JZBR - Iterated Belief Change for Conditional Ranking Constraints

An early report

Emil Weydert
Max-Planck-Institute for Computer Science
Im Stadtwald, D-66123 Saarbrücken, Germany
emil@mpi-sb.mpg.de

Abstract

We present JZBR, an extension of Spohn's iterated revision methodology which allows the handling of new information expressed by conditional ranking measure constraints. This is achieved by generalizing a canonical ranking construction technique originally developed for nonmonotonic entailment, notably system JZ. JZBR exploits the structure of the evidence base and is a robust relative of semi-qualitative cross entropy minimization.

1 INTRODUCTION

Rational belief change in the face of new, potentially conflicting information is a major characteristic of intelligent behaviour. Its formal modeling therefore has been - and still is - a major issue in artificial intelligence. Although many problems are far from being settled, important progress has been made in the past twenty years. Most of these efforts have been directly linked to or at least inspired from the work of Peter Gärdenfors. The famous AGM-postulates [AGM 1985], which he presented together with his colleagues Carlos Alchourron and David Makinson, are among the most influential principles in AI. His book *Knowledge in Flux* [Gärdenfors 88] has been - and still is - an important compendium and source of ideas for anyone interested in belief revision. In fact, ten years ago, a corresponding reading group in Bonn had quite an impact on my career, and my postdoc Richard Booth just bought a copy on the internet.

Of course, research has gone beyond his initial framework, and Peter Gärdenfors was among the first to acknowledge the necessity for doing so. I first experienced his friendly open-mindedness at the European Workshop for Logics in AI in 1990 (JELIA 90), where I gave a talk - my first one in AI - on doxastic preference logic [Weydert 90]. In this work, I criticized AGM's "seven flaws", namely its problems with absolute knowledge, iterated revision, introspection, believing defaults, nonlinear epistemic orderings, recovery, and practicality. After my presentation, he came to me and frankly told me, ... that I was right. Frankly, I was impressed.

In this paper, we are going to focus on iterated belief revision and default beliefs. Traditional approaches, as exemplified by the original AGM-paradigm, have been mainly concerned with one-step belief revision. That is, given - for instance - a set of ordered beliefs together with new propositional evidence, they only tell us how to revise the belief set, but not how to revise the the epistemic entrenchment ordering, i.e. the control structure. This, however, would be required for implementing iterated belief change in a reasonable way, obviously a necessity for real world agents.

Taking a broader perspective, we see that similar questions already have been addressed in probabilistic reasoning. According to the Bayesian paradigm, an epistemic state should be modeled by a probability distribution over worlds. If we assume that new evidence can be translated into a probabilistic equality constraint $P(A) = \alpha$, we may use Jeffrey conditionalization - enforcing the equality and preserving the conditional probabilities $P(.|A), P(.|\neg A)$ - to update the probabilistic belief state in a canonical way. However, this doesn't allow us to revise full belief, i.e. belief closed under conjunctions, which is expressed by $P(A) = 1$.

This problem can be attacked within the semi-qualitative context of Spohn's κ-ranking framework, which offers a popular and powerful procedure for iterated belief revision exploiting order-of-magnitude probabilities [Spohn 90]. More recently, there have been several other proposals for revising prioritized epistemic structures [Boutilier 92, Darwiche and Pearl 97, Lehmann 95, Williams 94]. One advantage of Spohn's revision paradigm - if we accept his model of belief - is the information-theoretic justification of probabilistic Jeffrey-conditionalization. In fact, among all those measures satisfying the probabilistic constraint encoding the evidence, it picks up the one with the least information content relative to the prior subjective probability measure. That is, it constitutes a special instance of cross-entropy minimization (MCE), a general, well-justified probabilistic update procedure [Shore and Johnson 80, Paris and Vencovska 90] which is particularly appropriate for linear constraints.

What is currently missing is a generalization of Spohn's revision strategy able to deal with more general epistemic ranking constraints, in particular with sets of conditional ranking constraints, like $R(A|B) \geq \alpha$. A brute force method would be to translate these conditions into infinitesimal probability constraints, like $P(A|B) \leq \varepsilon^{\alpha}$, to apply MCE in this nonstandard probability context [Weydert 95], and to translate the resulting nonstandard distribution back into an extended ranking framework. However, this approach is inadequate for several reasons.

First of all, the probabilistic translation procedure is inherently ambiguous. For instance, although the infinitesimals ε^{α} and $(2\varepsilon)^{\alpha}$ are practically indiscernible and have the same order of magnitude α, using one or the other bound for constraining $P(A|B)$ may produce completely different results. Similarly, the solution may also change if we refine the boolean propositional domain of the belief valuations. Secondly, nonstandard cross-entropy minimization is a rather cumbersome process, difficult to compute. It is not the canonical construction procedure we are looking for.

In this paper, we are going to motivate and describe a more direct revision procedure for epistemic rankings which extends Spohn's account and is able to deal with finite sets of conditional ranking constraints. Our approach is based on the generalization of a ranking construction technique first exploited in nonmonotonic entailment, namely system JZ [Weydert 98]. The idea is to construct a canonical model of the given constraints through iterated Jeffrey conditionalization trying to minimize the construction efforts.

The plan of the paper is as follows. We begin by introducing the basic ranking measure framework and the corresponding epistemic transformations. In this context, we present the epistemic construction paradigm and implement complementary ideas of minimality. Based on this, we develop our revision strategy JZBR. This is done in several steps, considering increasingly sophisticated constraint sets. After illustrating the JZBR-technique with some examples, we discuss its relationship with other approaches and indicate perspectives for future work.

2 STATES AND CONSTRAINTS

In the literature, epistemic states are often rudimentarily characterized by epistemic preference relations or valuations indicating the relative degree of plausibility attributed to different propositions. However, to guarantee the comparability of different belief states and to obtain a natural notion of independency (e.g. for constructing belief networks), we also need a reasonable notion of conditional measures. This leaves us with two basic alternatives. On the fine-grained level, standard and nonstandard (admitting infinitesimals) probability distributions, on the coarse-grained level, order of magnitude or ranking measures. In this paper, we are primarily going to be concerned with standard $\kappa\pi$-ranking measures.

Definition 2.1 (Standard $\kappa\pi$-measure)
Let $\mathcal{B} \subseteq 2^{\mathcal{W}}$ be a compact boolean algebra of propositions $\mathcal{B}$ over a world set $\mathcal{W}$. $R : \mathcal{B} \to \mathcal{V} = ([0, \infty], 0, \infty, +, <)$ is called a standard $\kappa\pi$-measure iff

- $R(A \cup B) = Min\{R(A), R(B)\}$,
- $R(\mathcal{W}) = 0$, $R(\emptyset) = \infty$.

The conditional $\kappa\pi$-measure $R(\ |\ )$ is defined by $R(B|A) = R(B \cap A) - R(A)$ for $R(A) \neq \infty$, and $R(B|A) = \infty$ for $R(A) = \infty$. R_0 denotes the uniform $\kappa\pi$-measure on $\mathcal{B}$, with $R_0(A) = 0$ for all $A \neq \emptyset$.

Standard $\kappa\pi$-measures constitute a necessary generalization of Spohn's discrete-valued κ-measures and Dubois and Prade's possibility measures. They allow a simple and robust modeling of graded full belief, while preserving many concepts, tools and features from classical probability theory. Following Spohn, $\kappa\pi$-measures are assumed to measure the degree of disbelief, implausibility or surprise. Accordingly, the degree of belief in A corresponds to the degree of surprise of $\neg A$. A is believed with strength α (at least) iff $R(\neg A) \geq \alpha$. It is believed iff $R(\neg A) > 0$. $\kappa\pi$-values have no obvious real world meaning - comparable to the frequentist reading of probabilities - but they may be interpreted as the order-of-magnitude of nonstandard probabilities. More precisely, $R(A) = \alpha < \infty$ would correspond to $-\log_{\varepsilon} P(A) \approx \alpha$, $R(A) = \infty$ to $P(A) = 0$.

For our present purposes, iterated belief revision in a mainly propositional context, we may impose two further simplifying conditions on epistemic $\kappa\pi$-measure states. First of all, we want them to be finitely representable. Although the domain $\mathcal{B}$ of R is allowed to be infinite, there should be a finite boolean subalgebra $\mathcal{B}_0$ of $\mathcal{B}$ s.t. for all $A \in \mathcal{B}$, $R(A) = inf\{R(B) \mid B \in Atom_{\mathcal{B}_0}, A \cap B \neq \emptyset\}$. Furthermore, we will also restrict ourselves to $\kappa\pi$-measures taking only rational values. Let $\mathcal{KP}_{\mathcal{B}}$ be the set of all finitely representable rational $\kappa\pi$-measures over the compact boolean algebra $\mathcal{B}$. In the following, we will assume that $\mathcal{B}$ is the model set algebra for propositional logic with infinitely many variables, although we could also have chosen first-order predicate logic. $\mathcal{KP}_{\mathcal{B}}$ will then be our set of epistemic states.

When modeling or discussing belief revision, it is important to distinguish between, on one hand, the incoming information, the general sensory or cognitive input, and on the other hand, the interpreted, processed information translated into an epistemic constraint. That is, we have a two phase process. First, the analysis, interpretation and deliberation of the epistemic input, which produces an epistemic constraint. Secondly, the actual revision process which transforms the initial epistemic valuation in a reasonable way so as to realize this constraint. We may illustrate this idea in the context of Spohn's framework. here, for instance, given a $\kappa\pi$-measure R and new propositional evidence A, we first translate the input A - by default - into the epistemic constraint $R(\neg A) \geq 1$ (interpretation), which we are then going to enforce by a minimal amount of Jeffrey-conditionalization for A and $\neg A$ (revision). Although both steps are equally important, in this paper, we concentrate on the proper revision process, i.e. the update of epistemic measures by epistemic constraints. Object-level ranking constraints are written $r(A|B) \geq \alpha$ or $r(A|B) = \alpha$.

Whereas the original approach only considered constraints of the form $r(A) \geq \alpha$, we are interested in arbitrary finite sets of conditional ranking constraints $\Delta = \{r(S_i|S'_i) \geq \alpha_i \mid i \leq n\} \cup \{r(S_i|S'_i) = 0 \mid i \leq n\}$, expressing the relative graded belief, or the non-belief, in the proposition $\neg S_i$ given S'_i. These constraints can also be written as $r(A) + \alpha \leq r(A')$ with $0 \leq \alpha$ and $A \cap A' = \emptyset$. That is, we may concentrate on knowledge bases of the form $\Delta = \{r(A_i) + \alpha_i \leq r(A'_i) \mid i \leq n\}$. Let $\mathcal{I}_{\mathcal{B}}$ be the set of all these Δ over $\mathcal{B}$, and for each $\Delta \in \mathcal{I}_{\mathcal{B}}$, let $\mathcal{KP}_{\mathcal{B}}(\Delta)$ be the set of Δ-models in $\mathcal{KP}_{\mathcal{B}}$. Where possible, we shall drop $\mathcal{B}$.

Our task now is to find a revision function which associates with every coherent pair $(R, \Delta) \in \mathcal{KP} \times \mathcal{I}$ - i.e. there is $R' \in \mathcal{KP}(\Delta)$ s.t. for all $B \in \mathcal{B}$, $R(B) = \infty$ implies $R'(B) = \infty$ - a natural update $R' \in \mathcal{KP}$, written $R[\Delta]$. Note that we have explicitly chosen to take into account the inner structure of the epistemic input by using Δ, and not $\mathcal{KP}(\Delta)$, as our argument. Whereas in base revision one considers the representational structure of the knowledge base, here we consider the representational structure of the incoming evidence, which may convey additional implicit information, e.g. default assumptions about different sources for different evidence items. For instance, we may want to distinguish $\Delta = \{r(A) \geq 1, r(A') \geq 1\}$ and $\Delta' = \{r(A \vee A') \geq 1\}$, although they are verified by the same $\kappa\pi$-measures. But whereas $R_0[\Delta][r(\neg A) \geq 1]$ would support belief in $\neg A'$, $R_0[\Delta'][r(\neg A) \geq 1]$ would not, both being instances of revision with A. A more appropriate semantic description of the epistemic inputs therefore might be $\{\mathcal{KP}(\{\delta\}) \mid \delta \in \Delta\}$.

3 RANKING CONSTRUCTIONS

Most approaches to belief revision are based on some kind of minimal change principle. That is, the revised belief state should be the closest, most similar or most easily accessible one satisfying the epistemic constraints encoding the new information. This will also be our strategy, but the question is what it means in our context. Let R and Δ be coherent. Roughly speaking, the idea is to choose for $R[\Delta]$ that $R' \in \mathcal{KP}(\Delta)$ which can be obtained from R by a minimal epistemic ranking model construction process. In particular, we are looking for an approach extending Spohn-type revision and inspired from the minimal information paradigm (MCE).

The most basic epistemic transformation for $\kappa\pi$-measures is shifting, also known as L-conditionalization [Goldszmidt and Pearl 96]. Given a prior $\kappa\pi$-measure R, shifting a proposition A upwards or downwards means increasing or decreasing its degree of surprise, i.e. strengthening or weakening belief in $\neg A$, without modifying the conditional structure $R(.|A)$ and $R(.|\neg A)$. To ensure that the resulting valuation is still a $\kappa\pi$-measure ($min\{R(A), R(\neg A)\} = 0$), if $R(\neg A) > 0$, we may first have to move $\neg A$ to the bottom before we can shift A to its intended value. Note that shifting steps are Jeffrey-conditionalization steps.

Definition 3.1 (L-conditionalization/shifting)
Let R be a $\kappa\pi$-measure on $\mathcal{B}$, $A \in \mathcal{B}$, and $a \in [-\infty, \infty]$.

If $R(\neg A) = \infty$, $R[A + \infty] = R$ *(blocking).*

If $R(\neg A) \neq \infty$ *and* $0 \leq a$, $R[A + a]$ *is the unique* $\kappa\pi$*-measure* R' *such that*

- $R'(B|A) = R(B|A)$, $R'(B|\neg A) = R(B|\neg A)$, *for all* $B \in \mathcal{B}$,
- $R'(\neg A) \leq R(\neg A)$ *and* $R(A) \leq R'(A)$,
- $(R(\neg A) - R'(\neg A)) + (R'(A) - R(A)) = \alpha$.

If $R(\neg A) \neq \infty$ *and* $a \leq 0$, $R[A + a] = R[\neg A - a]$.

Shifting is commutative as long as blocking does not occur. Let us call a $\kappa\pi$-measure R' *epistemically constructible* from R iff it is the result of an iterated shifting process starting at R, i.e. $R' = R[A_0 + a_0] \ldots [A_n + a_n]$. Note that the $\kappa\pi$-measures in $\mathcal{KP}$ are exactly those which are epistemically constructible from R_0 with rational parameters a_i. However, for arbitrary $\kappa\pi$-measures, constructibility may fail. On an informal level, our revision paradigm can now be described as follows.

Minimal epistemic constructibility. *For coherent* $(R, \Delta) \in \mathcal{KP} \times \mathcal{I}$, $R[\Delta]$ *should be that model of* Δ *which is minimally epistemically constructible from* R.

Minimality can be implemented in different ways. First of all, we may restrict the set of admissible transformations $[A + a]$, which is also necessary to avoid trivialization. To achieve this, we exploit the representational structure of Δ. In our framework, an epistemic constraint plays two roles. On one hand, as usual, it describes a set of $\kappa\pi$-models. On the other hand, it is also meant to sanction epistemic construction steps. The constraint $r(A) + \alpha \leq r(B) \in \mathcal{I}$, for instance, may be seen as a permission for shifting A downwards and B upwards. This guarantees that the corresponding transformations $[A + p]$, for $p \leq 0$, and $[B + s]$, for $0 \leq s$, will preserve the semantic condition, once it has been realized. Other shifts are not allowed, which keeps minimal the set of admissible transformations. R' is said to be *epistemically constructible* from R over Δ iff it can be obtained through transformations sanctioned by Δ.

Minimizing efforts also means preventing redundant, unnecessary, or unjustified construction steps. This calls for another basic requirement, first discussed in [Weydert 96].

Definition 3.2 (Justifiable constructibility)
R' *is justifiably constructible from* R *w.r.t.* $\Delta = \{r(A_i) + a_i \leq r(A'_i) \mid i \leq n\}$ *iff*

- R' *is epistemically constructible from* R *w.r.t.* Δ, *i.e.*
  $R' = R[A_0 + p_0][A'_0 + s_0] \ldots [A_n + p_n][A'_n + s_n]$ *for* $p_i \leq 0 \leq s_i$,
- $R'(A_i) + a_i < R'(A'_i)$ *implies* $p_i = s_i = 0$.

Let $Constr(R, \Delta)$ *be the set of* R' *which are justifiably constructible from* R *w.r.t.* Δ.

In other words, shifting only occurs if there is no overachievement of the corresponding sanctioning constraint, that is if it is realized as an equality constraint.

Justifiable constructibility.
$R[\Delta]$ *should be justifiably constructible from* R *w.r.t.* Δ.

This notion is supported by the following existence theorem, which will be validated by our JZBR-construction.

Theorem 3.3 (Accessibility)
Let $(R, \Delta) \in \mathcal{KP} \times \mathcal{I}$. *Then*

- (R, Δ) *coherent implies* $\mathcal{KP}(\Delta) \cap Constr(R, \Delta) \neq \emptyset$.

The set $\mathcal{KP}(\Delta) \cap Constr(R, \Delta)$ provides a prelimary set of candidates for $R[\Delta]$ from which we will pick up the closest or most easily accessible one.

4 JZBR - SINGLE CONSTRAINT

We are now going to built up our canonical revision strategy step by step, identifying and implementing basic principles of minimal epistemic constructions while considering increasingly complex epistemic constraint bases. This will be done by generalizing the JZ-shifting formalism sketched in [Weydert 98]. In the following, we assume that R is the prior state and Δ a default knowledge base coherent with R. We start with the simplest instance, when there is only a single constraint.

- $\Delta = \{r(A) + a \leq r(A')\}$.

If R already verifies Δ, minimal change enforces $R[\Delta] = R$. If R doesn't, we have several choices - shifting A' upwards, shifting A downwards, or both. This makes a difference if $A \cup A' \neq \mathcal{W}$. Whatever the exact procedure, minimality suggests that the shifting process should stop as soon as the constraint is satisfied, which is when $r(A) + a = r(A')$ becomes valid. That is, after shifting A or A', the corresponding equality constraint should be realized. Note that this is just what justifiable constructibility requires.

This leaves us with the question about which propositions to shift, A or A', and to what extent. We want to minimize the shifting efforts, but what does this mean in practice ? Here we are guided by two basic principles, which can be justified and validated by the cross entropy distance measure.

- **Bottom before top.**
  Making a more plausible A less plausible is more costly than making a less plausible A less plausible. Because it seems reasonable first to minimize the most expensive - but also relevant - tasks, and because we prefer an incremental approach without backtracking, we should start at the bottom.

- **Upwards before downwards.**
  Making A more implausible is less costly than reversing this act and making A again more plausible by the same amount. In particular, shifting A to ∞ is infinitely less costly than the converse, which is impossible.

It follows from these considerations that the best thing to do for a single constraint is shifting A' upwards (Note that $R[.]$ may take different meanings in our notation).

- $R[r(A) + a \leq r(A')] = R[A' + \alpha]$, where $\alpha = a + R(A) - R(A')$.

This solution corresponds to that proposed by cross entropy minimization translated into the ranking context.

5 JZBR - JZ-CONSTRAINTS

The handling of constraint bases with several elements is considerably more difficult because satisfying one constraint may interact with the satisfaction of other constraints. To begin with, we restrict ourselves to constraint sets of the form

- $\Delta = \{r(A_i) + a_i \leq r(A'_i) \mid i \leq n\}$ with $0 < a_i$ and $R(\mathcal{W} - \bigcup\{A'_i \mid i \leq n\}) = 0$.

This special instance of the problem is within the reach of the construction techniques used for system JZ, a powerful default formalism [Weydert 98]. The JZ-construction procedure combines the *minimal epistemic construction* paradigm with the *normality maximization* philosophy. For any consistent $\Delta \in \mathcal{I}$, normality maximization means picking up the - argument-wise - $\leq$-smallest $\kappa\pi$-measure in $\mathcal{KP}(\Delta)$. Here, we are more particularly interested in relative normality maximization, i.e. normality maximization above some reference measure R. If $\{R' \mid R \leq R' \in \mathcal{KP}(\Delta)\} \neq \emptyset$, we set

- $NM(R, \Delta) = Min\{R' \in \mathcal{KP} \mid R \leq R' \in \mathcal{KP}(\Delta)\}$.

$NM(R, \Delta)$ would be the ideal candidate for $R[\Delta]$ - if it were also justifiably constructible w.r.t. Δ. Most often, however, $NM(R, \Delta)$ is not even epistemically constructible w.r.t. Δ. Nevertheless, we may try to exploit and approximate normality maximization in the epistemic construction process. More precisely, we are going to use it to determine optimal reference marks guiding the construction steps.

Another principle is *minimal shifting* or *uniformity maximization*. Whereas normality maximization tries to minimize the absolute $\kappa\pi$-values, minimal shifting tries to minimize the shifting lengths for parallel shifting tasks with the same goal value, that is constructing the same level. This makes the shifting process more uniform.

JZ-Algorithm.
We are now going to sketch a streamlined version of the original JZ-procedure able to start from arbitrary *kp*-measures. It allows us to revise constraint sets of the above type. Let R and Δ be coherent. The idea is to pass from R to $R[\Delta]$ through an inductive justifiable bottom-up construction process, minimizing the shifting efforts at each level, i.e. maximizing normality and uniformity. More concretely, we construct a sequence of $\kappa\pi$-measures $(R^j \mid j \leq h)$ with

- $R^0 = R$, $R^{j+1} = R^j[A'_i + s_i \mid i \in I_{j+1}]$,

where $I_1 \cup \ldots \cup I_h$ is a partition of $I = \{i \mid i \leq n\}$ and the shifting parameters $0 \leq s_i$ are rational numbers. Let $\Delta_j = \{\delta_i \in \Delta \mid i \in I - I_1 \cup \ldots \cup I_j\}$, where $\delta_i = r(A_i) + a_i \leq r(A'_i)$, e.g. $\Delta_0 = \Delta$. Δ_j is the set of those constraints which have not yet been taken into account at step j. At step $j+1$, we first want to satisfy the weakest constraints in Δ_j, i.e. those affecting the lowest possible levels, and therefore the least dependent ones. The strength of a constraint δ_i in the context of R^j and Δ is measured by the minimal possible value A'_i could take in a model of Δ above R^j, which is just $NM(R^j, \Delta)(A'_i)$. Let

- $\alpha_{j+1} = Min\{NM(R^j, \Delta)(A'_i) \mid \delta_i \in \Delta_j\}$,
- $I_{j+1} = \{i \notin I_1 \cup \ldots \cup I_j \mid NM(R^j, \Delta)(A'_i) = \alpha_{j+1}$.

Our task is now to determine $R^{j+1} = R^j[A'_i + s_i \mid i \in I_{j+1}]$ in a way which guarantees $NM(R^{j+1}, \Delta_{j+1}) = NM(R^{j+1}, \Delta)$, justifiable constructibility for δ_i with $i \in I_{j+1}$, and which minimizes the required shifting effort. Since we assume $0 < a_i$, there is no need for shifting downwards any A_i. The s_i for $i \in I_{j+1}$ are chosen as follows (for existence, see [Weydert 98]). First, we set $s_i = s$ $(i \leq n)$, where s should verify for all $i \in I_{j+1}$,

- $NM(R^j[A'_i + s \mid i \in I_{j+1}], \Delta_{j+1})(A'_i) \geq \alpha_{j+1}$.

Secondly, we make all the s_i uniformly smaller until

- $NM(R^j[A'_i + s_i \mid i \in I_{j+1}], \Delta_{j+1})(A'_i) = \alpha_{j+1}$ or $s_i = 0$.

This is equivalent to minimize the largest s_i lexicographically. Let h be the first j where $\Delta_j = \emptyset$. Then we have

- $R^h = NM(R^h, \emptyset) = NM(R^h, \Delta_h) = NM(R^h, \Delta)$.

That is, R^h is a model of Δ. It is justifiably constructible because $NM(R^j, \Delta_j)(A'_i)$ is stable for shifted A'_i with $i \in I_1 \cup \ldots \cup I_j$. It follows from the construction that the s_i are rational and that R^h is finitely representable, i.e. $R^h \in \mathcal{KP}(\Delta)$. Under the above conditions, i.e. consistent belief strengthening and conditional ranking constraints with non-zero bounds, the relative JZ-model R^h will be the intended revised $\kappa\pi$-state $R[\Delta]$.

- $R[\Delta] = R^h$.

If $R = R_0$, $R[\Delta] = JZ[\Delta]$.

6 JZBR - ANY CONSTRAINTS

The standard JZ-revision procedure does not work for non-consistent revision, i.e. if $R(\mathcal{W} - \bigcup\{A'_i \mid i \leq n\}) > 0$, or for $\leq$-constraints, i.e. if $a_i = 0$. The problem with the former is that the $R(\neg A'_i)$ are then no longer guaranteed to be stable, which complicates the construction process. Fortunately, there exists an easy solution. The idea is to embed $\mathcal{W}$ into a larger universe $\mathcal{W}'$ and to introduce

- a propositional algebra $\mathcal{B}'$ over $\mathcal{W}'$ with $\mathcal{W} \in \mathcal{B}'$,
- a $\kappa\pi$-measure $R' : \mathcal{B}' \to [0, \infty]$ with $R = R'(.|\mathcal{W})$ and $R(\neg\mathcal{W}) = 0$.

For every constraint base Δ and proposition S, let Δ^S be obtained by relativizing Δ to S, that is by replacing $r(A) + a \leq r(A')$ with $r(A \wedge S) + a \leq r(A' \wedge S)$, or by relativizing the corresponding conditional ranking constraints to S. We may now proceed as follows. First, we observe that $R'(\mathcal{W}' - \bigcup\{A'_i \wedge \mathcal{W} \mid i \leq n\}) = 0$. Consequently, we may apply the standard relativized JZ-procedure to obtain $R'[\Delta^{\mathcal{W}}]$. A very reasonable choice for $R[\Delta]$ now seems to be $R'[\Delta^{\mathcal{W}}](.|\mathcal{W})$. In other words, we may always reformulate our revision problem so as to avoid the first difficulty. Equivalently, we could apply the algorithm to non-normalized pseudo-$\kappa\pi$-measures $(R(\mathcal{W}) > 0)$, and only normalize the result by shifting $\mathcal{W}$ downwards.

The problem with constraints of the form $r(A) \leq r(A')$ is that they may form loops, making the stratification for the inductive process more difficult to realize. Extending the JZ-procedure accordingly is much less straightforward. What can go wrong may be illustrated by the following simple example, with logically independent A, A', B.

- $\Delta_1 = \{1 \leq r(A), 1 \leq r(A'), r(A \wedge A') \leq r(B), r(B) \leq r(A')\}$.

If we restrict ourselves to upwards shifting and start from the uniform prior R_0, the only justifiably constructible model of Δ_1 is $R_0[A+1][A'+\infty][B+\infty]$. But this solution is inadequate because it violates a natural principle, which was proposed in a probabilistic context by Paris [94].

Open-mindedness. $NM(R, \Delta)(A) \neq \infty \ \to \ R[\Delta](A) \neq \infty$.

Because $NM(R_0, \Delta_1) = R_0[(A \vee A' \vee B) + 1]$, we should have $R_0[\Delta_1](A') \neq \infty$ and $R_0[\Delta_1](B) \neq \infty$. But this model is not even epistemically constructible over Δ_1. A natural candidate for $R_0[\Delta_1]$ would be $R_0[A+1][A'+1][B+1][A \wedge A' - 1]$. Here, A, A', B get the minimal possible values, as in the normality maximization model, whereas $A \wedge A'$ is shifted downwards to avoid making A' or B impossible. In fact, downwards shifting should only be permitted to prevent entangled situations risking to induce impossibility. It would not be justified in the next example.

- $\Delta_2 = \{1 \leq r(A), 1 \leq r(A'), r(A \wedge A') \leq r(B)\}$.

Here the appropriate model $R[\Delta_2]$ seems to be $R_0[A+1][A'+1][B+2]$. There is no need for a costly downwards shift. The distinction between these two examples is at the heart of the full JZBR-algorithm, which we are now going to present.

JZBR-Algorithm.
The extended strategy proceeds just as the original approach, with two exceptions. First of all, it determines a more fine-grained I_j-hierarchy, to reflect the $\leq$-relationships. Secondly, at each level, the standard minimal upwards shifting process is followed by a corresponding minimal downwards shifting process, for those $\leq$-constraints which have not yet been satisfied. The technical goal is that later shiftings should not invalidate constraints satisfied at earlier stages.

At step $j+1$, the definition of I_{j+1} is different. Let Φ_{j+1} be the set of all those unconsidered constraints with $NM(R^j, \Delta)(A'_i) = \alpha_{j+1}$. Within Φ_{j+1}, we want to single out those conditions which impose the weakest shifting tasks for their A'_i, and which are also the least dependent on the upwards shifting of other A'_i. The idea is to replace or strengthen as many constraints $r(A_i) + a_i \leq r(A'_i)$ in Φ_{j+1} as possible by $r(A_i) + a_i + \varepsilon \leq r(A'_i)$ $(\varepsilon > 0)$, while still having constraints δ_i with $NM(R^j, \Delta \cup \Phi'_{j+1})(A'_i) = \alpha_{j+1}$. Let Φ'_{j+1} be any such maximal transformation of Φ_{j+1}. Then, $\Phi_{j+1} \cap \Phi'_{j+1}$ is a minimal set of unaffected constraints. Let Ψ_{j+1} denote the union of all these minimal $\Phi_{j+1} \cap \Phi'_{j+1}$, whose elements may be interpreted as the least dependent or most stable ones. That's why we use Ψ_{j+1} to determine our new top-level constraints.

- $I_{j+1} = \{i \in I \mid \delta_i \in \Psi_{j+1}\}$.

Although this compact definition may look a bit technical, it still seems to represent the most natural prioritization for $\leq$-constraints, in line with the philosophy of system Z.

As in the original JZ-procedure, the corresponding A'_i are now uniformly and minimally shifted to level α_{j+1}. However, because the $\leq$-constraints may form entangled loops, as in above example, this process may also change the values of some A_i, thereby undermining the satisfaction of $r(A_i) \leq r(A'_i)$ at this stage. To repair this distortion, in a second phase, some A_i are shifted downwards, or backwards, to α_{j+1}. This is realized in the same uniform and minimal way as for the A'_i, only the direction changes. That is, we now consider negative p_i becoming larger and trying to reach α_{j+1} from above.

Theorem 6.1 (JZBR-Revision)
If R is coherent with Δ, the JZBR-procedure produces a unique model R of Δ. We write $R[\Delta] = JZ(R, \Delta)$.

We may illustrate the procedure with our previous examples. Let us denote the ith constraint by δ_i. For both examples, we have $\Phi_1 = \Delta$ and $\alpha_1 = 1$, 1 being the minimal possible value for the A'_i.

In the first example, there is only a single Φ'_1, $\Phi_1 \cap \Phi'_1 = \Psi_1 = \{\delta_1\}$, and A is shifted. Then $\Phi_2 = \{\delta_2, \delta_3, \delta_4\}$ and $\alpha_2 = 1$. Given that all these constraints are entangled, we get $\Phi_2 = \Phi'_2 = \Psi_2$ and have to shift A' and B to level 1. Because this pushes $A \wedge A'$ to level 2, the third constraint is violated and we have to repair it by shifting downwards $A \wedge A'$ by one unit. This gives us $R_0[A+1][A'+1][B+1][A \wedge A' - 1]$.

In the second example, we have $\Phi_1 \cap \Phi'_1 = \{\delta_1\}$ or $\Phi_1 \cap \Phi'_1 = \{\delta_2\}$, i.e. $\Psi_1 = \{\delta_1, \delta_2\}$. So, A, A' are shifted to level 1. Then, $\Phi_2 = \Psi_2 = \{\delta_3\}$ and B is shifted to $\alpha_2 = 2$. The result is $R_0[A+1][A'+1][B+2]$.

7 COMPARISONS AND CONCLUSIONS

JZBR interprets revision in the narrow sense, implicitly making a difference between the evaluation of new information and the realization of the resulting epistemic constraint by passing to a new epistemic measure. It is a rather singular approach, which differs from more traditional revision formalisms in several ways.

First of all, to compare JZBR with conventional proposition-centered accounts, we have to adopt a suitable translation policy mapping each incoming proposition to a $\kappa\pi$-measure constraint. The simplest solution is to associate with every new proposition A - as long as it is not declared impossible by the initial state R ($R(A) = \infty$), and thereby ignored - the constraint $r(\neg A) + 1 \leq r(A)$.

Secondly, it allows us to model multiple or parallel revision. In practice, evidence usually comes in packages of several items. The ad hoc strategy of imposing a more or less arbitrary order on the individual pieces of evidence is usually inappropriate because most powerful revision techniques, like Spohn's, are highly order-dependent. So, we need more direct approaches. JZBR does this by exploiting the top-level structure of constraint bases resulting from the evidential inputs. Whereas individual constraints are characterized by their semantic content, i.e. the set of $\kappa\pi$-measures satisfying them, this is no longer true for arbitrary constraint sets Δ. This may be seen as a feature insofar as it allows us to exploit implicit independencies (cf. corresponding discussion for default reasoning [Weydert 98]).

Thirdly, because JZBR not only modifies the set of beliefs but also the background preferences - to guide future revision processes or decision-theoretic considerations - it is difficult to grasp its behaviour by postulates mainly directed at the belief level. Concerning Pearl's and Darwiche's [97] postulates C1 - C6 , the situation is as follows. C1 and C2 do not work for complex belief states, like $\kappa\pi$-measures R. C3 and C4 hold without restrictions because they only make statements about individual beliefs. C5 and C6 always fail because they conflict with the Jeffrey conditionalization procedure.

Fourthly, the closest relative of JZBR may be MCE. In fact, JZBR can be seen as a semi-qualitative implementation of the information minimization paradigm. However, as we have already pointed out in the introduction, it should not to be confounded with $\kappa\pi$-MCE, the order-of-magnitude counterpart of MCE. JZBR is more robust because it is not dependent on the structure of the basic propositional algebra $\mathcal{B}$. This robustness also explains why MCE's problem with conditional probability constraints does not affect JZBR. Nevertheless, the Shore/Johnson axioms for discrete MCE [80] are also valid for JZBR, more precisely, for its extension to a slightly bigger class of constraints. Unfortunately, because the $\kappa\pi$-measure framework is much coarser than the probabilistic one, and semantic invariance is only partly ensured, this axiomatization does not characterize JZBR. Another interesting common property is open-mindedness.

The investigation of JZBR is still at its beginnings. Possibly, a deeper understanding of the multiple revision perspective may allow us to find postulates characterizing JZBR or variations thereof

REFERENCES

Boutilier 93 C. Boutilier. Revision sequences and nested conditionals. In *Proceedings of IJCAI 93*, 1993.

Darwiche and Pearl 97 A. Darwiche and J. Pearl. On the logic of iterated belief revision. In *Artificial Intelligence*, 89 : 1-29, 1997.

Gärdenfors 88 P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* Bradford Books, MIT-Press, Cambridge, Mass, 1988.

Goldszmidt and Pearl 96 M. Goldszmidt, J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. In *Artificial Intelligence*, 84 : 57-112, 1996.

Lehmann 95 D. Lehmann. Belief revision, revised. In *Proceedings of IJCAI 93.* Montreal, Canada 1995.

Paris and Vencovska 90 J.B. Paris and A. Vencovska. A note on the inevitability of maximum entropy.In *International Journal of Approximate Reasoning*, 4:183-223, 1990.

Paris 94 J.B. Paris. *The Uncertain Reasoner's Companion.* Cambridge University Press, 1994.

Shore and Johnson 80 J.E. Shore and R.W. Johnson. Axiomatic derivation of the principle of cross-entropy minimization. In *IEEE Transactions on Information Theory*, IT-26(1):26-37, 1980.

Spohn 90 W. Spohn. A general non-probabilistic theory of inductive reasoning. In *Uncertainty in Artificial Intelligence 4*, North-Holland, Amsterdam 1990.

Weydert 90 E. Weydert. Doxastic Preference Logic. A New Look at Belief Revision. In *Proceedings of JELIA 90*, ed. J.van Eijck. Springer, 1990.

Weydert 95 E. Weydert. Defaults and infinitesimals. Defeasible inference by non-archimdean entropy maximization. In *Proceedings of UAI 95.* Morgan Kaufmann, 1995.

Weydert 96 E. Weydert. System J - Revision entailment. In *Proceedings of FAPR 96*, Springer-Verlag, 1996.

Weydert 98 E. Weydert. System JZ - How to build a canonical ranking model of a default knowledge base. In *Proceedings of KR 98.* Morgan Kaufmann, 1998.

Williams 94 M. Williams. Transmutations of knowledge systems. In *Proceedings of KR94.* Morgan Kaufmann, 1994.