A thorough axiomatization of a principle of conditional preservation in belief revision

Gabriele Kern-Isberner
(gabriele.kern-isberner@fernuni-hagen.de)
FernUniversität Hagen, Department of Computer Science,
58084 Hagen, Germany

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 (Alchourrón et al., 1985) as a set of postulates outlining appropriate revision mechanisms in a propositional logical environment. Conditionals $(B|A)$, to be read as *"If A then B"*, seem to play an ambivalent role in belief revision: Although their dynamic power as *revision policies* has been appreciated (see e.g. (Ramsey, 1950; Boutilier and Goldszmidt, 1993)), Gärdenfors' *triviality result* (Gärdenfors, 1988) 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. (Levi, 1988)) 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 complex 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. (Ramsey, 1950; Katsuno and Mendelzon, 1991b)), it was only quite recently that first approaches extending the AGM-theory to that broader framework have been brought forth: Darwiche and Pearl (Darwiche and Pearl, 1997) 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 (Kern-Isberner, 1999a), 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 ($A \mathrel{|\!\sim}_{\Psi} B$), 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 (Makinson and Gärdenfors, 1991). 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 (Darwiche and Pearl, 1997) and (Kern-Isberner, 1999a) 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 includes any type of belief revision considered to date and generalizes the classical AGM-framework in three respects:

- observing conditional beliefs in the prior epistemic state
- handling revision by conditional beliefs
- handling simultaneous revision by a set of conditional beliefs

This principle to be developed here 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 (Kern-Isberner, 1999a) and hence also provides a high-level formalization of Darwiche and Pearl's ideas (Darwiche and Pearl, 1997).

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. Although the non-classical nature of conditionals has been widely recognized and emphasized, classical logical views have influenced (and limited) the handling of conditionals: Interactions of conditionals have been reduced to logical interactions, checking for logical inconsistencies (cf. e.g. (Goldszmidt and Pearl, 1996)), and one of the principal assumptions for *conditional events* is that their logic should extend classical logic (cf. e.g. (Dubois and Prade, 1991; Walker, 1994)). One of the main reasons for basing conditional and revision theories on classical logical approaches is to use the clear structure of classical logic as a guideline, in order not to get lost under revision. Indeed, the interactions of conditionals can become very complex, and an adequate structural means for handling sets of conditionals is urgently needed. Here we use the algebraic means of *conditional structures* to 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 by using the idea of *conditional indifference*. Conditional indifference proves to be more fundamental than the notion of conditional independence and is formalized in terms of conditionals. 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 and pointing out further applications and ongoing work.

This paper is an elaboration and extension of ideas presented in (Kern-Isberner, 2001c) and (Kern-Isberner, 2002a). All proofs can be found in the Appendix.

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 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 a 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. (Spohn, 1988)). 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, due to $\kappa^{-1}(0) \neq \emptyset$, 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 $\kappa(AB) + n < \kappa(A\overline{B})$ $(n \in \mathbb{N})$. Note that $\kappa \models (B|A)$ iff $\kappa \models (B|A)\,[0]$. OCF's are the qualitative counterpart of probability distributions. Their plausibility degrees may be taken as order-of-magnitude abstractions of probabilities (cf. (Goldszmidt et al., 1993; Goldszmidt and Pearl, 1996)).

– *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 (Spohn, 1988) to a possibilistic framework (see also (Kern-Isberner, 1999b) and (Benferhat et al., 1997)).

  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. (Dubois and Prade, 1994)). 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. (Benferhat et al., 1992)), but realize degrees of possibility (or plausibility) in a non-discrete, compact domain. They can be taken as fuzzy representations of epistemic states (cf. (Kruse et al., 1991;

Dubois et al., 1994)), and are closely related to belief functions (cf. (Benferhat et al., 2000)).

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 propositional 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 (Alchourrón et al., 1985) 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 states 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. (Darwiche and Pearl, 1997)):

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:

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

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.

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 $\mathcal{A} = (\mathcal{A}, \leqslant_{\mathcal{A}}, \oplus, \odot, 0^{\mathcal{A}}, 1^{\mathcal{A}})$ 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 $\leqslant_{\mathcal{A}}$ with minimum $0^{\mathcal{A}}$ and maximum $1^{\mathcal{A}}$, such that $\leqslant_{\mathcal{A}}$ 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 1. (conditional valuation function) A *conditional valuation function* is a (partial) function $V : \mathcal{L} \cup (\mathcal{L} \mid \mathcal{L}) \to \mathcal{A}$ from the sets of formulas and conditionals into the algebra $\mathcal{A}$ satisfying the following conditions:

1. $V|_{\mathcal{L}}$ is total such that $V(\bot) = 0^{\mathcal{A}}, V(\top) = 1^{\mathcal{A}}$, and for exclusive formulas A, B (i.e. $AB \equiv \bot$), it holds that $V(A \vee B) = V(A) \oplus V(B)$;

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 $\mathcal{A}$; for $V(A) = 0^{\mathcal{A}}$, $V(B|A)$ is undefined.

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. (Darwiche and Pearl, 1997)). There is, however, a close relationship between propositions and conditionals – propositions 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 $V : \Omega \to \mathcal{A}$. Note that, due to $1^{\mathcal{A}} = V(\top) = \sum_{\omega\in\Omega}{}^{\oplus} V(\omega)$, all $V(\omega)$ must "sum up" to $1^{\mathcal{A}}$. In general, 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, ((Friedman and Halpern, 1996; Freund, 1998)), that is, it fulfills $V(\perp) \leqslant_{\mathcal{A}} V(A)$ 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. Let V_u denote the uniform conditional valuation function.

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

$, +, \cdot, 0, 1)$ is a probability function. The uniform probability function is $P_u(\omega) = \dfrac{1}{|\Omega|}$.

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

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 truth and plausibility.

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 (DeFinetti, 1974):

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

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. (Calabrese, 1991)). 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* $(B|A)^+ := \mathit{Mod}(AB)$, and the *falsifying set* $(B|A)^- := \mathit{Mod}(A\overline{B})$ of a conditional $(B|A)$. $\mathit{Mod}(\overline{A})$ 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 $\mathit{Mod}(\overline{A}) = \emptyset$, 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 2. $(\overline{A}|A)$ 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 (Kern-Isberner, 2001b)).

Next, we introduce the notion of a subconditional:

Definition 2. (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 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 equivalence relation for conditionals can also be taken as to be induced by $\sqsubseteq$:

Lemma 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 3. (perpendicular conditionals, $\perp\!\!\!\perp$) Let $(B|A)$, $(D|C) \in (\mathcal{L} \mid \mathcal{L})$ be 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 3. 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 (Kern-Isberner, 2001c).

5. A principle of conditional preservation in a qualitative framework

In (Darwiche and Pearl, 1997), 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 (Alchourrón et al., 1985) 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 (Darwiche and Pearl, 1997)). 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 (Kern-Isberner, 1999a), 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)-(CR2) and (CR4)-(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, would never fulfill the precondition $\Psi * A \models \overline{A}$, as long as the revising proposition A is not inconsistent. 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 (Katsuno and Mendelzon, 1991a) 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 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 ((Darwiche and Pearl, 1997)) 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., (Kern-Isberner, 2001a; Kern-Isberner, 2000; Kern-Isberner, 2001d)). 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 (Kern-Isberner, 2001a); 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 $\mathcal{R}$; $\mathbf{a}_i^+$ symbolizes a positive effect for worlds verifying the respective conditional, whereas $\mathbf{a}_i^-$ symbolizes a negative effect for worlds falsifying it. Furthermore, in order 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. (Lyndon and Schupp, 1977; Fine and Rosenberger, 1999)). The commutativity of $\mathcal{F}_\mathcal{R}$ corresponds to a simultaneous application of the conditionals in $\mathcal{R}$, without assuming any order of application. Then the functions $\sigma_i = \sigma_{(B_i|A_i)}$, $1 \leqslant i \leqslant n$, defined by

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

represent the effects each conditional $(B_i|A_i)$ has on possible worlds ω. 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 group structure on $\mathcal{F}_\mathcal{R}$ allows us to form products and in this way, to make even complex interactions between the conditionals in $\mathcal{R}$ transparent.

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

Example 4. 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, $(c|b)$. For instance, the world abc verifies both conditionals, so we have $\sigma_\mathcal{R}(abc) = \mathbf{a}_1^+\mathbf{a}_2^+$. The following table shows the values of the function $\sigma_\mathcal{R}$ on arbitrary 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 $\sigma_\mathcal{R}(abc) \cdot \sigma_\mathcal{R}(\overline{a}\overline{b}c) \cdot \sigma_\mathcal{R}(a\overline{b}c)^{-1} \cdot \sigma_\mathcal{R}(\overline{a}bc)^{-1} = \mathbf{a}_1^+\mathbf{a}_2^+ \cdot 1 \cdot (\mathbf{a}_1^+)^{-1} \cdot (\mathbf{a}_2^+)^{-1} = 1$, which may be interpreted by saying that the sets of worlds $\{abc, \overline{a}\overline{b}c\}$ and $\{a\overline{b}c, \overline{a}bc\}$ show identical conditional effects – they are balanced with respect to the effects of the conditionals in $\mathcal{R}$. Although $abc, \overline{a}\overline{b}c, a\overline{b}c, \overline{a}bc$ all have different conditional structures, the relationships between them with respect to $\mathcal{R}$ are clearly revealed.

To comply with the group 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 words $\widehat{\omega} = {\omega_1}^{r_1} \ldots {\omega_m}^{r_m}$

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

$$\sigma_{\mathcal{R}}({\omega_1}^{r_1} \ldots {\omega_m}^{r_m}) = \sigma_{\mathcal{R}}(\omega_1)^{r_1} \ldots \sigma_{\mathcal{R}}(\omega_m)^{r_m}$$
$$\prod_{1 \leqslant i \leqslant n} (\mathbf{a}_i^+)^{\sum_{k:\sigma_i(\omega_k)=\mathbf{a}_i^+} r_k} \cdot \prod_{1 \leqslant i \leqslant n} (\mathbf{a}_i^-)^{\sum_{k:\sigma_i(\omega_k)=\mathbf{a}_i^-} r_k} \tag{5}$$

yielding a *homomorphism of groups* $\sigma_{\mathcal{R}} : \widehat{\Omega} \to \mathcal{F}_{\mathcal{R}}$. As for the elements of $\mathcal{F}_{\mathcal{R}}$, 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 $\widehat{\Omega}$:

$$\widehat{\omega}_1 \equiv_{\mathcal{R}} \widehat{\omega}_2 \text{ iff } \sigma_{\mathcal{R}}(\widehat{\omega}_1) = \sigma_{\mathcal{R}}(\widehat{\omega}_2) \text{ iff } \sigma_{\mathcal{R}}(\widehat{\omega}_1\widehat{\omega}_2^{-1}) = 1$$

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 $\sigma_{\mathcal{R}}$. 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 4 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 $\widehat{\Omega}_0 = ker\, \sigma_{(\top|\top)}$ of $\widehat{\Omega}$. Since $(\top|\top)(\omega) = 1$ for all $\omega \in \Omega$, we have

$$\sigma_{(\top|\top)}(\omega_1^{r_1} \ldots \omega_m^{r_m}) = (\mathbf{a}^+)^{r_1} \ldots (\mathbf{a}^+)^{r_m} = (\mathbf{a}^+)^{\sum_{j=1}^m r_j}$$

with some symbol $(\mathbf{a}^+)$ representing the positive effect of $(\top|\top)$ on possible worlds. Hence

$$\widehat{\Omega}_0 = \{\widehat{\omega} = {\omega_1}^{r_1} \cdot \ldots \cdot {\omega_m}^{r_m} \in \widehat{\Omega} \mid \sum_{j=1}^{m} r_j = 0\} \tag{6}$$

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* $\widehat{\Omega}_0$,

$$\widehat{\omega}_1 \equiv_{\top} \widehat{\omega}_2 \text{ iff } \widehat{\omega}_1\widehat{\Omega}_0 = \widehat{\omega}_2\widehat{\Omega}_0 \text{ iff } \sum_{1 \leqslant j \leqslant m} r_j = \sum_{1 \leqslant k \leqslant p} s_k$$

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 $\widehat{\Omega}_0$.

Example 5. In Example 4, we have $abc \cdot \overline{a}\overline{b}c \equiv_\top a\overline{b}c \cdot \overline{a}bc$, so $\dfrac{abc \cdot \overline{a}\overline{b}c}{a\overline{b}c \cdot \overline{a}bc}$ is not only an element of $ker\,\sigma_{\mathcal{R}}$, but also of $ker_0\ \sigma_{\mathcal{R}}$. Note that, although also $\sigma_{\mathcal{R}}(\dfrac{abc}{a\overline{b}c \cdot \overline{a}bc}) = 1$, $\dfrac{abc}{a\overline{b}c \cdot \overline{a}bc} \notin ker_0\ \sigma_{\mathcal{R}}$ because $\sigma_{(\top|\top)}(abc) = (\mathbf{a}^+) \neq (\mathbf{a}^+)^2 = \sigma_{(\top|\top)}(a\overline{b}c \cdot \overline{a}bc)$

Finally, we will show how to describe the relations $\sqsubseteq$ and $\perp\!\!\!\perp$ between conditionals, introduced in Definitions 2 and 3, respectively, by considering the kernels of the corresponding σ-homomorphisms. As a convenient notation, for each proposition $A \in \mathcal{L}$, we define

$$\widehat{A} := \{\widehat{\omega} = \omega_1^{r_1} \ldots \omega_m^{r_m} \in \widehat{\Omega} \mid \omega_i \models A \text{ for all } i, 1 \leqslant i \leqslant m\}$$

Proposition 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 $ker\,\sigma_{(D|C)} \cap \widehat{C} = ker\,\sigma_{(B|A)} \cap \widehat{C}$.

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 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 4. (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 $\sigma_{(B|A)}(\omega) \neq 1$ and $V(\omega') = 0^{\mathcal{A}}$ 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 4 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 $\mathcal{R}$.

A first simple, but important property of $\mathcal{R}$-indifferent valuation functions V is that $\equiv_{\mathcal{R}}$-equivalent worlds are mapped onto the same values under V:

Lemma 3. If the conditional valuation function V is indifferent with respect to $\mathcal{R}$, then $\sigma_{\mathcal{R}}(\omega_1) = \sigma_{\mathcal{R}}(\omega_2)$ implies $V(\omega_1) = V(\omega_2)$ for all worlds $\omega_1, \omega_2 \in \Omega$.

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

Proposition 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 4 holds, and $\mathit{ker}_0\, \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+ \subseteq \mathit{ker}_0\, V$.

The next theorem provides a clear characterization of probability functions, ordinal conditional functions and possibility distributions with indifference properties:

Theorem 2. Let $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\} \subseteq (\mathcal{L} \mid \mathcal{L})$ 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 α_0, $\alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+, \alpha_0 > 0$ such that, for all $\omega \in \Omega$,

$$P(\omega) = \alpha_0 \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{7}$$

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 $\omega \in \Omega$,

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

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 $\omega \in \Omega$,

$$\pi(\omega) = \alpha_0 \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{9}$$

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 2, 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 (7), (8) or (9), 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.

Definition 5. 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. (Kern-Isberner, 1998; Kern-Isberner, 2001a; Kern-Isberner, 2001e); see also Section 7.

Theorem 2 also shows, that most important and well-behaved inductive representation methods realize conditional indifference: Namely, the principle of maximum entropy in probabilistics (Paris, 1994), system-Z^* in the OCF-framework (Goldszmidt et al., 1993), and the LCD-functions of Benferhat, Saffiotti and Smets (Benferhat et al., 2000) all give rise to conditionally indifferent functions (cf. (Kern-Isberner, 1998; Kern-Isberner, 2001a; Kern-Isberner, 2001d)). 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. (Kern-Isberner, 2001d)). 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.

We close this section by establishing an interesting connection between conditional indifference and conditional independence, one of the most important means to support probabilistic reasoning in general and the crucial glue to build up Bayesian networks in particular (cf. e.g. (Pearl, 1988; Cowell et al., 1999)). The next proposition shows that conditional indifference is a more fundamental notion than conditional independence, reflecting more fine-grained structures.

Proposition 3. Let $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ be disjoint subsets of a set of propositional variables $\mathcal{V}$, and let P be a probability distribution over $\mathcal{V}$. Let $\mathcal{R}$ be the following set of conditionals:

$$\mathcal{R} = \{(\mathbf{x}|\mathbf{z}), (\mathbf{y}|\mathbf{z}) \mid \mathbf{x}, \mathbf{y}, \mathbf{z} \text{ instantiations of variables in } \mathbf{X}, \mathbf{Y}, \mathbf{Z}, \text{ resp.}\}$$

If P is indifferent with respect to $\mathcal{R}$, then $\mathbf{X}$ and $\mathbf{Y}$ are conditionally independent in P, given $\mathbf{Z}$.

Note that the converse of Proposition 3 does not hold. It is easy to build up a probability distribution over, e.g., four variables x, y, z, w such that x and y are conditionally independent given z, but $P(xyzw) \neq P(xyz\overline{w})$. However, for the conditional structure with respect to the respective $\mathcal{R}$ in this case, w does not matter, so $\sigma_{\mathcal{R}}(xyzw) = \sigma_{\mathcal{R}}(xyz\overline{w})$. Hence P can not be indifferent with respect to $\mathcal{R}$.

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 $\odot$ on 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 6. (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 $\mathcal{R}^{(*)}$; 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 $\sigma_{(B|A)}(\omega) \neq 1$ and $V^*(\omega') = 0^{\mathcal{A}}$ 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 $\widehat{\omega}_1 \equiv_{\top} \widehat{\omega}_2$ 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 6 is an appropriate generalization of Definition 4 for revisions. Indeed, it can easily be verified that conditional valuation functions are indifferent with respect to $\mathcal{R}$ iff they are indifferent with respect to R and the uniform conditional valuation function V_u.

Note that also for this extended notion of indifference, the quantifications of conditionals do not matter. For the revision process, however, quantifications if present have to be taken into account. So we use both symbols, $\mathcal{R}$ and $\mathcal{R}^*$, when considering indifferent revisions, and $\mathcal{R}$ will always denote the set of unquantified conditionals occurring in $\mathcal{R}^*$. Whereas in the probabilistic framework, quantifications of conditionals are essential, they may be omitted in the ordinal or possibilistic framework.

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

Definition 7. (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 2, 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 3. Let $\mathcal{R}^{(*)} = \{(B_1|A_1)([x_1]), \ldots, (B_n|A_n)([x_n])\} \subseteq (\mathcal{L} \mid \mathcal{L})$ 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 $P^*(A_i) \neq 0$, and 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 $\omega \in \Omega$,
$$P^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{10}$$

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 $\omega \in \Omega$,
$$\kappa^*(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \kappa_i^+ + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \kappa_i^- \tag{11}$$

3. A revision $\pi^* = \pi * \mathcal{R}^{(*)}$ satisfies the principle of conditional preservation with respect to $\mathcal{R}$ and π iff $\pi^*(A_i) \neq 0$, and there are non-negative real numbers $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+$ with $\alpha_0 > 0$ such that for all $\omega \in \Omega$,
$$\pi^*(\omega) = \alpha_0 \pi(\omega) \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^- \tag{12}$$

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 8. (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 3 and choosing the constants $\alpha_0, \alpha_i^+, \alpha_i^-$, and $\kappa_0, \kappa_i^+, \kappa_i^-$, respectively, appropriately so as to establish the necessary numerical relationships. Comparing Theorem 3 with Theorem 2 also shows clearly that c-representations of a set of conditionals $\mathcal{R}^{(*)}$ are c-revisions of uniform conditional valuation functions by $\mathcal{R}^{(*)}$. To illustrate this, we will go into this in more detail 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 (11), and the postulate $\kappa^* \models \mathcal{R}$ yields the following conditions for κ_i^+, κ_i^- in a straightforward way:

$$\begin{aligned}\kappa_i^- - \kappa_i^+ \; > \; & \min_{\omega \models A_i B_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \\ & - \min_{\omega \models A_i \overline{B}_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-)\end{aligned} \tag{13}$$

Moreover, quantifications of conditionals can be taken easily into account by modifying (13) slightly, so as to comply with the representation postulate $\kappa^* \models (B|A)\,[m_i]$:

$$\begin{aligned}\kappa_i^- - \kappa_i^+ \; > \; & m_i + \min_{\omega \models A_i B_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) \\ & - \min_{\omega \models A_i \overline{B}_i} (\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j B_j}} \kappa_j^+ + \sum_{\substack{j \neq i \\ \omega \models A_j \overline{B}_j}} \kappa_j^-)\end{aligned} \tag{14}$$

C-revisions exist for any finitely valued OCF κ and any consistent set $\mathcal{R}$ of conditionals; if κ also takes on infinite values, some basic demands for compatibility between κ and $\mathcal{R}$ have to be observed (cf. (Kern-Isberner, 2001c)). In the following, we will describe a procedure how to calculate such a c-revision for any finite OCF κ and any finite consistent set $\mathcal{R}$ of conditionals.

The consistency of a set $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ of conditionals in a qualitative framework can be characterized by the notion of *tolerance*. A conditional $(B|A)$ is said to be *tolerated* by a set of conditionals $\mathcal{R}$ iff there is a world ω such that ω verifies $(B|A)$ (i.e. $(B|A)(\omega) = 1$) and ω does not falsify any of the conditionals in $\mathcal{R}$ (i.e. $r(\omega) \neq 0$ for all $r \in \mathcal{R}$). $\mathcal{R}$ is consistent iff there is an ordered partition $\mathcal{R}_0, \mathcal{R}_1, \ldots, \mathcal{R}_k$ of $\mathcal{R}$ such that each conditional in $\mathcal{R}_m$ is tolerated by $\bigcup_{j=m}^{k} \mathcal{R}_j$, $0 \leqslant m \leqslant k$ (cf. (Goldszmidt and Pearl, 1996)).

Now suppose that $\mathcal{R}$ is consistent and such a partition $\mathcal{R}_0, \mathcal{R}_1, \ldots, \mathcal{R}_k$ of $\mathcal{R}$ is given. For all conditionals $r_i \in \mathcal{R}$, $1 \leqslant i \leqslant n$, set $\kappa_i^+ := 0$, and set successively, for each partitioning set $\mathcal{R}_m$, $0 \leqslant m \leqslant k$, starting with $\mathcal{R}_0$, and for each conditional $r_i = (B_i|A_i) \in \mathcal{R}_m$

$$\kappa_i^- := \min_{\substack{\omega \models A_i B_i \\ r(\omega) \neq 0 \forall r \in \cup_{l=m}^{k} \mathcal{R}_l}} (\kappa(\omega) + \sum_{\substack{r_j \in \cup_{l=0}^{m-1} \mathcal{R}_l \\ \omega \models A_j \overline{B}_j}} \kappa_j^-) + 1 \tag{15}$$

Finally, choose κ_0 appropriately to make $\kappa^*(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B}_i}} \kappa_i^-$ an ordinal conditional function. It is straightforward to check that indeed, $\kappa^* \models \mathcal{R}$, so κ^* is a c-revision of κ by $\mathcal{R}$. In the same way, by applying these ideas to the uniform OCF $\kappa_u(\omega) = 0$ (for all $\omega \in \Omega$), we obtain c-representations of $\mathcal{R}$.

We will illustrate the basic ideas and features of c-representations and c-revisions by an example.

Example 6. Epistemic knowledge about important relationships between the atoms f - *flying*, b - *birds*, p - *penguins*, w - winged animals, and k - *kiwis* is to be represented by an OCF. Let the set $\mathcal{R}$ consist of the following conditionals:

$\mathcal{R}$: 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: $(b|k)$ *kiwis are birds*

We will apply the procedure sketched above to compute an ordinal conditional function κ which is a c-representation of $\mathcal{R}$.

The conditionals r_1, r_4, and r_5 are tolerated by $\mathcal{R}$, whereas r_2 and r_3 are not; but both r_2 and r_3 are tolerated by the set $\{r_2, r_3\}$. This yields the partitioning $\mathcal{R}_0 = \{r_1, r_4, r_5\}, \mathcal{R}_1 = \{r_2, r_3\}$ of $\mathcal{R}$. In order to obtain a suitable c-representation of $\mathcal{R}$, we set $\kappa_i^+ = 0$ for all $i, 1 \leqslant i \leqslant 5$, and, according to (15),

$$\begin{aligned} \kappa_1^- = \kappa_4^- = \kappa_5^- &= 1, \\ \kappa_2^- = \kappa_3^- = \kappa_1^- + 1 &= 2 \end{aligned}$$

The resulting c-representation $\kappa(\omega) := \sum\limits_{\substack{1 \leqslant i \leqslant 5 \\ r_i(\omega) = 0}} \kappa_i^-$ of $\mathcal{R}$ is shown in the figure on page 28.

It is now easily checked that $\kappa \models (w|k)$ – 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 (11), 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 (13) 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 $\kappa^+ := 0, \kappa^- := 2$. No further normalization is necessary, so $\kappa_0 := 0$. The revised κ^* is shown in the figure on page 28, 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 $\kappa^*(bk) = \kappa^*(\overline{b}k) = 1$ – since *birds* and *wings* have been plausibly related by the conditional $(w|b)$, the property of not having wings casts (reasonably) doubt on kiwis being birds. This illustrates how conditional interrelationships are properly dealt with by c-revisions. One might wish, however, to state that *kiwi* and *birds* are more firmly related than *birds* and *wings*, in order to be able to accept $(b|k)$ still after revising κ by $(\overline{w}|k)$. This can be achieved by assigning an inferential strength $x \geqslant 1$ to $(b|k)$ (and – for reasons of symmetry – also to $(b|p)$, because *kiwis* and *penguins* both are *birds* by definition). Since no explicit quantification means assuming an inferential strength of 0, this amounts to considering the following set $\mathcal{R}'$ of quantified conditionals:

$$\mathcal{R}' = \{(f|b)[0], (b|p)[x], (\overline{f}|p)[0], (w|b)[0], (b|k)[x]\}$$

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

ω	$\kappa(\omega)$	$\kappa^*(\omega)$	$\kappa_1(\omega)$	$\kappa_1^*(\omega)$
$pbfwk$	2	4	2	4
$pbfw\overline{k}$	2	2	2	2
$pbf\overline{w}k$	3	3	3	3
$pbf\overline{w}\overline{k}$	3	3	3	3
$pb\overline{f}wk$	1	3	1	3
$pb\overline{f}w\overline{k}$	1	1	1	1
$pb\overline{f}\overline{w}k$	2	2	2	2
$pb\overline{f}\overline{w}\overline{k}$	2	2	2	2
$p\overline{b}fwk$	5	7	$2x+5$	$2x+7$
$p\overline{b}fw\overline{k}$	4	4	$x+4$	$x+4$
$p\overline{b}f\overline{w}k$	5	5	$2x+5$	$2x+5$
$p\overline{b}f\overline{w}\overline{k}$	4	4	$x+4$	$x+4$
$p\overline{b}\,\overline{f}wk$	3	5	$2x+3$	$2x+5$
$p\overline{b}\,\overline{f}w\overline{k}$	2	2	$x+2$	$x+2$
$p\overline{b}\,\overline{f}\overline{w}k$	3	3	$2x+3$	$2x+3$
$p\overline{b}\,\overline{f}\overline{w}\overline{k}$	2	2	$x+2$	$x+2$
$\overline{p}bfwk$	0	2	0	2
$\overline{p}bfw\overline{k}$	0	0	0	0
$\overline{p}bf\overline{w}k$	1	1	1	1
$\overline{p}bf\overline{w}\overline{k}$	1	1	1	1
$\overline{p}b\overline{f}wk$	1	3	1	3
$\overline{p}b\overline{f}w\overline{k}$	1	1	1	1
$\overline{p}b\overline{f}\overline{w}k$	2	2	2	2
$\overline{p}b\overline{f}\overline{w}\overline{k}$	2	2	2	2
$\overline{p}\overline{b}fwk$	1	3	$x+1$	$x+3$
$\overline{p}\overline{b}fw\overline{k}$	0	0	0	0
$\overline{p}\overline{b}f\overline{w}k$	1	1	$x+1$	$x+1$
$\overline{p}\overline{b}f\overline{w}\overline{k}$	0	0	0	0
$\overline{p}\overline{b}\,\overline{f}wk$	1	3	$x+1$	$x+3$
$\overline{p}\overline{b}\,\overline{f}w\overline{k}$	0	0	0	0
$\overline{p}\overline{b}\,\overline{f}\overline{w}k$	1	1	$x+1$	$x+1$
$\overline{p}\overline{b}\,\overline{f}\overline{w}\overline{k}$	0	0	0	0

Figure 1. OCF's κ and κ_1, and revised κ^* and κ_1^* for Example 6

A c-representation, κ_1, of $\mathcal{R}'$ and a c-revision $\kappa_1^* = \kappa_1 * \{(\overline{w}|k)\}$ (computed in the same way as above) are both shown in the figure on page 28, too. Now we find $\kappa_1^*(\overline{b}k) = x + 1 > 1 = \kappa_1^*(bk)$, so the conditional $(b|k)$ is still accepted in κ_1^*.

The idea of c-revisions can be recovered in well-known approaches to non-propositional revision. Actually, the so-called *J-conditioning* presented in (Goldszmidt and Pearl, 1996) to adjust an OCF κ to uncertain evidence is actually a c-revision. In a probabilistic framework, due to Theorem 3 it is easily seen that revisions following the *principle of minimum cross-entropy* (so-called *MINENT-principle* or, briefly, *ME-principle*) (Shore and Johnson, 1980; Paris and Vencovská, 1992; Paris, 1994; Kern-Isberner, 1998) 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-theoretical concept of *cross-entropy* $R(Q,P) = \sum\limits_{\omega \in \Omega} Q(\omega) \log \frac{Q(\omega)}{P(\omega)}$. 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. (Csiszár, 1975)), which can be written in the form

$$P^*(\omega) = \alpha_0 P(\omega) \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i B_i}} \alpha_i^{1-x_i} \prod_{\substack{1 \leqslant i \leqslant n \\ \omega \models A_i \overline{B_i}}} \alpha_i^{-x_i} \tag{16}$$

with the α_i's being exponentials of the Lagrange multipliers, appropriately chosen so as to satisfy all conditionals in $\mathcal{R}^*$ (cf. (Kern-Isberner, 2001c)). Comparing (16) to (10), it is obvious that $P *_{ME} \mathcal{R}^*$ satisfies the principle of conditional preservation, and hence is a c-revision.

An ME-revision realizes perfectly the idea of unique, minimal change in a probabilistic environment. For ordinal frameworks, the ideas underlying system-Z^* (Goldszmidt et al., 1993), or the LCD-functions (Benferhat et al., 2000), can now also be applied to make revisions "reasonably minimal", due to the structural similarity of c-representations and c-revisions (cf. Section 6). Basically, that is to say, that verification of conditionals should not change a world's degree of plausibility, hence setting $\kappa_i^+ = 0$ in (11), and $\alpha_i^+ = 1$ in (12), respectively, and worlds falsifying conditionals should be shifted minimally, which amounts to choosing κ_i^- in (11), and α_i^- in (12) as small as possible. Our Example 6 follows this idea, too.

By Definition 7, 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 (Kern-Isberner, 1998). Therefore, the axiomatization provided by Definition 7 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. (Goldszmidt et al., 1993)).

Now that we are able to carry out belief revision in a most general sense, namely by revising epistemic states by sets of (quantified) conditionals, an approach to give semantics to nested conditionals in epistemic states Ψ can be made via a straightforward generalization of the Ramsey test (cf. (1)):

$$\Psi \models ((D|C) \mid (B|A)) \quad \text{iff} \quad \Psi * \{(B|A)\} \models (D|C) \tag{17}$$

This is different from Goldszmidt & Pearl's suggestion made in (Goldszmidt and Pearl, 1996, p. 88) where nested conditionals are evaluated with respect to a knowledge base, not to an epistemic state. In our framework, their approach amounts to the following:

Given a set of conditionals (defaults) $\mathcal{R}$, the nested conditional $((D|C) \mid (B|A))$ is accepted iff $\Psi_0 * (\mathcal{R} \cup \{(B|A)\}) \models (D|C)$.

Here, Ψ_0 is the uniform epistemic state; note that in our framework, representations of conditional knowledge bases are obtained by revisions of uniform epistemic states. This approach simply adds the antecedent of $((D|C) \mid (B|A))$ to the current default base $\mathcal{R}$ and checks the consequences of this new default base. Here, Goldszmidt & Pearl emphasize the "essential distinction" between having a conditional explicitly represented in $\mathcal{R}$, or merely satisfied as a (nonmonotonic) consequence of $\mathcal{R}$. Indeed, if $(\overline{B}|A)$ is merely a default consequence of $\mathcal{R}$, then nevertheless $(B|A)$ might be consistent with $\mathcal{R}$ and $\mathcal{R} \cup \{(B|A)\}$ will yield reasonable inferences. Whereas, if $(\overline{B}|A) \in \mathcal{R}$, then $\mathcal{R} \cup \{(B|A)\}$ is definitely inconsistent and has no nonmonotonic consequences at all. This problem does not occur with the first definition (17) – even if $\Psi \models (\overline{B}|A)$, a revision $\Psi * \{(B|A)\}$ is always possible. Therefore, a distinction between explicit and implicit knowledge – a point the importance of which is pointed out by Goldszmidt & Pearl – seems to be impossible in our approach which uses basically epistemic states for inferences.

The difference between these two approaches to nesting conditionals is better understood from a more general point of view. In a framework as rich as ours, the epistemic state Ψ may be thought of as being formed by a combination of prior (or background) knowledge Ψ_1, and posterior (or evidential) default knowledge $\mathcal{R}$ via revision: $\Psi = \Psi_1 * \mathcal{R}$. Now there are two possible ways of incorporating new conditional knowledge, namely by *successive* revision, $(\Psi_1 * \mathcal{R}) * \{(B|A)\}$, which roughly corresponds to *updating* (Katsuno and Mendelzon, 1991a), or by *simultaneous* revision, $\Psi * (\mathcal{R} \cup \{(B|A)\})$, which is more in the sense of *AGM-revision* (Gärdenfors, 1988). Now it becomes clear that the difference between the two approaches above results from the difference between these two kinds of revision – in general, $\Psi * \mathcal{R} * \mathcal{S}$ and $\Psi * (\mathcal{R} \cup \mathcal{S})$, $\mathcal{R}, \mathcal{S} \subseteq (\mathcal{L} \mid \mathcal{L})$, will be found to differ, as the following example shows.

Example 7. We go back to Example 6. Here, the set $\mathcal{R}$ of conditionals can be split into two sets $\mathcal{S}_1, \mathcal{S}_2$ with $\mathcal{S}_1 = \{r_1, r_2, r_4, r_5\}$ and $\mathcal{S}_2 = \{r_3\}$: $\mathcal{R} = \mathcal{S}_1 \cup \mathcal{S}_2$. Suppose that first $\mathcal{S}_1$ is to be learnt and c-represented. Since all $r_i \in \mathcal{S}_1$ are tolerated by $\mathcal{S}_1$, we may choose $\lambda_i^+ = 0, \lambda_i^- = 1, i \in \{1, 2, 4, 5\}$, as appropriate revision constants in (11) (with $\kappa = \kappa_u$ being the uniform ordinal conditional function), thus arriving at $\kappa_2 := \kappa_u * \mathcal{S}_1$ as a c-representation of $\mathcal{S}_1$ (see the figure on page 32). A c-revision of κ_2 by $\mathcal{S}_2 = \{r_3\}$ according to the strategy described above can then be obtained by adding $\lambda_3^- = 2$ to all worlds falsifying r_3. The resulting $\kappa_2^* = \kappa_2 * \mathcal{S}_2 = (\kappa_u * \mathcal{S}_1) * \mathcal{S}_2$ is shown in the figure on page 32, too, and is clearly seen to be different from $\kappa = \kappa_u * \mathcal{R} = \kappa_u * (\mathcal{S}_1 \cup \mathcal{S}_2)$ in the figure on page 28.

Which type of revision – simultaneous or successive revision – is more appropriate will depend on the relation between already present knowledge, $\mathcal{R}$, and new incoming information, $\mathcal{S}$. If both pertain to the same situation, or the same world, respectively, simultaneous revision should be used; otherwise, successive revision seems to be the proper way to change beliefs. Actually, it needs this general framework for belief revision to understand this thoroughly, since successive and simultaneous revision cannot be distinguished in a purely propositional 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

ω	$\kappa_2(\omega)$	$\kappa_2^*(\omega)$	ω	$\kappa_2(\omega)$	$\kappa_2^*(\omega)$
$pbfwk$	0	2	$pbfw\overline{k}$	0	2
$pbf\overline{w}k$	1	3	$pbf\overline{w}\overline{k}$	1	3
$pb\overline{f}wk$	1	1	$pb\overline{f}w\overline{k}$	1	1
$pb\overline{f}\overline{w}k$	2	2	$pb\overline{f}\overline{w}\overline{k}$	2	2
$p\overline{b}fwk$	2	4	$p\overline{b}fw\overline{k}$	1	3
$p\overline{b}f\overline{w}k$	2	4	$p\overline{b}f\overline{w}\overline{k}$	1	3
$p\overline{b}\,\overline{f}wk$	2	2	$p\overline{b}\,\overline{f}w\overline{k}$	1	1
$p\overline{b}\,\overline{f}\overline{w}k$	2	2	$p\overline{b}\,\overline{f}\overline{w}\overline{k}$	1	1
$\overline{p}bfwk$	0	0	$\overline{p}bfw\overline{k}$	0	0
$\overline{p}bf\overline{w}k$	1	1	$\overline{p}bf\overline{w}\overline{k}$	1	1
$\overline{p}b\overline{f}wk$	1	1	$\overline{p}b\overline{f}w\overline{k}$	1	1
$\overline{p}b\overline{f}\overline{w}k$	2	2	$\overline{p}b\overline{f}\overline{w}\overline{k}$	2	2
$\overline{p}\overline{b}fwk$	1	1	$\overline{p}\overline{b}fw\overline{k}$	0	0
$\overline{p}\overline{b}f\overline{w}k$	1	1	$\overline{p}\overline{b}f\overline{w}\overline{k}$	0	0
$\overline{p}\overline{b}\,\overline{f}wk$	1	1	$\overline{p}\overline{b}\,\overline{f}w\overline{k}$	0	0
$\overline{p}\overline{b}\,\overline{f}\overline{w}k$	1	1	$\overline{p}\overline{b}\,\overline{f}\overline{w}\overline{k}$	0	0

Figure 2. OCF κ_2 and revised κ_2^* for Example 7

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

$$V^*(\omega) = \begin{cases} \alpha^+ \odot V(\omega) & \text{if} \quad \omega \models AB \\ \alpha^- \odot V(\omega) & \text{if} \quad \omega \models A\overline{B} \\ \alpha_0 \odot V(\omega) & \text{if} \quad \omega \models \overline{A} \end{cases} \tag{18}$$

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 4. 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 5. 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), is also in accordance with the qualitative principle of conditional preservation (as described by (CR5)-(CR7)):

Theorem 4. 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 4 identifies the principle of conditional preservation, as formalized in Definition 7, as a fundamental device to guide reasonable changes in the conditional structure of knowledge.

9. Conclusion and Outlook

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. In particular, the problem of nesting conditionals can be addressed and dealt with properly in our 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.

This rich, formal framework we used to develop our axiomatization of principles of conditional preservation, with the basic notions of conditional structures and conditional indifference, can also be used to study basically structural approaches to default reasoning (cf. (Kern-Isberner, 2002b)). The connections to group theory which might appear a bit strange at first sight can most efficiently be used to discover relevant conditional relationships in statistical data (Kern-Isberner, 2000). The implementations of these ideas as a computer system, CONDOR[2], are part of our ongoing work; a description of CONDOR as an abstract state machine can be found in (Beierle and Kern-Isberner, 2003).

Appendix

Proofs

Proof of Theorem 1. Let $A, B, C, D \in \mathcal{L}$.

Suppose $C \leqslant B$ or $C \leqslant \overline{B}$. Then $(D|C) \perp\!\!\!\perp (B|\top)$. (CR3) and (CR5) now imply (C1) and (C2).

(C3) and (C4) are direct consequences of (CR6) and (CR7) by using that $(B|A) \sqsubseteq (B|\top)$ and $(\overline{B}|A) \sqsubseteq (\overline{B}|\top)$, respectively, due to Lemma 1.

Proof of Proposition 1.

Proof of (1): Let $(D|C) \sqsubseteq (\dot{B}|A)$, where $\dot{B}$ is one of $B, \overline{B}$. Then $C \leqslant A$, by Lemma 1. Let $\widehat{\omega} = \omega_1^{r_1} \ldots \omega_m^{r_m} \in \mathit{ker}\ \sigma_{(D|C)} \cap \widehat{C}$, thus $\omega_k \models C$ for all $1 \leqslant k \leqslant m$, and hence $(D|C)(\omega_k) = (\dot{B}|A)(\omega_k) \in \{0,1\}$, using notation (2). So, $(D|C)(\omega_k) = 1$ iff $(\dot{B}|A)(\omega_k) = 1$, and $(D|C)(\omega_k) = 0$ iff $(\dot{B}|A)(\omega_k) = 0$. $1 = \sigma_{(D|C)}(\widehat{\omega})$ iff $\sum_{k:(D|C)(\omega_k)=1} r_k = \sum_{k:(D|C)(\omega_k)=0} r_k = 0$, due to (5). But this is equivalent to $\sum_{k:(B|A)(\omega_k)=1} r_k = \sum_{k:(B|A)(\omega_k)=0} r_k = 0$, too, and therefore to $\sigma_{(B|A)}(\widehat{\omega}) = 1$.

Conversely, suppose $\mathbf{C} \leqslant A$ and $\mathit{ker}\ \sigma_{(D|C)} \cap \widehat{C} = \mathit{ker}\ \sigma_{(B|A)} \cap \widehat{C}$. The case $(D|C)^+ = (D|C)^- = \emptyset$ is trivial, and also the case $|(D|C)^+| + |(D|C)^-| = 1$ is easily dealt with: For instance, let $(D|C)^+ = \{\omega_0\}$ and $(D|C)^- = \emptyset$. Then $\omega_0 \models A$, so that one of $\omega_0 \models AB$ or $\omega_0 \models A\overline{B}$ holds. Then clearly $(D|C) \sqsubseteq (B|A)$ or $(D|C) \sqsubseteq (\overline{B}|A)$.

[2] The development of CONDOR is supported by the DFG – Deutsche Forschungsgemeinschaft within the CONDOR-project under grant BE 1700/5-1.

So, let us now assume $|(D|C)^+| \geqslant 1, |(D|C)^-| \geqslant 1$, and let $\omega_1 \in (D|C)^+, \omega_2 \in (D|C)^-$. Then $\sigma_{(D|C)}(\frac{\omega_1}{\omega_2}) \neq 1$, so $\frac{\omega_1}{\omega_2} \notin \mathit{ker}\, \sigma_{(D|C)} \cap \widehat{C}$ and hence also $\frac{\omega_1}{\omega_2} \notin \mathit{ker}\, \sigma_{(B|A)} \cap \widehat{C}$. Therefore, $(B|A)(\omega_1) \neq (B|A)(\omega_2)$, so we have $\omega_1 \in (B|A)^+$ and $\omega_2 \in (B|A)^-$ or the other way round. On the other hand, for any ω_i' such that $(D|C)(\omega_i') = (D|C)(\omega_i)$ $(i \in \{1,2\})$, the presupposition $\mathit{ker}\, \sigma_{(D|C)} \cap \widehat{C} = \mathit{ker}\, \sigma_{(B|A)} \cap \widehat{C}$ implies that $(B|A)(\omega_i') = (B|A)(\omega_i)$. So $(D|C) \sqsubseteq (\dot{B}|A)$, as desired.

Proof of (2): Let $(D|C) \perp\!\!\!\perp (B|A)$, i.e. $C \leqslant AB, A\overline{B}$ or $\overline{A}$, respectively. Thus $\sigma_{(B|A)}(\omega)$ is the same for all $\omega \models C$. Due to cancellations, $\widehat{C} \cap \widehat{\Omega}_0 \subseteq \mathit{ker}\, \sigma_{(B|A)}$.

Conversely, suppose $(D|C) \perp\!\!\!\perp (B|A)$ does not hold. Then there are $\omega_1, \omega_2 \models C$ such that $\sigma_{(B|A)}(\omega_1) \neq \sigma_{(B|A)}(\omega_2)$, i.e. $\sigma_{(B|A)}\left(\frac{\omega_1}{\omega_2}\right) \neq 1$. So $\frac{\omega_1}{\omega_2} \in \widehat{C} \cap \widehat{\Omega}_0$, but $\frac{\omega_1}{\omega_2} \notin \mathit{ker}\, \sigma_{(B|A)}$.

Proof of Lemma 3. Let $\omega_1, \omega_2 \in \Omega$ such that $\sigma_{\mathcal{R}}(\omega_1) = \sigma_{\mathcal{R}}(\omega_2)$. If $V(\omega_1) = 0^{\mathcal{A}}$ then there is $(B|A) \in \mathcal{R}$ with $\sigma_{(B|A)}(\omega_1) \neq 1$ and $V(\omega') = 0^{\mathcal{A}}$ for all ω' with $\sigma_{(B|A)}(\omega') = \sigma_{(B|A)}(\omega)$. $\sigma_{\mathcal{R}}(\omega_1) = \sigma_{\mathcal{R}}(\omega_2)$ implies in particular $\sigma_{(B|A)}(\omega_1) = \sigma_{(B|A)}(\omega_2)$, and hence $V(\omega_2) = 0^{\mathcal{A}}$, too, by condition (i) of Definition 4(1).

Now suppose $V(\omega_1), V(\omega_2) \neq 0^{\mathcal{A}}$, i.e. $\omega_1, \omega_2 \in \widehat{\Omega}_+$. Moreover, we have $\frac{\omega_1}{\omega_2} \in \widehat{\Omega}_0$, so due to the presupposition $\sigma_{\mathcal{R}}(\omega_1) = \sigma_{\mathcal{R}}(\omega_2)$, we obtain $V(\omega_1) = V(\omega_2)$, by condition (ii) of Definition 4, (1) and (2).

Proof of Proposition 2. Suppose $V : \mathcal{L} \to \mathcal{A}$ is a conditional valuation function which is indifferent with respect to $\mathcal{R}$. Then, by definition, condition (i) of Definition 4 holds. Let $\widehat{\omega} \in \mathit{ker}_0\ \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+$, i.e. $\widehat{\omega} \in \widehat{\Omega}_0$, and $\sigma_{\mathcal{R}}(\widehat{\omega}) = 1 = \sigma_{\mathcal{R}}(\epsilon_\Omega)$, where ϵ_Ω is the empty word in $\widehat{\Omega}$. Because V is indifferent with respect to $\mathcal{R}$, we obtain $V(\widehat{\omega}) = V(\epsilon_\Omega) = 1$, so $\widehat{\omega} \in \mathit{ker}_0\ V$.

Conversely, let $V : \mathcal{L} \to \mathcal{A}$ be a conditional valuation function such that condition (i) of Definition 4 holds and $\mathit{ker}_0\ \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+ \subseteq \mathit{ker}_0\ V$. Suppose $\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \sigma_{\mathcal{R}}(\widehat{\omega}_2)$ for $\widehat{\omega}_1, \widehat{\omega}_2 \in \widehat{\Omega}_+$, $\widehat{\omega}_1 \equiv_\top \widehat{\omega}_2$. Then $\sigma_{\mathcal{R}}(\widehat{\omega}_1 \cdot \widehat{\omega}_2^{-1}) = 1$, i.e. $\widehat{\omega}_1 \cdot \widehat{\omega}_2^{-1} \in \mathit{ker}_0\ \sigma_{\mathcal{R}} \cap \widehat{\Omega}_+ \subseteq \mathit{ker}_0\ V$. This implies $V(\widehat{\omega}_1 \cdot \widehat{\omega}_2^{-1}) = 1$, and thus $V(\widehat{\omega}_1) = V(\widehat{\omega}_2)$. Therefore V is indifferent with respect to $\mathcal{R}$.

Proof of Theorem 2. We will give a detailed proof only for the case of probability functions. The proofs for ordinal conditional functions and possibility distributions, respectively, follow the same idea and are indeed quite analogous.

Let P be a probability function and $\mathcal{R} = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ be a set of conditionals. Suppose first that P is indifferent with respect to $\mathcal{R}$. Then $P(A_i) \neq 0$, due to the prerequisite in Definition 4. The equivalence relation $\equiv_{\mathcal{R}}$ induces a partitioning $\Omega_1, \ldots, \Omega_q$ of Ω so that, according to Lemma 3, $P(\omega)$ is constant on each equivalence class. Assume $P(\omega) = p_j$ for $\omega \in \Omega_j$. Let $\omega_1, \ldots, \omega_q \in \Omega$ be a representative system of $\Omega_1, \ldots, \Omega_q$.

For the sake of simplicity of notation, we suppose that $p_1, \ldots, p_{q'} > 0$, $p_{q'+1} = \ldots = p_q = 0$ with $q' \leqslant q$.

For all $P(\omega_j) = p_j = 0, q' < j \leqslant q$, there is $(B_{i_j}|A_{i_j}) \in \mathcal{R}$ such that $\sigma_{(B_{i_j}|A_{i_j})}(\omega_j) \neq 1$ and $P(\omega') = 0$ for all ω' with $\sigma_{(B_{i_j}|A_{i_j})}(\omega') = \sigma_{(B_{i_j}|A_{i_j})}(\omega_j)$. If $\sigma_{(B_{i_j}|A_{i_j})}(\omega_j) = \mathbf{a}_{i_j}^+$ then set $\alpha_{i_j}^+ = 0$ and $\alpha_{i_j}^- = 1$; if $\sigma_{(B_{i_j}|A_{i_j})}(\omega_j) = \mathbf{a}_{i_j}^-$ then set $\alpha_{i_j}^+ = 1$ and $\alpha_{i_j}^- = 0$. Without loss of generality, assume that those conditionals $(B_{i_j}|A_{i_j}) \in \mathcal{R}$ are the conditionals $(B_i|A_i)$, $n' < i \leqslant n$.

Let us now consider the constants $p_j \neq 0$. Finding positive factors $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_{n'}^+, \alpha_{n'}^-$ with $0 \neq P(\omega) = \alpha_0 \prod_{\substack{1 \leqslant i \leqslant n' \\ \omega \models A_i B_i = 1}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n' \\ \omega \models A_i \overline{B_i}}} \alpha_i^-$ amounts to solving the following system of q' equations

$$\alpha_0 \prod_{\substack{1 \leqslant i \leqslant n' \\ \omega_j \models A_i B_i}} \alpha_i^+ \prod_{\substack{1 \leqslant i \leqslant n' \\ \omega_j \models A_i \overline{B_i}}} \alpha_i^- = p_j,\ j = 1, \ldots, q', \tag{19}$$

which can be transformed into a linear equational system

$$\Theta \vec{\beta} = \vec{\lambda} \tag{20}$$

with $\vec{\beta} = (\log \alpha_0, \log \alpha_1^+, \log \alpha_1^-, \ldots, \log \alpha_{n'}^+, \log \alpha_{n'}^-)^T \in \mathbb{R}^{2n'+1}$, $\vec{\lambda} = (\log p_1, \ldots, \log p_{q'})^T \in \mathbb{R}^{q'}$ and a $q' \times (2n'+1)$-matrix Θ with elements in $\{0,1\}$, such that $\theta_{j,1} = 1$ for all j, $\theta_{j,2i} = 1$ iff $\sigma_i(\omega_j) = \mathbf{a}_i^+$, $\theta_{j,2i+1} = 1$ iff $\sigma_i(\omega_j) = \mathbf{a}_i^-$ for $1 \leqslant j \leqslant q'$, $1 \leqslant i \leqslant n'$. Let $\vec{\theta}_j$, $1 \leqslant j \leqslant q'$, denote the rows of Θ. The equational system (20) is solvable over $\mathbb{R}$ iff any linear dependencies (over the field of rationals, because each entry of Θ is either 0 or 1) between these rows correspond to relations between the $\lambda_j = \log p_j$, i.e. $\sum_k r_{m_k} \vec{\theta}_{m_k} = \sum_l s_{n_l} \vec{\theta}_{n_l}$ must imply $\sum_k r_{m_k} \lambda_{m_k} = \sum_l s_{n_l} \lambda_{n_l}$ with rationals r_{m_k}, s_{n_l}.

Arranging and multiplying both sums appropriately, we may assume $\sum_k r_{m_k} \vec{\theta}_{m_k} = \sum_l s_{n_l} \vec{\theta}_{n_l}$ with natural numbers r_{m_k}, s_{n_l}. By comparing the vector components, we obtain $\sum_k r_{m_k} \theta_{m_k,2i} = \sum_l s_{n_l} \theta_{n_l,2i}$, $\sum_k r_{m_k} \theta_{m_k,2i+1} = \sum_l s_{n_l} \theta_{n_l,2i+1}$, $1 \leqslant i \leqslant n'$. These equations imply $\sum_k r_{m_k} = \sum_l s_{n_l}$, $\sum_{k:\sigma_i(\omega_{m_k}) = \mathbf{a}_i^+} r_{m_k} = \sum_{l:\sigma_i(\omega_{n_l}) = \mathbf{a}_i^+} s_{n_l}$

and $\sum_{k:\sigma_i(\omega_{m_k})=\mathbf{a}_i^-} r_{m_k} = \sum_{l:\sigma_i(\omega_{n_l})=\mathbf{a}_i^-} s_{n_l}$. Therefore the elements $\prod_k \omega_{m_k}^{r_{m_k}}$ and $\prod_l \omega_{n_l}^{s_{n_l}}$ are $\equiv_\top$-equivalent and $\mathcal{R}$-equivalent by equations (5) and (6), and because P is assumed to be indifferent with respect to $\mathcal{R}$, we obtain

$$\prod_k p_{m_k}^{r_{m_k}} = \prod_k P(\omega_{m_k})^{r_{m_k}} = \prod_l P(\omega_{n_l})^{s_{n_l}} = \prod_l p_{n_l}^{s_{n_l}}.$$

Applying the logarithm function now yields

$$\sum_k r_{m_k}\lambda_{m_k} = \sum_l s_{n_l}\lambda_{n_l},$$

as desired. Thus the equational system (20), or (19), respectively, is solvable, yielding a solution $\vec{\beta} = (\beta_0, \beta_1^+, \beta_1^-, \ldots, \beta_{n'}^+, \beta_{n'}^-)^T \in \mathbb{R}^{2n'+1}$. Setting $\alpha_0 = \exp(\beta_0)$, $\alpha_i^+ = \exp(\beta_i^+)$ and $\alpha_i^- = \exp(\beta_i^-)$, $1 \leqslant i \leqslant n'$, we obtain $P(\omega) = \alpha_0 \prod_{\substack{1\leqslant i\leqslant n' \\ \omega\models A_iB_i}} \alpha_i^+ \prod_{\substack{1\leqslant i\leqslant n' \\ \omega\models A_i\overline{B_i}}} \alpha_i^-$ for $P(\omega) \neq 0$. Taking now also into account the conditionals $(B_{n'+1}|A_{n'+1}), \ldots, (B_n|A_n)$, belonging to $P(\omega_j) = 0$, we thus have $P(\omega) = \alpha_0 \prod_{\substack{1\leqslant i\leqslant n \\ \omega\models A_iB_i}} \alpha_i^+ \prod_{\substack{1\leqslant i\leqslant n \\ \omega\models A_i\overline{B_i}}} \alpha_i^-$ for all $\omega \in \Omega$ because the non-zero factors belonging to those conditionals are 1.

To prove the converse assume $P(\omega) = \alpha_0 \prod_{\substack{1\leqslant i\leqslant n \\ \omega\models A_iB_i}} \alpha_i^+ \prod_{\substack{1\leqslant i\leqslant n \\ \omega\models A_i\overline{B_i}}} \alpha_i^-$ is a probability distribution with $\alpha_0, \alpha_1^+, \alpha_1^-, \ldots, \alpha_n^+, \alpha_n^- \in \mathbb{R}^+$, $\alpha_0 > 0$. We have to show the indifference of P with respect to R.

If $P(\omega) = 0$ then there is $(B_i|A_i) \in \mathcal{R}$ such that $\omega \models A_iB_i$ and $\alpha_i^+ = 0$, or $\omega \models A_i\overline{B}_i$ and $\alpha_i^- = 0$. So, in any case $\sigma_i(\omega) \neq 1$ and $P(\omega') = 0$ for any $\omega' \in \Omega$ with $\sigma_i(\omega') = \sigma_i(\omega)$. This shows condition (i) of Definition 4.

Now consider two $\mathcal{R}$-equivalent elements

$$\widehat{\omega}_1 = \prod_{k=1}^{m_1} \omega_k^{r_k} \quad \text{and} \quad \widehat{\omega}_2 = \prod_{l=1}^{m_2} \omega_l^{s_l} \in \widehat{\Omega}_+$$

with identical conditional structures

$$\sigma_{\mathcal{R}}(\widehat{\omega}_1) = \prod_{1\leqslant k\leqslant m_1} \sigma_{\mathcal{R}}(\omega_k)^{r_k} = \prod_{1\leqslant l\leqslant m_2} \sigma_{\mathcal{R}}(\nu_l)^{s_l} = \sigma_{\mathcal{R}}(\widehat{\omega}_2)$$

which are also $\equiv_\top$-equivalent. Then $\sum_k r_{m_k} = \sum_l s_{n_l}$, $\sum_{k:\sigma_i(\omega_k)=\mathbf{a}_i^+} r_k = \sum_{l:\sigma_i(\nu_l)=\mathbf{a}_i^+} s_l$ and $\sum_{k:\sigma_i(\omega_k)=\mathbf{a}_i^-} r_k = \sum_{l:\sigma_i(\nu_l)=\mathbf{a}_i^-} s_l$ hold for all $i = 1, \ldots, n$

according to equation (5). Checking condition (ii) of Definition 4 is now an easy calculation:

$$\begin{aligned}
P(\widehat{\omega}_1) &= P(\omega_1)^{r_1} \ldots P(\omega_{m_1})^{r_{m_1}} = \\
&= \alpha_0^{\sum_k r_{m_k}} \prod_{1 \leqslant i \leqslant n} \left(\alpha_i^+\right)^{\sum\limits_{k:\omega_k \models A_i B_i} r_k} \prod_{1 \leqslant i \leqslant n} \left(\alpha_i^-\right)^{\sum\limits_{k:\omega_k \models A_i \overline{B}_i} r_k} \\
&= \alpha_0^{\sum_l s_{n_l}} \prod_{1 \leqslant i \leqslant n} \left(\alpha_i^+\right)^{\sum\limits_{l:\nu_l \models A_i B_i} s_l} \prod_{1 \leqslant i \leqslant n} \left(\alpha_i^-\right)^{\sum\limits_{l:\nu_l \models A_i \overline{B}_i} s_l} \\
&= P(\nu_1)^{s_1} \ldots P(\nu_{m_2})^{s_{m_2}} = P(\widehat{\omega}_2).
\end{aligned}$$

Proof of Proposition 3. The proof of this proposition is tedious and technical, but straightforward. We will exemplify it for the case that $\mathbf{X}, \mathbf{Y}, \mathbf{Z}$ each contain just one (binary) variable: $\mathbf{X} = \{a\}$, $\mathbf{Y} = \{b\}$, and $\mathbf{Z} = \{c\}$. The corresponding set $\mathcal{R}$ then consists of the following eight conditionals:

$$\begin{aligned}
\mathcal{R} = \{ \quad & (a|c),\ (\overline{a}|c),\ (b|c),\ (\overline{b}|c), \\
& (a|\overline{c}),\ (\overline{a}|\overline{c}),\ (b|\overline{c}),\ (\overline{b}|\overline{c})
\end{aligned}$$

Let P be indifferent with respect to $\mathcal{R}$. We have to show that a and b are conditionally independent in P, given c, i.e. $P(ab|\dot{c}) = P(a|\dot{c})P(b|\dot{c})$, $\dot{c} \in \{c, \overline{c}\}$, which is equivalent to $\dfrac{P(ab\dot{c})P(\overline{a}\overline{b}\dot{c})}{P(a\overline{b}\dot{c})P(\overline{a}b\dot{c})} = 1$. Let $\mathbf{a}_1^{\pm}, \mathbf{a}_2^{\pm}, \mathbf{a}_3^{\pm}, \mathbf{a}_4^{\pm}$ be the group generators of $\mathcal{F}_{\mathcal{R}}$ associated with $(a|\dot{c})$, $(\overline{a}|\dot{c})$, $(b|\dot{c})$, $(\overline{b}|\dot{c})$, respectively. Then

$$\sigma_{\mathcal{R}}\left(\frac{ab\dot{c} \cdot \overline{a}\overline{b}\dot{c}}{a\overline{b}\dot{c} \cdot \overline{a}b\dot{c}}\right) = \frac{\mathbf{a}_1^+ \mathbf{a}_2^- \mathbf{a}_3^+ \mathbf{a}_4^- \cdot \mathbf{a}_1^- \mathbf{a}_2^+ \mathbf{a}_3^- \mathbf{a}_4^+}{\mathbf{a}_1^+ \mathbf{a}_2^- \mathbf{a}_3^- \mathbf{a}_4^+ \cdot \mathbf{a}_1^- \mathbf{a}_2^+ \mathbf{a}_3^+ \mathbf{a}_4^-} = 1,$$

and due to the indifference of P with respect to $\mathcal{R}$, we also have $\dfrac{P(ab\dot{c})P(\overline{a}\overline{b}\dot{c})}{P(a\overline{b}\dot{c})P(\overline{a}b\dot{c})} = 1.$

Proof of Proposition 4. Let $V^* = V * \mathcal{R} = V * (B|A)$ denote a revision of the conditional valuation function $V : \mathcal{L} \to \mathcal{A}$ by $\mathcal{R} = \{(B|A)\}$, and assume $V^*(A) \neq 0^{\mathcal{A}}$.

V^* satisfies the principle of conditional preservation with respect to V and $\mathcal{R}$ iff V^* is indifferent with respect to V and $\mathcal{R}$. According to Definition 7, this means in particular that V^* is V-consistent, and

$$(V^*/V)(\omega_1) = (V^*/V)(\omega_2) \quad \text{if} \quad (B|A)(\omega_1) = (B|A)(\omega_2) \tag{21}$$

for $V(\omega_1), V(\omega_2) \neq 0^{\mathcal{A}}$. Due to the prerequisite $V^*(A) \neq 0^{\mathcal{A}}$ and the V-consistency of V^*, we have $V(A) \neq 0^{\mathcal{A}}$, too, so $V(AB) \neq 0^{\mathcal{A}}$ or $V(A\overline{B}) \neq 0^{\mathcal{A}}$. If $V(AB) = 0^{\mathcal{A}}$, then $V^*(AB) = 0^{\mathcal{A}}$ and $V(A\overline{B}), V^*(A\overline{B}) \neq 0^{\mathcal{A}}$. In this case, there is $\omega^- \in Mod(A\overline{B})$ such that $V(\omega^-), V^*(\omega^-) \neq 0^{\mathcal{A}}$; set $\alpha^+ := 1^{\mathcal{A}}, \alpha^- := (V^*/V)(\omega^-)$. If $V(A\overline{B}) = 0^{\mathcal{A}}$, then analogically, $\alpha^- := 1^{\mathcal{A}}$ and $\alpha^+ := (V^*/V)(\omega^+)$ for some suitable $\omega^+ \in Mod(AB)$. If both $V(AB), V(A\overline{B}) \neq 0^{\mathcal{A}}$, then choose worlds $\omega^+ \in Mod(AB), \omega^- \in Mod(A\overline{B})$ such that $V(\omega^+), V(\omega^-) \neq 0^{\mathcal{A}}$ and set $\alpha^+ := (V^*/V)(\omega^+), \alpha^- := (V^*/V)(\omega^-)$. Furthermore, we have $V^*(\overline{A}) = 0^{\mathcal{A}}$ iff $V(\overline{A}) = 0^{\mathcal{A}}$; in this case, set $\alpha_0 := 1^{\mathcal{A}}$. Otherwise, select $\omega_0 \in Mod(\overline{A})$ with $V(\omega_0), V^*(\omega_0) \neq 0^{\mathcal{A}}$ and set $\alpha_0 := (V^*/V)(\omega_0)$. Due to equation (21), we thus have

$$V^*(\omega) = \begin{cases} \alpha^+ \odot V(\omega) & \text{if} \quad \omega \models AB \\ \alpha^- \odot V(\omega) & \text{if} \quad \omega \models A\overline{B} \\ \alpha_0 \odot V(\omega) & \text{if} \quad \omega \models \overline{A} \end{cases} \tag{22}$$

with (at least) $\alpha_0 \neq 0^{\mathcal{A}}$.

Conversely, any revision V^* of type (22) is V-consistent and satisfies Definition 6. Let $\widehat{\omega} = \omega_1^{r_1} \cdot \ldots \cdot \omega_m^{r_m} \in \widehat{\Omega}_+^*$; then

$$\sigma_{(B|A)}(\widehat{\omega}) = (\mathbf{a}^+)^{\sum_{k:\omega_k \models AB} r_k} (\mathbf{a}^-)^{\sum_{k:\omega_k \models A\overline{B}} r_k},$$

and

$$(V^*/V)(\widehat{\omega}) = (\alpha^+)^{\sum_{k:\omega_k \models AB} r_k} (\alpha^-)^{\sum_{k:\omega_k \models A\overline{B}} r_k} (\alpha_0)^{\sum_{k:\omega_k \models \overline{A}} r_k}.$$

Thus we see that V^* of type (22) is indifferent with respect to V and $(B|A)$. Furthermore, by the remarks above, it is clear that if V^* is strictly V-consistent, then all constants $\alpha_0, \alpha^+, \alpha^-$ can be chosen $\neq 0^{\mathcal{A}}$. This completes the proof.

Proof of Lemma 4. Suppose the revision $V^* = V * (B|A)$ is strictly V-consistent and satisfies (CR5^{quant}). Let $(D|C)$ be a conditional such that $(D|C) \perp\!\!\!\perp (B|A)$ and with $V(C) \neq 0^{\mathcal{A}}$. Since $V^* = V * (B|A)$ is strictly V-consistent, we also have $V^*(C) \neq 0^{\mathcal{A}}$, and $V(C\dot{D}) = 0^{\mathcal{A}}$ iff $V^*(C\dot{D}) = 0^{\mathcal{A}}$, $\dot{D} \in \{D, \overline{D}\}$. If $V(CD) = V^*(CD) = 0^{\mathcal{A}}$, then $V(D|C) = V^*(D|C) = 0^{\mathcal{A}}$; if $V(C\overline{D}) = V^*(C\overline{D}) = 0^{\mathcal{A}}$, then $V(D|C) = V^*(D|C) = 1^{\mathcal{A}}$.

So assume now $V(CD), V(C\overline{D}) \neq 0^{\mathcal{A}}$. Then, by (CR5^{quant}),

$$V(C\overline{D}) \odot V(CD)^{-1} = V^*(C\overline{D}) \odot V^*(CD)^{-1},$$

and consequently,

$$V(D|C) \ = \ V(CD) \odot V(C)^{-1}$$

$$
\begin{aligned}
&= V(CD) \odot (V(CD) \oplus V(C\overline{D}))^{-1} \\
&= V(CD) \odot V(CD)^{-1} \odot (1^{\mathcal{A}} \oplus V(C\overline{D}) \odot V(CD)^{-1})^{-1} \\
&= (1^{\mathcal{A}} \oplus V^*(C\overline{D}) \odot V^*(CD)^{-1})^{-1} \\
&= V^*(CD) \odot V^*(CD)^{-1} \odot (1^{\mathcal{A}} \oplus V^*(C\overline{D}) \odot V^*(CD)^{-1})^{-1} \\
&= V^*(CD) \odot (V^*(CD) \oplus V^*(C\overline{D}))^{-1} \\
&= V^*(CD) \odot V^*(C)^{-1} \\
&= V^*(D|C).
\end{aligned}
$$

Proof of Proposition 5. Let $V^* = V * (B|A)$ denote a strictly V-consistent revision of V by $(B|A)$ satisfying $V^*(A) \neq 0^{\mathcal{A}}$ and (CR5^{quant}). Suppose $(D|C) \perp\!\!\!\perp (B|A)$. If $V(CD) = V^*(CD) = 0^{\mathcal{A}}$, then neither V nor V^* accepts $(D|C)$. So let $V(CD), V^*(CD) \neq 0^{\mathcal{A}}$. Then (CR5^{quant}) implies

$$V(C\overline{D}) \odot V(CD)^{-1} = V^*(C\overline{D}) \odot V^*(CD)^{-1}. \tag{23}$$

According to Section 3, we have

$$
\begin{aligned}
V \models (D|C) \quad &\Leftrightarrow \quad V(C\overline{D}) <_{\mathcal{A}} V(CD) \\
&\Leftrightarrow \quad V(C\overline{D}) \odot V(CD)^{-1} <_{\mathcal{A}} 1^{\mathcal{A}} \\
&\Leftrightarrow \quad V^*(C\overline{D}) \odot V^*(CD)^{-1} <_{\mathcal{A}} 1^{\mathcal{A}} \quad \text{(due to (23))} \\
&\Leftrightarrow \quad V^*(C\overline{D}) <_{\mathcal{A}} V^*(CD) \\
&\Leftrightarrow \quad V^* \models (D|C).
\end{aligned}
$$

Thus (CR5) holds.

Proof of Theorem 4. Let V be a conditional valuation function, and let $V^* = V * \{(B|A)\}$ denote a strictly V-consistent revision of V by $(B|A)$ fulfilling the postulates (CR1) (success) and (CR2) (stability). So in particular, we have $V^*(A) \neq 0^{\mathcal{A}}$, and by the strict V-consistency of the revision, we also have $V(A) \neq 0^{\mathcal{A}}$.

If V^* satisfies the principle of conditional preservation, then, by Proposition 4, there exist constants $\alpha_0, \alpha^+, \alpha^- \neq 0^{\mathcal{A}}$ in $\mathcal{A}$ such that

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

To prove (CR5^{quant}), suppose that $(D|C) \perp\!\!\!\perp (B|A)$ and $V^*(CD) \neq 0^{\mathcal{A}}$. So $Mod(C)$ is completely included in one of $Mod(AB), Mod(A\overline{B})$,

$Mod(\overline{A})$. Then for a suitable $\alpha \in \{\alpha_0, \alpha^+, \alpha^-\}$, we obtain

$$\begin{aligned}
V^*(C\overline{D}) \odot V^*(CD)^{-1} &= \left(\sum_{\omega \models C\overline{D}}{}^{\oplus} V^*(\omega)\right) \odot \left(\sum_{\omega \models CD}{}^{\oplus} V^*(\omega)\right)^{-1} \\
&= \left(\sum_{\omega \models C\overline{D}}{}^{\oplus} \alpha \odot V(\omega)\right) \odot \left(\sum_{\omega \models CD}{}^{\oplus} \alpha \odot V(\omega)\right)^{-1} \\
&= \left(\alpha \odot \sum_{\omega \models C\overline{D}}{}^{\oplus} V(\omega)\right) \odot \left(\alpha \odot \sum_{\omega \models CD}{}^{\oplus} V(\omega)\right)^{-1} \\
&= \alpha \odot V(C\overline{D}) \odot \alpha^{-1} \odot V(CD)^{-1} \\
&= V(C\overline{D}) \odot V(CD)^{-1}
\end{aligned}$$

This shows ($\mathrm{CR5}^{quant}$).

Suppose now $(D|C) \sqsubseteq (B|A)$, i.e. $CD \leqslant AB$ and $C\overline{D} \leqslant A\overline{B}$. Then, as in the calculations above, we obtain $V^*(AB) = \alpha^+ \odot V(AB), V^*(A\overline{B}) = \alpha^- \odot V(A\overline{B})$ and $V^*(CD) = \alpha^+ \odot V(CD), V^*(C\overline{D}) = \alpha^- \odot V(C\overline{D})$. Furthermore, $V(CD) \leqslant V(AB)$ and $V(C\overline{D}) \leqslant V(A\overline{B})$.

By prerequisite, $V^* \models (B|A)$, thus $V^*(A\overline{B}) <_{\mathcal{A}} V^*(AB)$. If $V \models (B|A)$, then, by (CR2), $V = V^*$, and (CR6), (CR7) are trivially fulfilled.

So assume now that $V \not\models (B|A)$, that is, $V(AB) \leqslant_{\mathcal{A}} V(A\overline{B})$. From $V^* \models (B|A)$, we have $\alpha^- \odot V(A\overline{B}) <_{\mathcal{A}} \alpha^+ \odot V(AB)$ which implies $\alpha^- <_{\mathcal{A}} \alpha^+$. If $V \models (D|C)$, this yields

$$V^*(C\overline{D}) = \alpha^- \odot V(C\overline{D}) <_{\mathcal{A}} \alpha^+ \odot V(C\overline{D}) <_{\mathcal{A}} \alpha^+ \odot V(CD) = V^*(CD),$$

hence $V^* \models (D|C)$. This shows (CR6).

To prove (CR7), suppose $(D|C) \sqsubseteq (\overline{B}|A)$, $V \not\models (B|A)$ and $V^* \models (D|C)$, i.e. $V^*(C\overline{D}) <_{\mathcal{A}} V^*(CD)$. Then $\alpha^+ \odot V(C\overline{D}) <_{\mathcal{A}} \alpha^- \odot V(CD)$, and consequently, by using $\alpha^- <_{\mathcal{A}} \alpha^+$, $V(C\overline{D}) <_{\mathcal{A}} V(CD)$, which means $V \models (D|C)$. This shows (CR7).

References

Alchourrón, C., P. Gärdenfors, and P. Makinson: 1985, 'On the logic of theory change: Partial meet contraction and revision functions'. *Journal of Symbolic Logic* **50**(2), 510–530.

Beierle, C. and G. Kern-Isberner: 2003, 'Modelling Conditional Knowledge Discovery and Belief Revision by Abstract State Machines'. In: E. Boerger, A. Gargantini, and E. Riccobene (eds.): *Abstract State Machines 2003 – Advances in Theory and Applications, Proceedings 10th International Workshop, ASM2003*. pp. 186–203, Springer, LNCS 2589.

Benferhat, S., D. Dubois, and H. Prade: 1992, 'Representing default rules in possibilistic logic'. In: *Proceedings 3th International Conference on Principles of Knowledge Representation and Reasoning KR'92*. pp. 673–684.

Benferhat, S., D. Dubois, and H. Prade: 1997, 'Nonmonotonic reasoning, conditional objects and possibility theory'. *Artificial Intelligence* **92**, 259–276.

Benferhat, S., A. Saffiotti, and P. Smets: 2000, 'Belief functions and default reasoning'. *Artificial Intelligence* **122**, 1–69.

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

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

Cowell, R., A. Dawid, S. Lauritzen, and D. Spiegelhalter: 1999, *Probabilistic networks and expert systems*. New York Berlin Heidelberg: Springer.

Csiszár, I.: 1975, 'I-divergence geometry of probability distributions and minimization problems'. *Ann. Prob.* **3**, 146–158.

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

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

Delgrande, J. and T. Schaub: 2000, 'A Consistency-Based Model for Belief Change: Preliminary Report'. In: *Proceedings AAAI-2000*. Austin, TX, pp. 392–398.

Dubois, D., J. Lang, and H. Prade: 1994, 'Possibilistic logic'. In: D. Gabbay, C. Hogger, and J. Robinson (eds.): *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3. Oxford University Press.

Dubois, D. and H. Prade: 1991, 'Conditioning, non-monotonic logic and non-standard uncertainty models'. In: I. Goodman, M. Gupta, H. Nguyen, and G. Rogers (eds.): *Conditional Logic in Expert Systems*. Elsevier, North Holland, pp. 115–158.

Dubois, D. and H. Prade: 1994, 'A survey of belief revision and updating rules in various uncertainty models'. *Intern. Journal of Intelligent Systems* **9**, 61–100.

Fine, B. and G. Rosenberger: 1999, *Algebraic Generalizations of Discrete Groups*. New York, Basel: Dekker.

Freund, M.: 1998, 'Preferential orders and plausibility measures'. *J. Logic Computat.* **8**, 147–158.

Friedman, N. and J. Halpern: 1996, 'Plausibility measures and default reasoning'. In: *Proceedings 13th National Conference on Artificial Intelligence, AAAI-96*, Vol. 2.

Gärdenfors, P.: 1988, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge, Mass.: MIT Press.

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

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

Katsuno, H. and A. Mendelzon: 1991a, '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. San Mateo, Ca., pp. 387–394, Morgan Kaufmann.

Katsuno, H. and A. Mendelzon: 1991b, '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, 'Postulates for conditional belief revision'. In: *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI-99*. pp. 186–191, Morgan Kaufmann.

Kern-Isberner, G.: 1999b, *A unifying framework for symbolic and numerical approaches to nonmonotonic reasoning and belief revision*. Department of Computer Science, FernUniversität Hagen. Habilitation thesis.

Kern-Isberner, G.: 2000, 'Solving the inverse representation problem'. In: *Proceedings 14th European Conference on Artificial Intelligence, ECAI'2000*. Berlin, pp. 581–585, IOS Press.

Kern-Isberner, G.: 2001a, 'Conditional indifference and conditional preservation'. *Journal of Applied Non-Classical Logics* **11**(1-2), 85–106.

Kern-Isberner, G.: 2001b, 'Conditionals in knowledge representation and belief revision'. In: *Proceedings Fifth Dutch-German Workshop on Nonmonotonic Reasoning Techniques and their applications, DGNMR'01*. Potsdam, Germany.

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

Kern-Isberner, G.: 2001d, 'Handling conditionals adequately in uncertain reasoning'. In: *Proceedings European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, ECSQARU'01*. pp. 604–615, Springer LNAI 2143.

Kern-Isberner, G.: 2001e, 'Representing and learning conditional information in possibility theory'. In: *Proceedings 7th Fuzzy Days, Dortmund, Germany*. pp. 194–217, Springer LNCS 2206.

Kern-Isberner, G.: 2002a, 'The principle of conditional preservation in belief revision'. In: *Proceedings of the Second International Symposium on Foundations of Information and Knowledge Systems, FoIKS 2002*. pp. 105–129, Springer LNCS 2284.

Kern-Isberner, G.: 2002b, 'A structural approach to default reasoning'. In: *Proceedings of the Eighth International Conference on Principles of Knowledge Representation and Reasoning, KR'2002*. San Francisco, Ca., pp. 147–157, Morgan Kaufmann.

Kruse, R., E. Schwecke, and J. Heinsohn: 1991, *Uncertainty and Vagueness in Knowledge Based Systems*. Berlin Heidelberg New York: Springer.

Levi, I.: 1988, 'Iteration of conditionals and the Ramsey test'. *Synthese* **76**, 49–81.

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

Makinson, D. and P. Gärdenfors: 1991, 'Relations between the logic of theory change and nonmonotonic logic'. In: *Proceedings Workshop The Logic of Theory Change, Konstanz, Germany, 1989*. Berlin Heidelberg New York, pp. 185–205, Springer.

Paris, J.: 1994, *The uncertain reasoner's companion – A mathematical perspective*. Cambridge University Press.

Paris, J. and A. Vencovská: 1992, 'A method for updating that justifies minimum cross entropy'. *International Journal of Approximate Reasoning* **7**, 1–18.

Pearl, J.: 1988, *Probabilistic Reasoning in Intelligent Systems*. San Mateo, Ca.: Morgan Kaufmann.

Ramsey, F.: 1950, 'General propositions and causality'. In: R. Braithwaite (ed.): *Foundations of Mathematics and other logical essays.* New York: Routledge and Kegan Paul, pp. 237–257.

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

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

Walker, E.: 1994, 'Stone Algebras, Conditional Events, and Three Valued Logic'. *IEEE Transactions on Systems, Man, and Cybernetics (Special Issue on The Conditional Event Algebra)* **24**(12), 1699–1707.