotti and Smets [4] all give rise to conditionally indifferent functions (cf. [21, 25, 28]). The system-Z^* approach and that of LCD-functions can easily be derived by postulating conditional indifference and further plausibility assumptions (for a more detailed discussion, cf. [28]). Indeed, the crucial meaning of all these formalisms for adequate inductive knowledge representation is mainly due to this indifference property. It should be emphasized, that, to study interactions of conditionals, conditionals here are not reduced to material implications, as for system-Z^*, or for LCD-functions. Instead, the full dynamic, non-classical power of conditionals is preserved, and highly complex conditional interactions can be dealt with.

Therefore, the theory of conditional structures and conditional indifference presented so far proves to be of fundamental importance both for theoretical and practical issues in inductive knowledge representation. In the next section, we will show that it also provides an appropriate framework for revising quantified beliefs.

7 A principle of conditional preservation in a (semi-)quantitative framework

When we revise an epistemic state Ψ – which is supposed to be represented by a conditional valuation function V – by a set of (quantified) conditionals $\mathcal{R}$ to obtain a posterior epistemic state $\Psi * \mathcal{R} \equiv V^* = V * \mathcal{R}$, conditional structures and/or interactions must be observed with respect to the prior state Ψ as well as to the new conditionals in $\mathcal{R}$. The theory of conditional structures can only be applied with respect to $\mathcal{R}$, since we usually do not know anything about the history of Ψ, or V, respectively. Conditional relationships within Ψ, however, are realized via the operation $\odot$ of V, so we base our definition of a principle of conditional preservation on an indifference property of the *relative change function* $V^* \odot V^{-1}$, in the following written as V^*/V. Taking into regard prior knowledge V and the worlds ω with $V(\omega) = 0^{\mathcal{A}}$ appropriately, this gives rise to the following definitions:

Definition 7.1 (V-consistency, indifference wrt $\mathcal{R}$ and V). *Let $V : \mathcal{L} \to \mathcal{A}$ be a conditional valuation function, and let $\mathcal{R}$ be a finite set of (quantified) conditionals. Let $V^* = V * \mathcal{R}$ denote the result of revising V by $\mathcal{R}$; in particular, suppose that $V^*(A) \neq 0^{\mathcal{A}}$ for all $(B|A) \in \mathcal{R}$.*

1. *V^* is called* V-consistent *iff $V(\omega) = 0^{\mathcal{A}}$ implies $V^*(\omega) = 0^{\mathcal{A}}$; V^* is called* strictly V-consistent *iff $V(\omega) = 0^{\mathcal{A}} \Leftrightarrow V^*(\omega) = 0^{\mathcal{A}}$;*
2. *If V^* is V-consistent, then the* relative change function *$(V^*/V) : \Omega \to \mathcal{A}$ is defined by*

$$(V^*/V)(\omega) = \begin{cases} V^*(\omega) \odot V(\omega)^{-1} & \text{if} \quad V(\omega) \neq 0^{\mathcal{A}} \\ 0^{\mathcal{A}} & \text{if} \quad V(\omega) = 0^{\mathcal{A}} \end{cases}$$

3. *V^* is* indifferent with respect to $\mathcal{R}$ and V *iff V^* is V-consistent and the following two conditions hold:*
   (i) If $V^*(\omega) = 0^{\mathcal{A}}$ then $V(\omega) = 0^{\mathcal{A}}$, or there is $(B|A) \in \mathcal{R}$ such that $V^*(\omega') = 0^{\mathcal{A}}$ for all ω' with $\sigma_{(B|A)}(\omega') = \sigma_{(B|A)}(\omega)$.
   (ii) $(V^*/V)(\widehat{\omega}_1) = (V^*/V)(\widehat{\omega}_2)$ whenever $\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \sigma_{\mathcal{R}}(\widehat{\omega}_2)$ and $\widehat{\omega}_1 \widehat{\Omega}_0 = \widehat{\omega}_2 \widehat{\Omega}_0$ for $\widehat{\omega}_1, \widehat{\omega}_2 \in \widehat{\Omega}^*_+$, where $\widehat{\Omega}^*_+ = \langle \omega \in \Omega \mid V^*(\omega) \neq 0^{\mathcal{A}} \rangle$.

Although the relative change function (V^*/V) is not a conditional valuation function, it may nevertheless be extended to a homomorphism $(V^*/V) : \widehat{\Omega}^*_+ \to (\mathcal{A}, \odot)$. Therefore, Definition 7.1 is an appropriate modification of Definition 6.1 for revisions.

We are now ready to formalize appropriately a principle of conditional preservation for belief revision in a (semi-)quantitative framework:

Definition 7.2 (principle of conditional preservation wrt $\mathcal{R}$ and V). *A revision $V^* = V * \mathcal{R}$ of a conditional valuation function by a set $\mathcal{R}$ of (quantified) conditionals is said to satisfy the* principle of conditional preservation with respect to $\mathcal{R}$ and V *iff V^* is indifferent with respect to $\mathcal{R}$ and V.*

Thus in a numerical framework, the principle of conditional preservation is realized as an indifference property.

From Theorem 6.1, we immediately obtain a concise characterization of revisions preserving conditional beliefs, which may also serve in practice as a schema to set up appropriate revision formalisms:

Theorem 7.1. *Let $\mathcal{R} = \{(B_1|A_1)([x_1]), \ldots, (B_n|A_n)([x_n])\} \subset (\mathcal{L} \mid \mathcal{L})$ be a (finite) set of (quantified) conditionals. Let P be a probability distribution, κ an ordinal conditional function, and π a possibility distribution, all serving as prior knowledge.*

1. *A probability distribution* $P^* = P * \mathcal{R}$ *satisfies the principle of conditional preservation with respect to* $\mathcal{R}$ *and* P *if and only if there are real numbers* $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^-$ *with* $\alpha_0 > 0$ *and* $\alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^-$ *satisfying* $\alpha_i^+, \alpha_i^- \geqslant 0$, $\alpha_i^+ = 0$ *iff* $x_i = 0$, $\alpha_i^- = 0$ *iff* $x_i = 1$, $1 \leqslant i \leqslant n$, *such that, for all* $\omega \in \Omega$,
$$P^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{8}$$
2. *A revision* $\kappa^* = \kappa * \mathcal{R}$ *satisfies the principle of conditional preservation with respect to* $\mathcal{R}$ *and* κ *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* $\omega \in \Omega$,
$$\kappa^*(\omega) = \kappa(\omega) + \kappa_0 + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^- \tag{9}$$
3. *A revision* $\pi^* = \pi * \mathcal{R}$ *satisfies the principle of conditional preservation with respect to* $\mathcal{R}$ *and* π *iff there are positive real numbers* $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+$ *such that for all* $\omega \in \Omega$,
$$\pi^*(\omega) = \alpha_0 \pi(\omega) \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{10}$$

Note that the principle of conditional preservation is based only on observing conditional structures, without using any acceptance conditions or taking quantifications of conditionals into account. It is exactly this separation of numerical from structural aspects that results in a wide applicability of this principle within a quantitative framework. Revisions of epistemic states Ψ by sets $\mathcal{R}$ of (quantified) conditionals that also fulfill the so-called *success postulate* $\Psi * \mathcal{R} \models \mathcal{R}$ are termed *c-revisions*:

Definition 7.3 (c-revision). *A revision* $V^* = V * \mathcal{R}$ *of a conditional valuation function by a set* $\mathcal{R}$ *of (quantified) conditionals is called a* c-revision *iff* V^* *satisfies the principle of conditional preservation with respect to* V *and* $\mathcal{R}$, *and* $V^* \models \mathcal{R}$.

C-revisions can easily be obtained by using the schemata provided by Theorem 7.1 and choosing the constants $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^-$, and $\kappa_0, \kappa_i^+, \kappa_i^-$, respectively, appropriately so as to establish the necessary numerical relationships. For instance, for ordinal conditional functions, a c-revision $\kappa^* = \kappa * \mathcal{R}$ of an OCF κ by $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ has the form (9), and the postulate $\kappa^* \models \mathcal{R}$ yields the following conditions for κ_i^+, κ_i^- in a straightforward way:
$$\begin{aligned} \kappa_i^- - \kappa_i^+ > \min_{\omega \models A_i B_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \\ - \min_{\omega \models A_i \overline{B}_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \end{aligned} \tag{11}$$
Moreover, quantifications of conditionals can be taken easily into account by modifying (11) slightly, so as to comply with the representation postulate $\kappa^* \models (B|A)\,[m_i]$:
$$\begin{aligned} \kappa_i^- - \kappa_i^+ > m_i + \min_{\omega \models A_i B_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \\ - \min_{\omega \models A_i \overline{B}_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \end{aligned} \tag{12}$$
We will illustrate this by an example.

Example 7.1. Let the ordinal conditional function κ, as shown in Table 7.1, represent epistemic knowledge about the atoms f - *flying*, b - *birds*, p - *penguins*, w - winged animals, and k - *kiwis*. Among others, κ represents the conditionals $(f|b)$

ω	$\kappa(\omega)$	$\kappa^*(\omega)$	ω	$\kappa(\omega)$	$\kappa^*(\omega)$	ω	$\kappa(\omega)$	$\kappa^*(\omega)$	ω	$\kappa(\omega)$	$\kappa^*(\omega)$
$pbfwk$	2	3	$\overline{p}bfwk$	0	1	$pbfw\overline{k}$	2	2	$\overline{p}bfw\overline{k}$	0	0
$pbf\overline{w}k$	3	3	$\overline{p}bf\overline{w}k$	1	1	$pbf\overline{w}\overline{k}$	3	3	$\overline{p}bf\overline{w}\overline{k}$	1	1
$pb\overline{f}wk$	1	2	$\overline{p}b\overline{f}wk$	1	2	$pb\overline{f}w\overline{k}$	1	1	$\overline{p}b\overline{f}w\overline{k}$	1	1
$pb\overline{f}\overline{w}k$	2	2	$\overline{p}b\overline{f}\overline{w}k$	2	2	$pb\overline{f}\overline{w}\overline{k}$	2	2	$\overline{p}b\overline{f}\overline{w}\overline{k}$	2	2
$p\overline{b}fwk$	5	6	$\overline{p}\overline{b}fwk$	1	2	$p\overline{b}fw\overline{k}$	4	4	$\overline{p}\overline{b}fw\overline{k}$	0	0
$p\overline{b}f\overline{w}k$	5	5	$\overline{p}\overline{b}f\overline{w}k$	1	1	$p\overline{b}f\overline{w}\overline{k}$	4	4	$\overline{p}\overline{b}f\overline{w}\overline{k}$	0	0
$p\overline{b}\,\overline{f}wk$	3	4	$\overline{p}\overline{b}\,\overline{f}wk$	1	2	$p\overline{b}\,\overline{f}w\overline{k}$	2	2	$\overline{p}\overline{b}\,\overline{f}w\overline{k}$	0	0
$p\overline{b}\,\overline{f}\overline{w}k$	3	3	$\overline{p}\overline{b}\,\overline{f}\overline{w}k$	1	1	$p\overline{b}\,\overline{f}\overline{w}\overline{k}$	2	2	$\overline{p}\overline{b}\,\overline{f}\overline{w}\overline{k}$	0	0

Fig. 1. OCF κ, and revised $\kappa^* = \kappa * \{(\overline{w}|k)\}$ for Example 7.1

(*birds fly*), $(b|p)$ (*penguins are birds*), $(\overline{f}|p)$ (*penguins do not fly*), $(w|b)$ (*birds have wings*), and $(b|k)$ (*kiwis are birds*). This is checked easily by comparing the rankings $\kappa(\omega)$ of verifying and falsifying worlds ω, respectively, for each of the conditionals. (Actually, κ is a c-representation of these conditionals.)

Moreover, we have $\kappa \models (w|k)$ – from their superclass *birds*, kiwis inherit the property of having wings. Suppose now that we come to know that this is false – kiwis do *not* possess wings, – and we want to revise our knowledge κ by this new information. The revised epistemic state $\kappa^* = \kappa * \{(\overline{w}|k)\}$ should be a c-revision of κ by $\{(\overline{w}|k)\}$. Then due to (9), κ^* has the form

$$\kappa^*(\omega) = \begin{cases} \kappa_0 + \kappa(\omega) + \kappa^+ & \text{if } \omega \models k\overline{w} \\ \kappa_0 + \kappa(\omega) + \kappa^- & \text{if } \omega \models kw \\ \kappa_0 + \kappa(\omega) & \text{if } \omega \models \overline{k} \end{cases}$$

and (11) yields $\kappa^- - \kappa^+ > \min\limits_{\omega \models k\overline{w}} \kappa(\omega) - \min\limits_{\omega \models kw} \kappa(\omega) = 1 - 0 = 1$, i.e. $\kappa^- > \kappa^+ + 1$. Any such pair of κ^+, κ^- will give rise to a c-revision, but, in order to keep numerical changes minimal, we choose $\kappa^+ := 0, \kappa^- := 1$. No further normalization is necessary, so $\kappa_0 := 0$. The revised κ^* is shown in Table 7.1, too[1]. κ^* still represents the conditionals $(f|b), (b|p), (\overline{f}|p)$ and $(w|b)$, but it no longer satisfies $(b|k)$, since $\kappa^*(bk) = \kappa^*(\overline{b}k) = 1$ – the property of not having wings casts (reasonably) doubt on kiwis being birds. Note that no external knowledge, such as "*typically*, birds have wings", and no quantifications of the inferential strength of conditionals is needed to bring forth this reasonable behavior. So, this example illustrates strikingly, how adequately conditional interrelationships are dealt with by c-revisions.

In a probabilistic framework, Theorem 7.1 proves that revisions following the *principle of minimum cross-entropy* (so-called *MINENT-principle* or, briefly, *ME-principle*) [37, 35, 34, 21] are also c-revisions. This principle is a method to revise a prior distribution P by a set $\mathcal{R}^* = \{(B_1|A_1)\,[x_1], \ldots, (B_n|A_n)\,[x_n]\}$ of probabilistic conditionals, so that the "dissimilarity" between P and the resulting distribution $P^* \models \mathcal{R}^*$ is minimal. A measure for this dissimilarity is given by the information-

[1] κ^* can also be regarded as the result of an *update process*, following evolution.

theoretical concept of *cross-entropy* $R(Q,P) = \sum_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)}$. If $\mathcal{R}^*$ is compatible with the prior P, in the sense that there is a P-consistent distribution Q representing $\mathcal{R}^*$, this optimization problem has a unique solution $P^* = P *_{ME} \mathcal{R}^*$ (cf. [7]), which can be written in the form

$$P^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^{1-x_i} \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^{-x_i} \tag{13}$$

with the α_i's being exponentials of the Lagrange multipliers, appropriately chosen so as to satisfy all conditionals in $\mathcal{R}^*$ (cf. [27]). Comparing (13) to (8), it is obvious that $P *_{ME} \mathcal{R}^*$ satisfies the principle of conditional preservation, and hence is a c-revision.

By Definition 7.2, we obtain a technically clear and precise formalization of the intuitive idea of conditional preservation in a very general framework, making it applicable to probabilistic, possibilistic and OCF-revisions. Note that, as abstract and technical as it appears, this principle is not a formal artifact but has been effectively guiding probabilistic revisions via the principle of minimum cross-entropy for many decades. Indeed, the first steps towards formalizing this principle have been taken when extracting the most basic and crucial properties of minimum cross-entropy methods in [21]. Therefore, the axiomatization provided by Definition 7.2 allows us to carry over a most successful information-theoretical idea from probabilistics to other frameworks when designing adequate revision methods. No explicit reference to ME-probability distributions is needed, as was done for system-Z^* (cf. [17]). Moreover, generalizations of system-Z^* and LCD-functions for belief revision are straightforward and can be theoretically justified within the corresponding framework.

8 Linking qualitative and quantitative approaches

In Sections 5 and 7, the idea of preserving conditional beliefs under revision have been formalized in two (apparently) different ways: In Section 5, we made use of the two relations $\sqsubseteq$ and $\perp\!\!\!\perp$, describing quite simple ways of conditional interactions. In Section 7, we based our formalization upon observing conditional structures. In any case, the principal idea was to focus on conditional (not logical) interactions, considering the effects conditionals may exert when being established. We will now show, that both approaches essentially coincide in the case that a conditional valuation function (as a quantitative representation of epistemic beliefs, like e.g. ordinal conditional functions or possibility distributions) is revised by only one conditional. More exactly, we will prove that a revision following the quantitative principle of conditional preservation (see Definition 7.2 in Section 7) satisfy the postulates (CR5)-(CR7) in Section 5, describing a qualitative principle of conditional preservation.

We begin by characterizing revisions $V^* = V * \mathcal{R} = V * (B|A)$ of a conditional valuation function V which satisfy the (quantitative) principle of conditional preservation with respect to $\mathcal{R} = \{(B|A)\}$ and V. As a basic requirement for such revisions, we will only presuppose that $V^*(A) \neq 0^{\mathcal{A}}$, instead of the (stronger) success postulate $V^* \models (B|A)$. This makes the results to be presented independent of acceptance conditions and helps concentrating on conditional structures; in particular, it will be possible to make use of these results even when conditionals are assigned numerical degrees of acceptance.

Proposition 8.1. *Let $V : \mathcal{L} \to \mathcal{A}$ be a conditional valuation function, and let $\mathcal{R} = \{(B|A)\}, (B|A) \in (\mathcal{L} \mid \mathcal{L})$, consist of only one conditional. Let $V^* = V * \mathcal{R} =$*

$V * (B|A)$ denote a revision of V by $(B|A)$ such that $V^*(A) \neq 0^{\mathcal{A}}$. V^* satisfies the principle of conditional preservation with respect to V and $\mathcal{R}$ iff there are constants $\alpha_0, \alpha^+, \alpha^- \in \mathcal{A}$ such that

$$V^*(\omega) = \begin{cases} \alpha^+ \odot V(\omega) & \textit{if} \quad \omega \models AB \\ \alpha^- \odot V(\omega) & \textit{if} \quad \omega \models A\overline{B} \\ \alpha_0 \odot V(\omega) & \textit{if} \quad \omega \models \overline{A} \end{cases} \tag{14}$$

If V^* is strictly V-consistent, then all constants $\alpha_0, \alpha^+, \alpha^- \in \mathcal{A}$ may be chosen $\neq 0^{\mathcal{A}}$.

As an obvious link between the qualitative and the quantitative frameworks, we now strengthen the central postulate (CR5) to comply with the numerical information provided by conditional valuation functions V:

(CR5^{quant}) If $(D|C) \perp\!\!\!\perp (B|A)$ and $V(CD), (V * (B|A))(CD) \neq 0^{\mathcal{A}}$, then

$$V(C\overline{D}) \odot V(CD)^{-1} = (V * (B|A))(C\overline{D}) \odot (V * (B|A))(CD)^{-1}.$$

(CR5^{quant}) ensures that essentially, the values assigned to conditionals which are perpendicular to the revising conditional are not changed under revision:

Lemma 8.1. *Suppose the revision $V * (B|A)$ is strictly V-consistent and satisfies (CR5^{quant}). Then for any conditional $(D|C) \perp\!\!\!\perp (B|A)$ with $V(C) \neq 0^{\mathcal{A}}$, it holds that $V(D|C) = (V * (B|A))(D|C)$.*

The next proposition shows that indeed, (CR5^{quant}) is stronger than its qualitative counterpart (CR5):

Proposition 8.2. *Let $V^* = V * \mathcal{R} = V * \{(B|A)\}$ denote a strictly V-consistent revision of V by $(B|A)$ such that $V^*(A) \neq 0^{\mathcal{A}}$. If V^* fulfills (CR5^{quant}), then it also satisfies (CR5).*

The following theorem states that essentially, any revision of a conditional valuation function which satisfies the quantitative principle of conditional preservation (as specified by Definition 7.2), is also in accordance with the qualitative principle of conditional preservation (as described by (CR5)-(CR7)):

Theorem 8.1. *Let $V : \mathcal{L} \to \mathcal{A}$ be a conditional valuation function, and let $\mathcal{R} = \{(B|A)\}, (B|A) \in (\mathcal{L} \mid \mathcal{L})$, consist of only one conditional. Let $V^* = V * \mathcal{R}$ denote a strictly V-consistent revision of V by $\mathcal{R}$ fulfilling the postulates (CR1) (success) and (CR2) (stability).*

If V^* satisfies the principle of conditional preservation, then the revision also satisfies postulate (CR5^{quant}) and the postulates (CR6) and (CR7); in particular, it satisfies all of the postulates (CR5)-(CR7).

Therefore, Theorem 8.1 identifies the principle of conditional preservation, as formalized in Definition 7.2, as a fundamental device to guide reasonable changes in the conditional structure of knowledge.

9 Conclusion

In this paper, we presented axiomatizations of a principle of conditional preservation for belief revision operations in qualitative as well as in (semi-)quantitative settings. In both cases, we dealt with revisions of epistemic states by sets of conditional beliefs, thus studying belief revision in a most general framework. As the inductive representation of a set of conditionals (or default rules, respectively) can be considered as a special instance of a revision problem, this paper also provides an approach for adequate knowledge induction.

The crucial point in preserving conditional beliefs is to observe conditional interactions, which can be described by two relations, *subconditionality* and *perpendicularity*, in the qualitative framework, and are based on the algebraic notion of conditional structures in the quantitative framework. Since subconditionality and perpendicularity can also be defined via conditional stuctures, the theory of conditional structures developed in this paper proves to be a most basic and powerful tool for handling conditionals in knowledge representation and belief revision. We applied this theory to *conditional valuation functions* as basic representations of (semi-) quantitative epistemic states, covering probability distributions, ranking functions (ordinal conditional functions), and possibility distributions. Therefore, the results presented in this paper are of relevance for a wide range of revision problems in very different environments. Moreover, apart from theoretical aspects, our approach also yields practical schemata for setting up revision and representation operations in probabilistic, possibilistic and ordinal frameworks.

As the main result of this paper, we showed that the quantitative principle of conditional preservation implies the qualitative principle in semi-quantitative settings. This not only closes the gap between qualitative and quantitative approaches to belief revision, but also may give new impetus to classical belief revision theory.

References

1. C.E. Alchourrón, P. Gärdenfors, and P. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
2. S. Benferhat, D. Dubois, and H. Prade. Representing default rules in possibilistic logic. In *Proceedings 3th International Conference on Principles of Knowledge Representation and Reasoning KR'92*, pages 673–684, 1992.
3. S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence*, 92:259–276, 1997.
4. S. Benferhat, A. Saffiotti, and P. Smets. Belief functions and default reasoning. *Artificial Intelligence*, 122:1–69, 2000.
5. C. Boutilier and M. Goldszmidt. Revision by conditional beliefs. In *Proceedings 11th National Conference on Artificial Intelligence (AAAI'93)*, pages 649–654, Washington, DC., 1993.
6. P.G. Calabrese. Deduction and inference using conditional logic and probability. In I.R. Goodman, M.M. Gupta, H.T. Nguyen, and G.S. Rogers, editors, *Conditional Logic in Expert Systems*, pages 71–100. Elsevier, North Holland, 1991.
7. I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Ann. Prob.*, 3:146–158, 1975.
8. A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
9. B. DeFinetti. *Theory of Probability*, volume 1,2. John Wiley and Sons, New York, 1974.
10. J.P. Delgrande and T. Schaub. A consistency-based model for belief change: Preliminary report. In *Proceedings AAAI-2000*, pages 392–398, Austin, TX, 2000.
11. D. Dubois, J. Lang, and H. Prade. Possibilistic logic. In D.M. Gabbay, C.H. Hogger, and J.A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford University Press, 1994.
12. D. Dubois and H. Prade. A survey of belief revision and updating rules in various uncertainty models. *Intern. Journal of Intelligent Systems*, 9:61–100, 1994.
13. B. Fine and G. Rosenberger. *Algebraic Generalizations of Discrete Groups.* Dekker, New York, Basel, 1999.
14. M. Freund. Preferential orders and plausibility measures. *J. Logic Computat.*, 8:147–158, 1998.

15. N. Friedman and J.Y. Halpern. Plausibility measures and default reasoning. In *Proceedings 13th National Conference on Artificial Intelligence, AAAI-96*, volume 2, 1996.
16. P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge, Mass., 1988.
17. M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):220–232, 1993.
18. M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84:57–112, 1996.
19. H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.
20. H. Katsuno and A.O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings Second International Conference on Principles of Knowledge Representation and Reasoning, KR'91*, pages 387–394, San Mateo, Ca., 1991. Morgan Kaufmann.
21. G. Kern-Isberner. Characterizing the principle of minimum cross-entropy within a conditional-logical framework. *Artificial Intelligence*, 98:169–208, 1998.
22. G. Kern-Isberner. Postulates for conditional belief revision. In *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI-99*, pages 186–191. Morgan Kaufmann, 1999.
23. G. Kern-Isberner. *A unifying framework for symbolic and numerical approaches to nonmonotonic reasoning and belief revision.* Department of Computer Science, FernUniversität Hagen, 1999. Habilitation thesis.
24. G. Kern-Isberner. Solving the inverse representation problem. In *Proceedings 14th European Conference on Artificial Intelligence, ECAI'2000*, pages 581–585, Berlin, 2000. IOS Press.
25. G. Kern-Isberner. Conditional preservation and conditional indifference. *Journal of Applied Non-Classical Logics*, 11(1-2):85–106, 2001.
26. G. Kern-Isberner. Conditionals in knowledge representation and belief revision. In *Proceedings Fifth Dutch-German Workshop on Nonmonotonic Reasoning Techniques and their applications, DGNMR'01*, Potsdam, Germany, 2001.
27. G. Kern-Isberner. *Conditionals in nonmonotonic reasoning and belief revision.* Springer, Lecture Notes in Artificial Intelligence LNAI 2087, 2001.
28. G. Kern-Isberner. Handling conditionals adequately in uncertain reasoning. In *Proceedings European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU'01*, pages 604–615. Springer LNAI 2143, 2001.
29. G. Kern-Isberner. Representing and learning conditional information in possibility theory. In *Proceedings 7th Fuzzy Days, Dortmund, Germany*, pages 194–217. Springer LNCS 2206, 2001.
30. R. Kruse, E. Schwecke, and J. Heinsohn. *Uncertainty and Vagueness in Knowledge Based Systems.* Springer, Berlin Heidelberg New York, 1991.
31. I. Levi. Iteration of conditionals and the Ramsey test. *Synthese*, 76:49–81, 1988.
32. R.C. Lyndon and P.E. Schupp. *Combinatorial group theory.* Springer, Berlin Heidelberg New York, 1977.
33. D. Makinson and P. Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In *Proceedings Workshop The Logic of Theory Change, Konstanz, Germany, 1989*, pages 185–205, Berlin Heidelberg New York, 1991. Springer.
34. J.B. Paris. *The uncertain reasoner's companion – A mathematical perspective.* Cambridge University Press, 1994.
35. J.B. Paris and A. Vencovská. A method for updating that justifies minimum cross entropy. *International Journal of Approximate Reasoning*, 7:1–18, 1992.
36. F.P. Ramsey. General propositions and causality. In R.B. Braithwaite, editor, *Foundations of Mathematics and other logical essays*, pages 237–257. Routledge and Kegan Paul, New York, 1950.
37. J.E. Shore and R.W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, IT-26:26–37, 1980.
38. W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics, II*, pages 105–134. Kluwer Academic Publishers, 1988.