Non-prioritised ranked belief change

Thomas Meyer
Decision Systems Laboratory
Department of Information Systems
University of Wollongong
Wollongong, NSW 2522, Australia
e-mail: tmeyer@uow.edu.au

Aditya Ghose
Decision Systems Laboratory
Department of Information Systems
University of Wollongong
Wollongong, NSW 2522, Australia
e-mail: aditya@uow.edu.au

Samir Chopra
Knowledge Systems Group
School of Computer Science and Engineering
University of New South Wales
Sydney, NSW 2052, Australia
e-mail: schopra@cse.unsw.edu.au

Abstract

Traditional accounts of belief change have been criticised for placing undue emphasis on the new belief provided as input. A recent proposal to address such issues is the framework of Ghose and Goebel for non-prioritised belief change based on default theories [8]. A novel feature of this approach is the introduction of *disbeliefs* alongside beliefs which allows for a view of belief contraction as independently useful, instead of just being seen as an intermediate step in the process of belief revision. The Ghose-Goebel approach is, however, restrictive in assuming a linear ordering of reliability on the received inputs. In this paper, we extend the Ghose-Goebel framework by replacing the linear ordering with a *preference ranking* on inputs from which a total preorder on inputs can be induced. This extension brings along with it the problem of dealing with inputs of equal rank, a situation that cannot occur in the Ghose-Goebel setup. We provide a semantic solution to this problem which contains, as a special case, AGM belief change on closed theories.

1 Introduction

The AGM theory of belief change [1] focuses on the process through which beliefs are added to, or removed from, belief states (represented as closed theories) while maintaining consistency. Subsequent studies have considered belief states represented as *belief bases*, which are sets of sentences not necessarily closed under logical entailment [7, 11, 21].

Non-prioritised belief change refers to the class of belief change operations that do not come with a guarantee of success. Thus a non-prioritised belief revision operation does not guarantee the entailment of the new belief from the resulting belief state and a non-prioritised belief contraction operation does not guarantee the non-entailment of the retracted belief from the resulting belief state. Studies of non-prioritised belief change are motivated by the observation that traditional accounts of belief change place undue emphasis on the new belief provided as input [5, 10]. A recent proposal to address such concerns is the framework for non-prioritised belief change of Ghose and Goebel [8]. In addition to generalising classical base change this approach also allows for prior revision and contraction steps to be recorded and brought to bear in deciding the outcome of subsequent belief change steps. As in, for example, [3, 4], belief change operations are trivial, with new inputs simply being added to the current belief state. Instead, it is the entailment relation by which information is extracted from the current belief state which is sophisticated. This is in line with the principles of *lazy evaluation* and *deferred commitment* [8]. It also ensures that *iterated belief change* is dealt with appropriately. The novel feature of the Ghose-Goebel approach is the introduction of explicit *disbeliefs*, on par with explicit beliefs, as allowable inputs. This allows for a view of belief contraction as independently useful, instead of merely acting as an intermediate step, via the Levi identity [15], in a belief revision operation.

The main drawback of the Ghose-Goebel framework is its unrealistic assumption of a linear ordering of reliability on inputs which is unlikely to be realised in many practical situations. For example, the assumption of linearity renders their system unable to deal with any situation in which two or more inputs are regarded as equally reliable. In this paper we extend the Ghose-Goebel framework by replacing the linear orderings on inputs with *preference rankings* on inputs. These rankings have more structure than linear orderings since they permit an expression of the *strength of preferences*. They are also less restrictive than linear orderings since they induce *total preorders* on inputs. The freedom brought about by the move from linear orderings to rankings brings along with it the problem of dealing in an adequate fashion with inputs of equal rank, a situation that cannot occur in the Ghose-Goebel setup. We provide a semantic solution to this problem by replacing all inputs with logically weakened versions, from which an appropriate entailment relation is defined. Our account of non-prioritised ranked belief change satisfies desirable properties. An interesting feature of our approach is that it contains, as a special case, AGM belief change on closed theories.

1.1 Formal preliminaries

We assume a finitely generated propositional language L closed under the usual propositional connectives and with a classical model-theoretic semantics. Classical entailment is denoted by $\models$ and U is the set of interpretations of L. Logical equivalence is denoted by $\equiv$. For $X \subseteq L$, $M(X)$ is the set of models of X. For $\alpha \in L$ we write $M(\alpha)$ instead of $M(\{\alpha\})$. A *belief set* is a set of sentences closed under $\models$. For any total preorder $\preceq$ (i.e. reflexive transitive relation) on U and $\alpha \in L$ we let $M_{\preceq}(\alpha) = \{u \in M(\alpha) \mid u \preceq v \ \forall v \in M(\alpha)\}$. Thus $M_{\preceq}(\alpha)$ are the $\preceq$-minimal models of α. Examples involving interpretations are phrased in the language with two atoms, p and q. Interpretations are represented as sequences of 0s and 1s representing falsity and truth respectively. The convention is that the first digit represents the truth value of p and the second one the truth value of q. We use $X \setminus Y$ to represent the difference of two sets X, Y. The set of natural numbers is denoted by $\mathbb{N}$.

2 Non-prioritised belief change

This section is our account of the Ghose-Goebel approach to non-prioritised belief change described in [8]. We emphasise the role played by the linear orderings on inputs, an issue that is glossed over by Ghose and Goebel. This leads us to a formal description of the logic of beliefs and disbeliefs which is inherent in their original description. Although Ghose and Goebel base their proposal for non-prioritised belief change on default theories [20, 18] and use the THEORIST system [19] as the underlying default reasoning formalism, our account of their approach mentions neither the THEORIST system, not does it refer to default reasoning.

The *information* available to an agent is represented in the form of *beliefs* and *disbeliefs*. Beliefs are *positive* inputs which an agent will, under suitable conditions, commit to for purposes such as query-answering and acting. Disbeliefs are *negative* inputs which an agent will, under suitable conditions, *refuse* to commit to. The novel feature of this approach is the introduction of explicit disbeliefs which are used to record and store sentences that an agent refuses to commit to. We shall see below that this allows for a level of expression that cannot be obtained with sets just containing beliefs. Intuitively, the assertion of a disbelief is analogous to belief contraction. In fact, we shall see that base contraction can be recovered from the Ghose-Goebel setup.

Both beliefs and disbeliefs are represented as sentences of L. We use overlining to distinguish disbeliefs from beliefs. Thus p is a belief, while $\overline{p}$ is a disbelief. We usually use α and β to denote beliefs, $\overline{\gamma}$ and $\overline{\delta}$ to denote disbeliefs, and ϕ and ψ to refer to pieces of information (which can be either beliefs or disbeliefs).

Definition 2.1 An *information state* I is a finite set of information, i.e. beliefs and disbeliefs. The *belief state* consisting of all the beliefs in I is denoted by B_I. Similarly, the *disbelief state* consisting of all the disbeliefs in I is denoted by D_I. ∎

The notion of inconsistency for information states is crucial to the development of the rest of the account. For beliefs, consistency coincides with classical consistency. For disbeliefs, the situation is somewhat different.

Definition 2.2 A disbelief $\overline{\delta}$ is *consistent with* a belief state B iff $B \cup \{\neg\delta\} \nvDash \bot$. An information state I is *consistent* iff $B_I \nvDash \bot$ and every $\overline{\delta} \in D_I$ is consistent with B_I. ∎

Observe that for an information state I to be consistent it has to be the case that the belief state B_I is consistent in the classicial sense. Observe also that there is no notion of inconsistency between disbeliefs. Thus, it is perfectly possible for a disbelief state to contain the disbeliefs $\overline{\alpha}$ and $\overline{\neg\alpha}$. It simply means that the agent refuses to commit itself to either α or $\neg\alpha$. It is possible, however, for an information state to be inconsistent even if the belief state is consistent. For example the information state $I = \{p, \overline{p}\}$ is inconsistent.

The explicit expression of beliefs and disbeliefs is reminiscent of the handling of knowledge and belief in a modal logic setup, specifically in S5-models [12]. The statement $\Box\alpha$ in an S5-model corresponds to the assertion, in our setup, that α is believed. Similarly, the statement $\Diamond\neg\gamma$ in an S5-model corresponds to the assertion $\overline{\gamma}$ in our setup; that γ is disbelieved. It can easily be verified that the notions of consistency in these two frameworks coincide.

The challenge, then, is to obtain acceptable methods of extracting consistent information from possibly inconsistent information states. Ghose and Goebel's solution to this problem is to consider the *maximally consistent subsets* of an information state I as the foundation for defining appropriate entailment relations for I (where J is a maximally consistent subset of I iff (1) $J \subseteq I$, (2) J is consistent, and (3) for every K such that $J \subset K \subseteq I$, K is inconsistent). These maximally consistent subsets can be thought of as being generated by linear orders of reliability on I, with information higher up in the ordering being more reliable. Suppose, for example, that $I = \{p, p \rightarrow q, \overline{q}\}$. It should be clear that I is inconsistent. Now let us suppose that p is more reliable than $\overline{q}$ which, in turn, is more reliable than $p \rightarrow q$. The maximally consistent subset generated by this reliability ordering is $\{p, \overline{q}\}$.

Definition 2.3 A maximally consistent subset J of I is *<-generated* by a linear order $<$ on I iff

1. every $\overline{\delta} \in D_I \setminus J$ is inconsistent with $\{\beta \in B_J \mid \beta > \overline{\delta}\}$, and
2. $\forall \alpha \in B_I \setminus J$, either $\{\beta \in B_J \mid \beta > \alpha\} \cup \{\alpha\}$ is inconsistent or there is some $\overline{\delta} \in D_J$ such that $\overline{\delta} > \alpha$ and $\overline{\delta}$ is inconsistent with $\{\beta \in B_J \mid \beta > \alpha\} \cup \{\alpha\}$.

■

We denote by $I^<$ the $<$-generated maximally consistent subset of the information state I. Obtaining $I^<$ is simply a matter of considering the elements of I one by one in the order prescribed by $<$ and adding the element iff its inclusion ensures that the set built up thus far is consistent.

We are now in a position to define what it means for information to be entailed by an information state.

Definition 2.4 Let I be an information state I and $<$ a linear order on I. A belief α is *<-entailed* by I, denoted by $I \models^< \alpha$, iff $B_{I^<} \models \alpha$. A disbelief $\overline{\delta}$ is *<-entailed* by I, denoted by $I \models^< \overline{\delta}$ iff either $\delta \equiv \bot$ or there is a $\overline{\gamma} \in D_{I^<}$ such that $\delta \models \gamma$. An agent with I as its information state and $<$ as its linear order of reliability is said to be *committed* to the information $<$-entailed by I. ■

It is perhaps necessary to expand somewhat on the notion of $<$-entailment for disbeliefs. Disbeliefs are, in a sense, dual to beliefs. Committing to a belief α means that we also commit to any belief that is logically weaker than α, while committing to a disbelief $\overline{\delta}$ means that we are also committing to every disbelief $\overline{\gamma}$ where γ is logically *stronger* than δ. As extreme cases, an agent is always committed to believing tautologies and disbelieving contradictons. Weakening a disbelief $\overline{\delta}$, as we shall do in section 4, therefore means that we replace $\overline{\delta}$ with a disbelief $\overline{\gamma}$ such that $\gamma \models \delta$.

Observe also that committing to a disbelief $\overline{\delta}$ is quite different from committing to the belief $\neg\delta$. A commitment to the belief $\neg\delta$ does not *imply* a commitment to the disbelief $\overline{\delta}$, and neither does the converse hold. The following example serves as an illustration of the use of $<$-entailment.

Example 2.1 Let $I = \{p, q, \overline{q}\}$ and let $<$ be such that $q < \overline{q} < p$. Then $I^< = \{p, \overline{q}\}$ and therefore $I \models^< p$, $I \models^< p \vee q$, $I \models^< \overline{q}$ and $I \models^< \overline{p \wedge q}$. ■

Much of classical base change [7, 11, 21] is based on maximal consistency. The result of contracting a base B with a sentence α is taken to be any maximally consistent subset of B which does not entail α. The revision of B by α is defined, via the Levi identity [15], as the result of contracting B with $\neg\alpha$ and then adding α. It is easily verified that base change, as described above, can be recovered from the Ghose-Goebel setup. The revision $*$ of a belief base B by α is obtained by regarding $B' = B \cup \{\alpha\}$ as an information state, imposing on it any linear order $<$ such that $\beta < \alpha$

for every $\beta \in B' \setminus \{\alpha\}$ and setting $B * \alpha = (B')^<$. Similarly, the contraction $-$ of B by δ is obtained by viewing $B' = B \cup \{\overline{\delta}\}$ as an information state, imposing on it any linear order $<$ such that $\beta < \overline{\delta}$ for every $\beta \in B' \setminus \{\overline{\delta}\}$ and setting $B - \delta = (B')^< \setminus \{\overline{\delta}\}$.

The introduction of explicit disbeliefs ensures that belief contraction is placed on an equal footing with revision instead of being seen as an intermediate step in the process of performing revision. Ghose and Goebel provide the following example to illustrate the point.

Example 2.2 Consider a situation in which a totally ignorant agent first revises by $p \rightarrow q$, then contracts with q and finally revises by p. In a framework which only caters for the explicit representation of beliefs there is only one feasible result for this sequence of base change operations. Let $B_0 = \emptyset$ represent the initial beliefs of the agent. We then have $B_1 = B_0 * p \rightarrow q = \{p \rightarrow q\}$, $B_2 = B_1 - q = \{p \rightarrow q\}$, and $B_3 = B * p = \{p, p \rightarrow q\}$. While this ought to be *one* of the permissible outcomes of such a sequence of steps, the argument is that some of the relevant inputs, specifically the contraction of q, have not been taken into account. Contrast this with the framework described above which proceeds as follows: $I_0 = \emptyset$, $I_1 = I_0 * p \rightarrow q = \{p \rightarrow q\}$, $I_2 = I_1 - q = \{p \rightarrow q, \overline{q}\}$, and $I_3 = I_2 * p = \{p, p \rightarrow q, \overline{q}\}$. $<$-entailment for I_3 would then be based on one of the maximally consistent subsets $\{p, p \rightarrow q\}$, $\{p, \overline{q}\}$, or $\{p \rightarrow q, \overline{q}\}$, depending on the linear ordering imposed on I_3. ∎

In contrast with traditional approaches to belief change, *iterated belief change* has a solution built into this setup, as can be seen from the previous example. Sequential inputs, in the form of beliefs and disbeliefs, are simply added to the current information state as they are made available to the agent. The problems associated with iterated belief change are now transformed into the question of how to extract consistent information from the updated information state. This question, in turn, is deferred until it needs to be resolved, at which time an appropriate linear ordering is determined from the context and other potential relevant issues, such as the order in which inputs were received, and $<$-entailment is applied.

3 Ranked information states

While $<$-entailment provides an interesting deployment of non-prioritised belief change its applicability is severely restricted by the insistence that a linear order of reliability on the elements of an information state be supplied. We propose, instead, to impose reliability by *ranking* the elements of information states in the spirit of Spohn [22]. These rankings take on the form of an assignment of natural numbers to the elements of information states; the *higher* the rank of a piece of information, the more reliable it is deemed to be.

Definition 3.1 A *ranked belief* is a pair (α, r) where $\alpha \in L$ and $r \in \mathbb{N}$. A *ranked disbelief* is a pair $(\overline{\delta}, r)$ where $\delta \in L$ and $r \in \mathbb{N}$. A *ranked information state RI* is a finite *multiset* of ranked information (i.e. ranked beliefs and ranked disbeliefs). The *ranked belief state* consisting of all the ranked beliefs in *RI* is denoted by B_{RI}. Similarly, the *ranked disbelief state* consisting of all the disbeliefs in *RI* is denoted by D_{RI}. ∎

Since a ranked information state may assign the same rank to different pieces of information, it is possible for different bits of information to be *equally reliable*, a much weaker and more realistic restriction than the Ghose-Goebel insistence on bits of information being linearly ordered in terms of reliability. As an extreme case, it allows for the possibility of regarding *all* the information contained in an information state as equally reliable. A ranked information state is also more expressive than an information state equipped with a linear ordering. The former can be used to indicate reliability *strength*, something that the latter is unable to do.

4 Dealing with equal reliability

The possibility of information in a ranked information state being equally reliable introduces a problem which does not crop up in the Ghose-Goebel setup; what to do with mutually inconsistent pieces of information that are equally reliable. This section is devoted to finding a solution to this problem. We shall, in this section, be concerned with multisets of information (not ranked information) and shall refer to these as *information multisets*. The multiset of beliefs contained in an information multiset I is denoted by B_I. Similarly, the multiset of disbeliefs contained in I is denoted by D_I. In section 5 the results obtained in this section will be applied to ranked information states.

One way of dealing with inconsistent *beliefs* of equal reliability is to employ the *infobases* of Meyer et al. [16, 17]. An infobase is a finite multiset of sentences.[1] Intuitively an infobase is a structured representation of the beliefs of an agent with a foundational flavour. It is assumed that every belief in an infobase is obtained independently. Meyer exploits the structure of infobases to generate total preorders on U with an interpretation *lower* down being regarded as *more* plausible.

Definition 4.1 For $u \in U$, the *B-number* of u, denoted by $num_B(u)$, is the number of beliefs α in an infobase B such that $u \in M(\alpha)$. The total preorder $\preceq_B$ on U *generated* by num_B is defined as follows: $u \preceq_B v$ iff $num_B(v) \leq num_B(u)$. ∎

Since every sentence in an infobase B is seen as independently obtained it is reasonable to regard as more plausible those interpretations which are models of more of the sentences in B. It should now also be clear why infobases need to be multisets instead of sets. Multiple copies of the same belief in an infobase will lead to the generation of a total preorder that differs from the one generated by an infobase with just a single copy of the belief, as demonstrated in the example below.

Example 4.1 For the infobase $B = \{p, q\}$ we have $num_B(11) = 2, num_B(10) = num_B(01) = 1$, and $num_B(00) = 0$. For $B' = \{p, p, q\}$ we have $num_B(11) = 3$, $num_B(10) = 2$, $num_B(01) = 1$, and $num_B(00) = 0$. ∎

For our purposes the most important feature of an infobase is that it allows for the extraction of a consistent knowledge base from *any* infobase B, even if beliefs occurring in it are inconsistent. The idea is simple; take the minimal elements of U, with respect to $\preceq_B$, to be the models of a sentence $K(\preceq_B)$ where, in the spirit of [13], $K(\preceq_B)$ is regarded as a finite representation of its own logical closure.

[1] Actually, Meyer defines an infobase as a list, or sequence, of sentences. However, the reason for this is to allow for multiple copies of the same sentence; the order of the sentences in an infobase is irrelevant and a definition in terms of multisets is therefore also possible.

Definition 4.2 The *knowledge base extracted* from a total preorder $\preceq$ on U is some sentence $K(\preceq)$ s.t. $M(K(\preceq)) = M_{\preceq}(\top)$. ■

The intuition is in line with the typical *minimal-model* semantics approach in the nonmonotonic reasoning and belief revision literature [9, 14, 13].

Example 4.2 Consider the infobase $B = \{p, q, \neg p\}$. It is easily verified that $num_B(11) = 2$, $num_B(10) = 1$, $num_B(01) = 2$, $num_B(00) = 1$, and thus that $M_{\preceq_B}(\top) = \{11, 01\}$. In other words, the knowledge base extracted from $\preceq_B$ is logically equivalent to q. ■

As would be expected, knowledge extraction from infobases is compatible with consistency.

Proposition 4.1 *Suppose that the infobase B is consistent. Then* $K(\preceq_B) \equiv \bigwedge_{\alpha \in B} \alpha$.

The use of infobases supplies us with the formal machinery to extract consistency from a possibly inconsistent multiset of beliefs. The next step is to expand this to information multisets and to bring *disbeliefs* into the picture as well.[2] The definition for consistency of information multisets is an obvious extension of definition 2.2.

Definition 4.3 Let B be a multiset of beliefs. A disbelief $\overline{\delta}$ is *consistent* with $K(\preceq_B)$ iff $K(\preceq_B) \wedge \neg\delta \nvDash \bot$. An information multiset I is *consistent* iff every disbelief in D_I is consistent with $K(\preceq_B)$. ■

To resolve the inconsistency of an information multiset I in a fair and equitable manner we shall effect a weakening of all the elements of I to ensure consistency. In the first step for doing so $\preceq_{B_I}$ is modified so that the minimal models, with respect to $\preceq_{B_I}$, of $\neg\delta$ for every disbelief $\overline{\delta} \in D_I$ are taken to be as plausible as the models of $K(\preceq_{B_I})$. This ensures the consistency of each of the elements of D_I with the knowledge base extracted from the modified total preorder. This process is reminiscent of, and indeed was inspired by, the semantic description of AGM theory contraction where a contraction with the sentence α results in the addition of the minimal models of $\neg\alpha$, defined relative to some appropriate total preorder on U, to the models of the current belief set.

Definition 4.4 Let I be an information multiset. The *I-number* of an interpretation u is defined as follows:

$$num_I(u) = \begin{cases} \max\{num_{B_I}(v) \mid v \in U\} \text{ if} \\ \quad u \in M_{\preceq_{B_I}}(\neg\delta) \text{ for some } \overline{\delta} \in D_I, \\ num_{B_I}(u) \text{ otherwise.} \end{cases}$$

The total preorder $\preceq_I$ on U generated by num_I is defined as: $u \preceq_I v$ iff $num_I(v) \leq num_I(u)$. The knowledge base extracted from $\preceq_I$ is denoted by $K(\preceq_I)$. ■

So, while $K(\preceq_{B_I})$ is a weakening of the elements of B_I to ensure the consistency of B_I, $K(\preceq_I)$ amounts to a further weakening to ensure the consistency of I.

[2]Actually, the use of *sets* of disbeliefs poses no problem, but since it is necessary to collect beliefs in multisets it simplifies matters to do the same with disbeliefs.

Example 4.3 For the information multiset $I = \{p, q, \overline{p}, \overline{q}\}$ it can be verified that $num_{B_I}(11) = 2$, $num_{B_I}(10) = num_{B_I}(01) = 1$, and $num_{B_I}(00) = 0$. Therefore $M(K(\preceq_{B_I})) = \{11\}$, $M_{\preceq_{B_I}}(\neg p) = \{01\}$, and $M_{\preceq_{B_I}}(\neg q) = \{10\}$. So $num_I(11) = num_I(10) = num_I(01) = 2$, $num_I(00) = 0$ and thus $K(\preceq_I)$ is a sentence for which $M(K(\preceq_I)) = \{11, 10, 01\}$. And since $M(p \vee q) = \{11, 01, 01\}$ we can set $K(\preceq_I) = p \vee q$. It is easily verified that $\{p \vee q, \overline{p}, \overline{q}\}$ is consistent. ■

It is no accident in the example above that $\{p \vee q, \overline{p}, \overline{q}\}$ is consistent.

Proposition 4.2 *For any information multiset I, $\{K(\preceq_I)\} \cup D_I$ is consistent.*

There is, however, a kind of asymmetry built into this construction. Observe that while B_I has been weakened considerably, first to $K(\preceq_{B_I})$ in order to obtain consistency *within* B_I, and then to $K(\preceq_I)$ to ensure consistency with D_I as well, there has been no corresponding weakening of D_I. The second step in resolving the inconsistency of I is thus to replace each of the elements of D_I with suitably weakened disbeliefs.

Definition 4.5 Given a knowledge base K, the *K-weakened version* $\overline{W^K(\delta)}$ of a disbelief $\overline{\delta}$ is a disbelief such that

$$W^K(\delta) = \begin{cases} \delta \text{ if } \overline{\delta} \text{ is consistent with } K, \\ \delta \wedge \neg K \text{ otherwise.} \end{cases}$$

■

The next example, which is a continuation of example 4.3, demonstrates the use of weakening.

Example 4.4 For $I = \{p, q, \overline{p}, \overline{q}\}$ we can set $K(\preceq_{B_I}) = p \wedge q$. Since both $\overline{p}$ and $\overline{q}$ are inconsistent with $K(\preceq_{B_I})$ we have that $W^{K(\preceq_{B_I})}(p) = p \wedge \neg(p \wedge q)$ and $W^{K(\preceq_{B_I})}(q) = q \wedge \neg(p \wedge q)$. Observe that $W^{K(\preceq_{B_I})}(p) \equiv p \wedge \neg q$ and $W^{K(\preceq_{B_I})}(q) \equiv \neg p \wedge q$. ■

In summary then, consistency can be extracted from an inconsistent information multiset I by replacing B_I with $K(\preceq_I)$ and replacing each disbelief in D_I with its weakened version. Of course, weakening should only occur when I is inconsistent.

Proposition 4.3 *If I is consistent, then $K(\preceq_{B_I}) \equiv \bigwedge_{\alpha \in B_I} \alpha$ and $W^{K(\preceq_{B_I})}(\delta) = \delta$ for every $\overline{\delta} \in I$.*

5 Non-prioritised ranked belief change

We are now in a position to describe our proposal for non-prioritised ranked belief change. Given a ranked information state RI, information of equal rank is collected into different information multisets, with each information multiset corresponding to a particular rank. From these, a consistent information multiset is constructed, using the formal tools developed in section 4, which respects the ranks assigned to information.

Consider a ranked information state RI. Let $\max = \max\{i \mid (\phi, i) \in RI\}$ and let $I^r = \{\phi \mid (\phi, r) \in RI\}$ for $r \in \{0, \ldots, \max\}$. So I^r is the multiset containing all information in RI with rank r. To aid with readability we refer to the multiset of beliefs contained in I^r as B^r and the multiset of disbeliefs contained in I^r as D^r. As a first step we view each B^r as an infobase and define a lexicographic refinement of the total preorders generated by these infobases. The higher the rank of an infobase, the higher the prominence of its generated total preorder in the lexicographic ordering. This is similar to the lexicographic orderings generated from stratified knowledge bases, such as in [2].

Definition 5.1 For $r \leq \max$, the total preorder $\preceq^r$ *generated* by RI is defined as follows: $u \preceq^r v$ iff $\forall i \in \{r, \ldots, \max\}$, $v \prec_{B^i} u$ implies $\exists j \in \{j+1, \ldots, \max\}$ s.t. $u \prec_{B^j} v$. The knowledge base *extracted* from $\preceq^r$ is denoted by $K(\preceq^r)$. ∎

If $B = \bigcup_{i=r}^{i \leq \max} B^i$ is inconsistent, then, analogous to the case for infobases, $K(\preceq^r)$ represents an appropriate way of extracting maximal consistency from the beliefs in B, but with the ranks of the beliefs taken into account in this case. If B is consistent we get the expected result.

Proposition 5.1 *If* $B = \bigcup_{i=r}^{i \leq \max} B^i$ *is consistent then* $\bigwedge_{\alpha \in B} \alpha \equiv K(\preceq^r)$.

Observe that $K(\preceq^0)$ is the knowledge base extracted when infobases of *all* ranks have been taken into account.

Example 5.1 Let $RI = \{(p, 2), (p \rightarrow q, 1), (\neg q, 0)\}$. Then $B^0 = \{\neg q\}$, $B^1 = \{p \rightarrow q\}$ and $B^2 = \{p\}$. It can be verified that $11 \preceq^0 10$, $10 \preceq^0 00$, $00 \preceq^0 01$ and thus that $K(\preceq^0) \equiv p \wedge q$. ∎

The next step is to bring ranked disbeliefs into the picture. Roughly, the idea is to retain intact those disbeliefs which are consistent with higher ranked beliefs, to discard those which are inconsistent with higher ranked beliefs, and to weaken those which are consistent with higher ranked beliefs but inconsistent with equally ranked beliefs.

Definition 5.2 For every $r \leq \max$, the multiset RD^r of disbeliefs to be *retained* from D^r is defined as follows: $\overline{\delta} \in RD^r$ iff $\overline{\delta} \in D^r$ and $\overline{\delta}$ is consistent with $K(\preceq^r)$. Furthermore, we let let $RD = \bigcup_{i=0}^{i \leq \max} RD^r$. ∎

Observe that $\overline{\delta}$ is consistent with $K(\preceq^r)$ iff it is consistent with $K(\preceq^i)$ for every $i \in \{r, \ldots, \max\}$.

Definition 5.3 For every $r \leq \max$, the multiset WD^r of disbeliefs in D^r to be *weakened* is defined as follows: $\overline{\delta} \in WD^r$ iff $\overline{\delta} \in D^r$ and $\overline{\delta}$ is consistent with $K(\preceq^i)$ for every $i \in \{r+1, \ldots, \max\}$ but inconsistent with $K(\preceq^r)$. Also, we let $WD = \bigcup_{i=0}^{i \leq \max} WD^r$. ∎

So RD contains all the disbeliefs to be retained and WD contains all the disbeliefs to be weakened.

As with the definition of $K(I)$ in definition 4.4, the aim is now, firstly, to weaken $K(\preceq^0)$ by adding to the models of $K(\preceq^0)$ the minimal models, with respect to $\preceq^r$, of the negation of every disbelief of every rank r which is to be retained and weakened.

Definition 5.4 The knowledge base *extracted* from the information state RI is a sentence $K(RI)$ such that

$$M(K(RI)) = M(K(\preceq^0)) \cup \bigcup_{r=0}^{r \leq \max} \left(\bigcup_{\overline{\delta} \in WD^r} M_{\preceq^r}(\neg\delta) \right) \cup \bigcup_{r=0}^{r \leq \max} \left(\bigcup_{\overline{\delta} \in RD^r} M_{\preceq^r}(\neg\delta) \right)$$

∎

Secondly, for every rank r, the disbeliefs in WD^r are weakened in the manner described in definition 4.5 and added to the disbeliefs to be retained.

Definition 5.5 The set of disbeliefs $D(RI)$ *associated* with RI is defined as

$$D = \bigcup_{r=0}^{r\leq\max} RD^r \cup \bigcup_{r=0}^{r\leq\max} \{\overline{W^{K(\preceq^r)}(\delta)} \mid \overline{\delta} \in WD^r\}.$$

■

This brings us to the definition of rank-entailment.

Definition 5.6 A belief α is *rank-entailed* by RI, denoted by $RI \Vdash \alpha$, iff $K(RI) \models \alpha$. A disbelief $\overline{\delta}$ is *rank-entailed* by RI, denoted by $RI \Vdash \overline{\delta}$, iff $\delta \equiv \bot$ or $\delta \models \gamma$ for some $\overline{\gamma} \in D$. An agent with RI as its ranked information state is *committed* to the information rank-entailed by RI. ■

The next example illustrates the use of rank-entailment.

Example 5.2 Let $RI = \{(p \rightarrow q, 2), (\overline{p}, 2), (p, 1), (\overline{p \wedge q}, 1), \overline{p \rightarrow q}, 1), (\neg q, 0)\}$. It can be verified that $K(\preceq^2) \equiv p \rightarrow q$, $K(\preceq^1) \equiv p \wedge q$ and $K(\preceq^0) \equiv p \wedge q$. Now, $\overline{p} \in D^2$ is consistent with $K(\preceq^2)$ and should therefore be retained, $\overline{p \wedge q} \in D^1$ is consistent with $K(\preceq^2)$ but inconsistent with $K(\preceq^1)$ and should therefore be weakened, and $\overline{p \rightarrow q} \in D^1$ is inconsistent with $K(\preceq^2)$ and should therefore be removed. Since $M_{\preceq^2}(\neg p) = M_{\preceq^1}(\neg(p \wedge q)) = \{01, 00\}$, it follows that $K(RI) \equiv (p \rightarrow q)$. Moreover, the disbeliefs associated with RI are $\overline{p}$ and $\overline{W^{K(\preceq^0)}(p \wedge q)}$ where $W^{K(\preceq^0)}(p \wedge q) \equiv \bot$. So, for example, we have that $RI \Vdash p \rightarrow q$, $RI \Vdash \overline{p}$, and $RI \Vdash \overline{p \wedge \neg q}$. ■

Since our framework is an extension of the Ghose-Goebel framework, many of its desirable properties, such as deferred commitment, lazy evaluation, and an adequate handling of iterated belief change, are carried over. Another interesting feature of rank-entailment is that AGM belief change on belief sets can be recovered from it. Informally, the idea is to simulate AGM revision by ensuring that the belief to be added has a higher rank than any of the elements in the current ranked information state. Similarly, AGM contraction is simulated by ensuring that the belief to be contracted, represented as a *disbelief* to be added, has a higher rank than any of the elements in the current ranked information state. Of course, this can be achieved in the fashion described above because the inputs added to a ranked information state are more informative than the inputs in the AGM style. The next result makes these ideas precise.

It is well-known [13] that the different AGM revision and contraction operations can be obtained from the total preorders on U.

Theorem 5.1 *Let K be a knowledge base.*

1. *Let $*$ be an AGM belief revision operation $*$, and let $\preceq$ be the total preorder $\preceq$ on U corresponding to $*$ and K. Now construct an infobase B such that $\preceq_B = \preceq$. (This is always possible. See [16].) Then $K * \alpha \models \beta$ iff β is rank-entailed by $\{(\rho, 0) \mid \rho \in B\} \cup \{(\alpha, 1)\}$.*
2. *Let $-$ be an AGM belief contraction operation $-$, and let $\preceq$ be the total preorder $\preceq$ corresponding to $-$ and K. Now construct an infobase B such that $\preceq_B = \preceq$. (Again, this is always possible. See [16].) Then $K - \alpha \models \beta$ iff β is rank-entailed by $\{(\rho, 0) \mid \rho \in B\} \cup \{(\overline{\alpha}, 1)\}$.*

6 Conclusion and future work

In this paper we have provided a different angle on the framework of Ghose and Goebel for non-prioritised belief change based on default theories described in [8]. We have retained the novel feature of their approach i.e., the introduction of *disbeliefs* alongside beliefs, which places belief contraction on an equal footing with belief revision. We also retain revision and contraction histories, but relax the linear ordering of reliability in the Ghose-Goebel framework to a preference ranking on inputs. Our model was shown to deal satisfactorily with inputs of equal rank by replacing inputs with logically weakened versions, from which an appropriate entailment relation is defined. Our account of non-prioritised ranked belief change was also shown to incorporate AGM belief change.

The similarities between the handling of knowledge and belief in modal logic and the explicit expression of beliefs and disbeliefs were pointed out in section 2. The connection breaks down, however, when the ranking, or ordering, of sentences is taken into account. An investigation into the incorporation of such rankings, or orderings, into the modal logic setup might prove to be useful, especially since it might lead to results concerning computational complexity. And in a similar vein, there seems to be a connection between our framework and possibilistic logic [6], in which necessity-valued statements may be seen as beliefs, and possibility-valued statements as disbeliefs.

References

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] Salem Benferhat, C. Cayrol, Didier Dubois, Jerome Lang, and Henri Prade. Inconsistency management and prioritized syntax-based entailment. In R. Bajcsy, editor, *IJCAI-93. Proceedings of the 13th International Joint Conference on Artificial Intelligence held in Chambery, France, August 28 to September 3, 1993*, pages 640–645, San Mateo, CA, 1993. Morgan Kaufmann.

[3] Gerhard Brewka. Belief revision in a framework for default reasoning. In Andre Fuhrmann and N. Morreau, editors, *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence, Number 465, pages 206–222. Springer, Berlin, 1991.

[4] Samir Chopra, Konstantinos Georgatos, and Rohit Parikh. Relevance sensitive nonmonotonic inference on belief sequences. In *Proceedings of NMR'2000*, 2000.

[5] C. B. Cross and R. H. Thomason. Conditionals and knowledge-base-update. In Peter Gärdenfors, editor, *Belief Revision*, volume 29 of *Cambridge Tracts in Theoretical Computer Science*, pages 247–275. Cambridge University Press, Cambridge, 1992.

[6] Didier Dubois, Jerome Lang, and Henri Prade. Possibilistic logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, pages 439–513. 1994.

[7] André Fuhrmann. Theory contraction through base contraction. *Journal of Philosophical Logic*, 20:175–203, 1991.

[8] Aditya Ghose and Randy Goebel. Belief states as default theories: Studies in non-prioritised belief change. In Henri Prade, editor, *ECAI 98. 13th European Conference on Artificial Intelligence*, pages 8–12, New York, 1998. Johm Wiley & Sons, Ltd.

[9] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

[10] Sven Hansson. Semi-revision. *Journal of Applied Non-Classical Logics*, 7(1–2):151–175, 1997.

[11] Sven-Ove Hansson. Changes of disjunctively closed bases. *Journal of Logic, Language and Information*, 2(4):255–284, 1993.

[12] G.E. Hughes and M.J. Cresswell. *An introduction to Modal Logic*. Methuen, 1972.

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

[14] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[15] Isaac Levi. Subjunctives, dispositions and changes. *Synthese*, 34:423–455, 1977.

[16] Thomas Meyer. Basic Infobase Change. In Norman Foo, editor, *Advanced Topics in Artificial Intelligence*, volume 1747 of *Lecture Notes In Artificial Intelligence*, pages 156–167, Berlin, 1999. Springer-Verlag.

[17] Thomas A. Meyer, Willem A. Labuschagne, and Johannes Heidema. Infobase Change: A First Approximation. *Journal of Logic, Language and Information*, 9(3):353–377, 2000.

[18] David Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

[19] David Poole, Randy Goebel, and Romas Aleliunas. Theorist: a logical reasoning system for defaults and diagnosis. In N. J. Cercone and G. McCalla, editors, *The Knowledge Frontier: Essays in the Representation of Knowledge*, pages 331–352. Springer Verlag, 1987.

[20] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[21] Hans Rott. *Change, Choice and Inference*. Oxford University Press, 1996.

[22] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision: Belief, Change and Statistics: Proceedings of the Irvine Conference on Probability and Causation: Volume II*, volume 42 of *The University of Western Ontario Series in Philosophy of Science*, pages 105–134, Dordrecht, 1988. Kluwer Academic Publishers.