Non-prioritized ranked belief change

Samir Chopra (schopra@sci.brooklyn.cuny.edu)
Department of Computer and Information Science
Brooklyn College of the City University of New York
Brooklyn, NY 11210, USA

Aditya Ghose (aditya@uow.edu.au)
Decision Systems Laboratory
Department of Information Systems
University of Wollongong
Wollongong, NSW 2522, Australia

Thomas Meyer (tmeyer@cs.up.ac.za)
Department of Computer Science
University of Pretoria
Pretoria 0002, South Africa

Abstract. Traditional accounts of belief change have been criticized for placing undue emphasis on the new belief provided as input. A recent proposal to address such issues is a framework for *non-prioritized* belief change based on default theories (Ghose and Goebel, 1998). 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. This approach is, however, restrictive in assuming a linear ordering of reliability on the received inputs. In this paper, we replace 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. We provide a semantic solution to this problem which contains, as a special case, AGM belief change on closed theories.

1. Introduction

An adequate formal modeling of the process of *belief change* must conform to *methodological principles* and *rationality constraints*. In this area of philosophical inquiry, the interplay of descriptive and prescriptive theorizing is particularly intense. This is due to the rich inter-disciplinary nature of work done therein, attracting the attention of mathematicians, logicians, philosophers and computer scientists. Perhaps even more fundamentally, this is because the issue of belief change is linked with studies of rationality, a field marked by its lack of agreement—as any sufficiently interesting philosophical field should be. In this paper, we present a *formal* model that aims to do justice to several intuitions about rational belief change. These intuitions are: the need for a model of belief revision to adequately address the issue of *judgement* of beliefs, the reduction of the problem of belief revision to the construction of an adequate *non-monotonic inference relation* and the problem of the resolution of *mutually inconsistent, yet equally reliable* inputs. By addressing these issues, we obtain as well, a treatment of *contraction*, which enables a view of it as an independent epistemic operation rather than as a merely derivative one.

The AGM theory of belief change (Alchourrón et al., 1985) focuses on the process through which beliefs are added to, or removed from, belief states—represented as closed theories—while maintaining consistency. Revision and contraction operations are axiomatized so as to guarantee success: after revision, the new input is included in the revised belief state, and after contraction the contracted input is no longer entailed by the contracted belief state. The explicit process of investigation and judgement of the beliefs is left aside. Subsequent studies have considered belief states represented as *belief bases*, which are finite sets of sentences not necessarily closed under logical entailment (Fuhrmann, 1991; Nebel, 1992; Hansson, 1993; Rott, 1996).

Non-prioritized belief change refers to the class of belief change operations that do not come with a guarantee of success. Thus a non-prioritized belief revision operation does not guarantee the entailment of the new belief from the resulting belief state and a non-prioritized belief contraction operation does not guarantee the non-entailment of the retracted belief from the resulting belief state. Studies of non-prioritized belief change are motivated by the observation that traditional accounts of belief change place undue emphasis on the new belief provided as input (Cross and Thomason, 1992; Hansson, 1997; Hansson, 1999).[1] Furthermore, implausibly enough, without any talk of epistemic entrenchments or the like, traditional models treat the standing of a belief as being rapidly in flux. Upon receipt, a new epistemic input is more important than any other member of the belief corpus (in order to restore consistency, anything other than tautologies may be given up). Upon receipt of another epistemic input, however, the older belief is now just as susceptible to removal as anything else in the corpus. This rapid changing of the status of the beliefs is not in accordance with any reasonable descriptive or prescriptive model of belief change.

A non-prioritized model is also more likely to do justice to a full-blown model of inquiry since in traditional models, the process of *investigating* the new information is essentially abstracted away. The formal operations simply concern themselves with the mechanics of constructing the revised belief states. But a more interesting model would pay attention to the process of investigating the new belief, weighing it up, and then carrying out the revision if warranted. In non-prioritized models of belief revision, a guiding principle is that some notion of *epistemic value* is at hand, which is used to adjudicate in the case of conflicting beliefs. (Hansson, 1997) describes this model as

[1] These models are also termed 'autonomous' (Galliers, 1992) since an agent is allowed to decide for itself whether it wants to revise its beliefs or not.

consisting first of expansion by the new input followed by the restoration of consistency by either dropping the new input or by giving up other beliefs that conflict with it. Since expansion by the new belief can result in an inconsistent belief state, the representation of the belief state in a non-prioritized model cannot be a closed theory—belief bases are a natural choice instead. Note too, that since epistemic inputs (in the form of sentences) are to be retained with the decision of which ones to ignore to obtain consistency postponed until necessary, it seems almost inevitable that belief states be represented as bases in this model.

Other than the arguments given above, epistemic scenarios not accomodated in traditional belief change models can now be handled. Consider the recalcitrant religious fundamentalist (Hansson, 1999), unwilling to give up anything that is contained in a suitable religious text. This scenario cannot be handled by a belief change model based on epistemic entrenchment, since the only maximally entrenched sentences are the tautologies. But this person is not all that fundamentalist if their beliefs can simply be changed by carrying out revision by the negation of a sentence that contradicts something in the religious text. However, using Hansson's consolidation operator, we can represent this easily: the operator just is one that rejects anything that contradicts the text.

A framework for non-prioritized belief change using belief bases presented in (Ghose and Goebel, 1998) generalizes classical base change and 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, (Brewka, 1991; Chopra et al., 2000), the *vertical* or *direct* mode of belief revision is employed. In the direct mode, 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; a form of *deferred epistemic commitment* as it were. This ensures that *iterated belief change* is dealt with appropriately: new epistemic inputs are simply added to the belief state and the problem of iterated revision reduces to the problem of determining which inputs are to be used—via a consideration of context, history and other parameters—in obtaining consistency to guide the associated inference relation. There are other desirable features. First, derived beliefs or implicit ones change *only* if changes take place to the information base (the model is foundationalist rather than coherentist). Furthermore, *minimal change* is easily achieved. Adding ϕ to a base, sequence or list (as in (Chopra et al., 2000) and this paper) is a minimal change to effect the acceptance of ϕ; the set-theoretic deletion of ϕ is a minimal change that removes ϕ. The latter is not the only way to model loss of beliefs. A more sophisticated model would make contraction *reason-based*: to lose belief in a proposition would be to revise by information that renders it underivable from the belief corpus. This is the method followed in the Ghose-Goebel framework since there are no explicit deletions, but lack of commitment to a belief corresponds to the non-entailment of that belief via the associated inference relation—the same method is employed in (Chopra et al., 2000). Lastly, the AGM operation of *expansion* disappears: all belief change is recorded via changes in the set of consequences of the inference relation.

The inference relation that enables the mapping between the agent's information base and its set of accepted beliefs is understood to be non-monotonic and/or paraconsistent. However such an inference operation does not imply the use of a paraconsistent *logic* since the resolution of inconsistencies takes place via the use of parameters such as the syntactic structure or prioritization of the base. Epistemic states in this scenario are the non-monotonic consequences of the inference relation that is devised. This relationship is depicted in the figure below (Rott, 1996). Note that the epistemic states are *constructed by the inference operation* and that change operations can be carried out on these if so required.

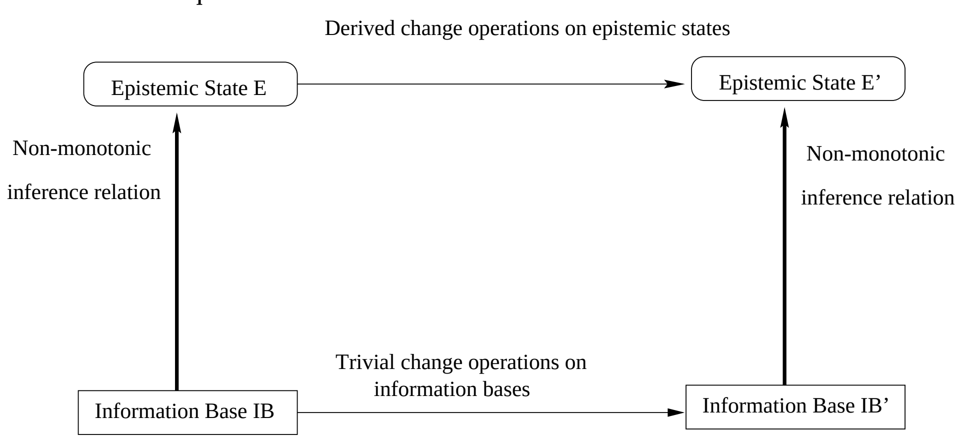

Figure 1. The vertical mode of belief change

The novel feature of the Ghose-Goebel approach is the introduction of explicit *disbeliefs*, on par with explicit beliefs, as allowable epistemic inputs. The explicit disbeliefs are used to record and store sentences that an agent refuses to commit to. This allows for a level of expression that cannot be obtained with sets just containing beliefs. Why is it useful to think of representing disbeliefs explicitly? Often this is a normal part of belief ascriptions ("Ricardo does not believe that it is the case that he will lose at the races"); an explicit representation of disbeliefs then, does justice to pretheoretic talk of belief ascription. Talk of disbeliefs is plausible in some famous paradoxes of belief such as the Lottery and Preface paradoxes. For example, in the lottery paradox, let $t_1, \ldots, t_n$ represent the tickets in the lottery. It seems a better expression of an agent's epistemic attitude towards a particular ticket's chances of winning the lottery to say that the agent disbelieves the proposition "t_i will win" rather than to say that the agent believes the proposition that "t_i will not win". A similar case could be made for disbelief ascription in the Preface paradox ("I disbelieve the proposition that there are no errors in the book"). As another example, I know both murder suspects A and B well, and am unwilling to believe that A committed the murder (and similarly B). Yet I know that these were the only two people at the scene of the crime and must believe the disjunction. It seems a better expression of my epistemic attitudes to say that I refuse to commit myself to the belief that A did it, rather than to say that I believe that A did not do it.

Intuitively, the assertion of a disbelief is analogous to belief contraction; standard base contraction operations can be recovered in this model. The disbeliefs considered for explicit representation in this model are the analogs of explicit beliefs; they are foundational in some sense. This is also necessary if the use of disbeliefs is to not be rendered a purely derivative notion based on beliefs: there is an interesting set of explicit failures of epistemic commitment that are worth recording and representing and using in a model of belief change. If this were not the case, then all of an agent's disbeliefs could simply be recovered from the agent's belief set.

Interestingly, the explicit representation of disbeliefs allows a view of belief contraction as independently useful, instead of merely acting as an intermediate step, via the Levi identity (Alchourrón et al., 1985), in a belief revision operation. In the belief revision literature, contraction is taken to be a primary operation, and postulates characterizing it are offered. However, there is an understanding at hand that pure contraction is hard to actually come by; contraction is often understood as revision by inputs that cause the non-entailment of the contracted belief from the agent's corpus. Standard examples of contraction, then, are easily reinterpreted as simply examples of revision. However, contraction *can* be understood as an independent epistemic operation. One way of understanding contraction is in the forms of constraints ("modify your belief corpus in such a way so as to render a particular belief underivable"). Another is to provide a formal way to deal with discredited sources ("all information received up till now from Carlos is untrustworthy") and yet another is to understand contraction as *hypothetical* belief change. In the latter, a belief is given up in order to provide a forum for a sympathetic assessment of a contradicting belief ("I do not believe as you do, that we will face an energy crisis in this century. However, for the sake of argument, assume that neither of us knows what will transpire"). Such a scenario is however, 'merely' hypothetical and does not constitute 'real' belief change. Is it worth formalizing then? Our concern is simpler: given that we can understand contraction as a peculiar idealized form of belief change, is a formalization and representation of it possible in which it emerges as an independent operation, not reducible to revision and not understood as an intermediary operation? We think the answer to this question is in the affirmative.

The framework of Ghose and Goebel assumes a *linear ordering* of reliability on epistemic inputs. This ordering is sufficient to obtain consistency—but not necessary—and leads to an associated class of non-monotonic inference relations. But the world that a reasoner finds itself in is unlikely to be so accomodating of epistemic dilemmas: the assumption of linearity renders their system unable to deal with any situation in which two or more epistemic inputs are regarded as equally reliable. In this paper we extend their 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*. Rather than simply assessing the relative positions of beliefs in rankings, it enables more granular distinctions to be made. These rankings are also less restrictive than linear orderings since they induce *total preorders* on inputs. These preorders can be thought of as composite orderings which take into account the various factors that influence an agent's preference ranking of its epistemic inputs: temporal ordering, source reliability and so on. The freedom brought about by the move from linear orderings to rankings introduces the problem of dealing in an adequate fashion with inputs of *equal* rank. We provide a semantic solution to this problem by replacing all inputs with logically weakened versions, from which an appropriate entailment relation is defined. The intuitive correspondence to this operation is an agent weakening the foundational inputs on which its whole corpus is based in order to recover a coherent picture of the world. **To formally capture all of the intutions expressed in the discussion above, we provide a theorem—the central result in this paper—that shows that our model is a true generalization of the AGM model, since it contains as a special case, AGM belief change on closed theories. This result also demonstrates formally, the capturing of AGM models of belief change by those based on the provision of an appropriate non-monotonic inference operation on the agent's beliefs.**

The outline of this paper is as follows. We first establish some formal preliminaries. In Section 2 we describe the Ghose-Goebel approach to non-prioritized belief change and point out its strengths and weakenesses. In Section 3 we establish the notion of *ranked belief bases* and deal with the problem of inputs of equal reliability. In Section 4 we bring all this material together to describe our model, and show how AGM belief change on closed theories can be embedded in it. We then draw some conclusions and point to directions for future work.

We assume a finitely generated propositional language L closed under the usual propositional connectives $(\wedge, \neg, \vee, \rightarrow)$ and equipped with a classical model-theoretic semantics; the constants $\top, \bot$ are in L. 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. a reflexive and 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-prioritized belief change

This section is our account of the Ghose-Goebel approach to non-prioritized belief change. We emphasise the role played by the linear orderings on epistemic inputs, an issue previously glossed over. This leads us to a formal description of the logic of beliefs and disbeliefs which is implicit in their original description. Although this proposal for non-prioritized belief change is based on default theories (Reiter, 1980; Poole, 1988) our account does not refer to default reasoning.

The *information* available to an agent is represented in the form of *beliefs* and *disbeliefs*. Beliefs are *positive* epistemic inputs which an agent will, under suitable conditions, commit to for purposes such as query-answering and acting. Disbeliefs are *negative* epistemic inputs which an agent will, under suitable conditions, *refuse* to commit to. 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. This leads to a language syntactically richer than the normal propositional ones by allowing for explicit representation of disbeliefs . Note that overlining does not apply to sub-formulas, but only to sentences considered as atomic wholes i.e., the agent's disbelief extends to propositions.

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. □

Given a set of beliefs and disbeliefs, a unique information state is determined. 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 $B \cup \{\neg\delta\} \nvdash \bot$ or alternatively if $B \not\models \delta$. An information state I is *consistent* iff $B_I \nvdash \bot$ and every $\overline{\delta} \in D_I$ is consistent with B_I. □

For an information state I to be consistent it has to be the case that the belief state B_I is consistent in the classical sense. 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. This distinction in the kinds of consistency being discussed will be useful in the discussion to follow.

The explicit expression of beliefs and disbeliefs is reminiscent of the handling of knowledge and belief in modal logics, specifically in S5-models (Hughes and Cresswell, 1972; Fagin et al., 1995). The statement $\Box\alpha$ in an S5-model corresponds to the assertion, in our framework, that α is believed. Similarly, the statement $\Diamond\neg\gamma$ in an S5-model corresponds to the assertion $\overline{\gamma}$, in our framework: that γ is disbelieved. It can easily be verified that the notions of consistency in these two frameworks coincide. Hansson (Hansson, 1999) draws a similar connection between modal notions of necessity and possibility and the formal framework required to express non-prioritized belief change operations.

But how is one to not commit the agent to incoherence if it possesses an inconsistent information state? The challenge 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.

DEFINITION 2.3. J is a maximally consistent subset of I iff (i) $J \subseteq I$ (ii) J is consistent (iii) 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 $I = \{p, p \to q, \overline{q}\}$. 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 $p \to q$. The maximally consistent subset generated by this reliability ordering is $\{p, \overline{q}\}$. This leads us to the following:

DEFINITION 2.4. A maximally consistent subset J of I is $<$*-generated* by a linear order $<$ on I iff (i) every $\overline{\delta} \in D_I \setminus J$ is inconsistent with $\{\beta \in B_J \mid \beta > \overline{\delta}\}$, and (ii) $\forall \alpha \in B_I \setminus J$, 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 (unique, as proved by Ghose and Goebel) $<$-generated maximally consistent subset of the information state I; $B_{I^<}$ and $D_{I^<}$ are the associated sets of beliefs and disbeliefs respectively. Obtaining $I^<$ is simply a matter of considering the elements of I one by one in the order prescribed by $<$, starting from the top, 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 beliefs to be entailed by an information state.

DEFINITION 2.5. Let I be an information state 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 reliability ordering is said to be *committed* to the information $<$-entailed by I. □

The following example illustrates 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 $I \models^< \overline{p \wedge q}$. □

It is perhaps necessary to expand somewhat on the notion of $<$-entailment for disbeliefs. Disbeliefs are, dual to beliefs—in the following partial sense. 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 commit to every disbelief $\overline{\gamma}$ where γ is logically *stronger* than δ. An agent that believes α is also committed to $\alpha \vee \beta$. If I believe that the world economy will be hit by recession then I am also committed to believing that either the world economy will be hit by recession or that I will continue to retain my job. An agent that is committed to a disbelief $\overline{\delta}$ is also committed to the disbelief $\overline{\delta \wedge \gamma}$. If I do not believe that there will be an energy crisis next year, I do not believe either that there will be an energy crisis and that the Netherlands will win the Soccer World Cup. As extreme cases, an agent is always committed to believing tautologies and disbelieving contradictions. Weakening a disbelief $\overline{\delta}$, as we do in Section 3, 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.

Much of classical base change (Fuhrmann, 1991; Hansson, 1993; Rott, 1996) is based on maximal consistency. The contraction of a base B by a sentence α is carried out by generating all maximally consistent subsets of B which do not entail α; selection functions can then be used to pick out the 'best' of these subsets. The revision of B by α is defined, via the Levi identity (Alchourrón et al., 1985), 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 framework. 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 as an illustration.

EXAMPLE 2.2. A totally ignorant agent (i.e., one beginning with an epistemic *tabula rasa*) 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, some of the relevant inputs, specifically the contraction of q, have not been taken into account: part of the agent's intellectual history has been ignored. 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. Note then, that the explicit representation of disbeliefs provides for a way to keep track, in the belief representation, of not just revision, but contraction histories as well. □

In contrast with traditional approaches to belief change, *iterated belief change* has a solution built into this framework, 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.

Note that in this model, one of the important conceptual weaknesses of models of non-prioritized belief change is overcome: the fact that every new piece of information that is logically compatible with older information is accepted. This is true at the level of simply updating the information base of the agent, but it is certainly not true of what can be *inferred* from the agent's information state.

3. The problem of equal reliability

While $<$-entailment provides an interesting deployment of non-prioritized 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, 1988). 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 (Spohn's original construction assigned ordinal ranks to beliefs).

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.

The possibility of information in a ranked information state being equally reliable introduces a problem which does not crop up in the Ghose-Goebel framework; what to do with mutually inconsistent pieces of information that are equally reliable. For example, I am told by my friend that it is the case that my favorite sporting team won last night. I am also told that they lost by another equally reliable source. Who do I believe till receiving confirmation one way or the other? How do I treat this information?[2]

This section is devoted to finding a solution to this problem. We will be be concerned with multisets of information (not ranked information) and will 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 4 the results obtained in this section will be applied to ranked information states. One way of dealing with inconsistent *beliefs* of equal reliability in information multisets is to employ *infobases* (Meyer, 1999; Meyer et al., 2000). These are defined as a list, or sequence, of sentences. However, the reason for this is to allow for multiple copies of the same sentence. Since in our context, the order of the sentences in an infobase is irrelevant, we define infobases as finite multisets of sentences. 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* in the ranking being regarded as *more* plausible.

DEFINITION 3.2. An infobase B is a finite multiset of sentences. 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 $num_B(v) \leq num_B(u)$. □

We shall extend the terminology used for sets to apply to multisets as well e.g., $M(B)$, where B is a multiset, refers to the set of models of all the wffs in B, $|B|$ refers to the cardinality of B, and so on. 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 3.1. For the infobase $B = \{p, q\}$, $num_B(11) = 2$, $num_B(10) = num_B(01) = 1$, 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 infobases is that they allow for the extraction of a consistent knowledge base from *any* infobase B, even

[2] As an aside, note that use of linear orders also places a certain kind of syntax sensitivity on the model. Receiving $\alpha \wedge \beta$ and receiving α, β separately has different effects on the agent's epistemic state since different linear orders are imposed on these formulas.

if beliefs occurring in it are inconsistent. The idea is simple; take the minimal elements of U, with respect to $\preceq_B$ (i.e., the minimal models of the tautologies), to be the models of a sentence $K(\preceq_B)$ where, in the spirit of (Katsuno and Mendelzon, 1991), $K(\preceq_B)$ is regarded as a finite representation of its own logical closure.

DEFINITION 3.3. 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 (Grove, 1988; Kraus et al., 1990; Katsuno and Mendelzon, 1991).

EXAMPLE 3.2. 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 3.1. *An infobase B is consistent iff $K(\preceq_B) \equiv \bigwedge_{\alpha \in B} \alpha$.*

Proof: The only-if direction is straighforward. For the other direction, note that since B is consistent it follows for every $u \in U$ that $num(u) = |B|$ iff $u \in M(B)$. And since $num(u) \leq |B|\ \forall u \in U$, it follows immediately that $M_{\preceq}(\top) = M(B)$, and therefore $K(\preceq_B) \equiv \bigwedge_{\alpha \in B} \alpha$. □

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. 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. The definition for consistency of information multisets is an obvious extension of definition 2.2.

DEFINITION 3.4. Let B be an infobase. 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 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 3.5. 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 conjunction of all 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.

EXAMPLE 3.3. 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. □

The following result shows that it is no accident in the example above that $\{p \vee q, \overline{p}, \overline{q}\}$ is consistent.

PROPOSITION 3.2. *I is an information multiset such that $\forall\overline{\delta} \in D_I, \delta \not\equiv \top$ iff $\{K(\preceq_I)\} \cup D_I$ is consistent.*

Proof: The only-if direction is straightforward. For the other direction, it suffices to show that every disbelief in I is consistent with $K(\preceq_I)$. So pick a $\overline{\delta} \in D_I$. From definition 3.5 it follows directly that $\forall u \in U$, $num_I(u) \leq \max\{num_{B_I}(v) \mid v \in U\}$, and therefore that, $\forall v \in U$ and $\forall u \in M_{\preceq_I}(\neg\delta)$, $num_I(v) \geq num_I(u)$. So $M_{\preceq_I}(\neg\delta) \subseteq M_{\preceq_I}(\top)$. And since $\delta \not\equiv \top$ it follows that $\overline{\delta}$ is consistent with $K(\preceq_I)$. □

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. It seems necessary to weaken these disbeliefs, since they are judged to be no more important or reliable than the beliefs in B_I, but have not been subject to the same kind of weakening than the elements of B_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 3.6. 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 3.3, demonstrates the use of weakening.

EXAMPLE 3.4. 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 3.3. *I is consistent iff:*

- $K(\preceq_{B_I}) \equiv \bigwedge_{\alpha \in B_I} \alpha$ *and,*
- $W^{K(\preceq_{B_I})}(\delta) = \delta$ *for every* $\overline{\delta} \in I$.

Proof: The only-if part is trivial (by definition 3.6). For the other direction, suppose I is a consistent information multiset. It follows from proposition 3.1 that $K(\preceq_{B_I}) \equiv \bigwedge_{\alpha \in B_I} \alpha$. That $W^{K(\preceq_{B_I})}(\delta) = \delta$ for every $\overline{\delta} \in I$ follows directly from definition 3.6. □

4. Non-prioritized ranked belief change

We are now in a position to describe our proposal for non-prioritized 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

3, which respects the ranks assigned to information. Classical entailment of a belief from this multiset establishes its rank-entailment from the original ranked information state.

DEFINITION 4.1. For 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. 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 (Benferhat et al., 1993).

DEFINITION 4.2. For $r \leq \max$, the total preorder $\preceq^r$ on U *generated* by RI is defined as follows: $u \preceq^r v$ iff $\forall i \in \{r, \ldots, \max\}$, $v \prec_{B^i} u$ implies $\exists j \in \{i+1, \ldots, \max\}$ s.t. $u \prec_{B^j} v$. The knowledge base *extracted* from $\preceq^r$ is denoted by $K(\preceq^r)$. □

In the discussions below, we follow the convention of using primes to denote information states deviating slightly from RI, and we decorate constructions obtained from these information states with an appropriate number of primes. So, for example, $\preceq'^r$ refers to the total preorder $\preceq^r$ obtained from the information state RI' and WD' refers to the set of disbeliefs to be weakened from the information state RI'.

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

Proof: The only-if direction is straightforward. For the other direction, the proof is by induction on s where $r = \max - s$. If $s = 0$ then $r = \max$ and the result follows immediately from definition 4.2 and proposition 3.1. Now assume that the result holds for s where $s < \max$. We show that the result holds for $s + 1$. Pick any $u \in M(\bigwedge_{\alpha \in B}) \cap M(\bigwedge_{\alpha \in B^{r-1}})$. By the induction hypothesis $u \in M(K(\preceq^r))$ and so $u \preceq^r v$ $\forall v \in U$. To show that $u \in M(K(\preceq^{r-1}))$ it suffices to show that $u \preceq^{r-1} v$ $\forall v \in U$. And to do this it suffices, by definition 4.2 and the induction hypothesis, to show that $u \preceq_{B^{r-1}} v, \forall v \in U$, which follows from proposition 3.1, since $u \in M(\bigwedge_{\alpha \in B^{r-1}})$. Conversely, pick any $u \in M(K(\preceq^{r-1}))$ and assume that $u \notin M(\bigwedge_{\alpha \in D})$ where $D = \bigcup_{i=r-1}^{i \leq \max} B^i$. So, by proposition 3.1, there is a $j \in$

$\{r-1,\ldots,\max\}$ such that $u \notin M(K(\preceq_{B^j}))$. And since D is consistent there is a $v \in M(\bigwedge_{\alpha\in D})$. It then follows easily from definition 4.2 that $v \prec^{r-1} u$, contradicting the assumption that $u \in M(K(\preceq^{r-1}))$. So $v \preceq_{B^i} u, \forall i \in \{r-1,\max\}$ and $v \prec_{B^j} u$. □

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

EXAMPLE 4.1. Let $RI = \{(p,2),(p\to q,1),(\neg q,0)\}$. Then $B^0 = \{\neg q\}$, $B^1 = \{p\to q\}$ 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. The motivation for this should be obvious: we seek to retain as much of our original belief state as possible while being respectful of the conflicts that exist between beliefs and disbeliefs.

DEFINITION 4.3. 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 4.4. 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 3.5, 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 $\preceq^0$, of the negation of every disbelief of every rank r which is to be retained and weakened.

DEFINITION 4.5. The knowledge base *extracted* from the information state RI is a sentence $K(RI)$ such that

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

□

Secondly, for every rank r, the disbeliefs in WD^r are weakened in the manner described in definition 3.6 and added to the disbeliefs to be retained.

DEFINITION 4.6. The set of disbeliefs $D(RI)$ *associated* with RI is defined as

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

□

This brings us to the definition of rank-entailment.

DEFINITION 4.7. A belief α is *rank-entailed* by RI, denoted by $RI \Vdash \alpha$, iff $K(RI) \vDash \alpha$. A disbelief $\overline{\delta}$ is *rank-entailed* by RI, denoted by $RI \Vdash \overline{\delta}$, iff $\delta \equiv \bot$ or $\delta \vDash \gamma$ for some $\overline{\gamma} \in D$. An agent with RI as its ranked information state is *committed* to the information rank-entailed by RI. □

Since under the approach described here, epistemic states are simply the consequences of the entailment relation associated with an information state, we have the following:

DEFINITION 4.8. The epistemic state, E_{RI}, associated with a ranked information state RI is given by $E_{RI} = \{\alpha | RI \Vdash \alpha\} \cup \{\overline{\delta} | RI \Vdash \overline{\delta}\}$. □

The definitions above, then, make it possible to speak of the agent as being committed to its epistemic state. The next example illustrates the use of rank-entailment.

EXAMPLE 4.2. Let $RI = \{(p \to q, 2), (\overline{p}, 2), (p, 1), (\overline{p \wedge q}, 1), \overline{p \to q}, 1),$ $(\neg q, 0)\}$. It can be verified that $K(\preceq^2) \equiv p \to q$, $K(\preceq^1) \equiv p \wedge q$ 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 $\overline{p \to q} \in D^1$ is inconsistent with $K(\preceq^2)$ and should therefore be removed. Since $M_{\preceq^0}(\neg p) = M_{\preceq^0}(\neg(p \wedge q)) = \{00\}$, it follows that $K(RI) \equiv (p \leftrightarrow q)$. 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 $RI \Vdash p \leftrightarrow q$, $RI \Vdash \overline{p}$, 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. It is well-known (Katsuno and Mendelzon, 1991) that the different AGM revision and contraction operations can be obtained from the total preorders on U. In our framework, 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.

THEOREM 4.1. *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 (*Meyer, 1999*).) Then $K * \alpha \vDash \beta$ 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 (*Meyer, 1999*).) Then $K - \alpha \vDash \beta$ iff β is rank-entailed by $\{(\rho, 0) \mid \rho \in B\} \cup \{(\overline{\alpha}, 1)\}$.*

Proof:

1. From (Katsuno and Mendelzon, 1991) we have that $M(K*\alpha) = M_{\preceq_B}(\alpha)$. Since RI does not contain any disbeliefs, it follows that $M(K(RI)) = M(K(\preceq^0))$. So it suffices to show that $M_{\preceq_B}(\alpha) = M(K(\preceq^0))$. Now observe that it follows directly from definition 4.2 that $M_{\preceq^1}(\top) = M(\alpha)$, and furthermore, that $M_{\preceq^0}(\top) = M_{\preceq_B}(\alpha)$, from which the result follows directly.

2. Observe that $RD^1 = \{\overline{\alpha}\}$ is the set of disbeliefs to be retained from D^1, and so, by definition, $M(K(RI)) = M(K(\preceq^0)) \cup M_{\preceq^0}(\neg\alpha)$. Also, from (Katsuno and Mendelzon, 1991) we have that $M(K - \alpha) = M(K) \cup M_{\preceq_B}(\alpha)$. Now, since $B^1 = \emptyset$, it follows that $\preceq^1 = \preceq^0 = \preceq_B$. So $K(\preceq_0) \equiv K$ and $M_{\preceq^0}(\neg\alpha) = M_{\preceq_B}(\neg\alpha)$, from which the result follows directly.

□

The result above brings together the various intutions that have served as the motivation of our model and demonstrates that it represents a richer, more philosophically satisfying model of belief change. The model provides a richer language—via the introduction of explicit disbeliefs—that allows for expression of finer grained expressions of epistemic commitments and an understanding of contraction as a genuinely independent epistemic operation Its notion of belief change addresses the issue of judgement of beliefs since new inputs do not possess some artificially

privileged status. Previous models of belief change are plausibly viewed as incomplete since they do not address this issue, choosing instead to begin their modeling at the point that a new epistemic input has been accepted. Our model confirms the the reduction of the problem of belief revision to the construction of an adequate non-monotonic inference relation, an intuitive connection that finds further formal grounding here. Lastly the model addresses an important part of epistemic judgement carried out by rational agents: the resolution of mutually inconsistent, yet equally reliable epistemic inputs. The semantic nature of model is also satisfying since it does justice to the intutive notion of agents ranking alternative states of affairs as being more or less plausible according to some agreed upon scale.

5. Properties of Iterated Revision Operator

In this section we briefly examine the iterated revision properties introduced in our framework: we do this by translating two sets of axioms suggested as desirable properties for iterated revision. Before we begin, it is worth pointing out the commutativity of our operations. That is, adding (α, r) to a ranked information state RI followed by the addition of (β, s), yields the same results as the case where (β, s) is added before (α, r).

5.1. Conformance with Lehmann Axioms

In Lehmann's sequences model (Lehmann, 1995), a belief state results from a sequence of revisions—the individual revisions themselves are just those by consistent formulas. A concatenation of these sequences is denoted by $\circ$ and sequences of length one are formulas.[3] We can think of the sequence $\sigma \circ \alpha \circ \beta \circ \rho$ as denoting, first, the formulas of the sequence σ, then the result of the revision of σ by α, then by β and then by the formulas of the sequence ρ. The belief set that results from a sequence σ of individual revisions is denoted $[\sigma]$; $\circ$ denotes a revision procedure. The belief set $[\sigma \circ \alpha]$ denotes the result of revising σ by the formula α. We do not identify $[\sigma \circ \alpha]$ with $[\sigma] * \alpha$ where $*$ is an AGM operator: such an identification causes problems since Lehmann considers operations on belief *states*, while AGM considers operations on belief *sets*. The AGM postulates require that if $[\sigma] = [\rho]$, then $[\sigma]*\alpha = [\rho]*\alpha$, but in this framework $[\sigma] = [\rho]$ does not imply $[\rho \circ \alpha] = [\sigma \circ \alpha]$. Lehmann's framework allows a reasoning agent to base its revision not just on the belief set $[\sigma]$ but also on the sequence itself i.e., on the history of the agent. In this way, the temporal ordering of the beliefs plays a part in the revisions that the agent carries out. Given the definitions above, a set of axioms is

[3] Lehmann's original notation has been modifi ed for greater readability.

presented for revisions of belief sequences (Cn is the classical operator for logical consequence):

(**I1**) $[\sigma]$ is a consistent theory

(**I2**) $\alpha \in [\sigma \circ \alpha]$

(**I3**) If $\beta \in [\sigma \circ \alpha]$, then $\alpha \to \beta \in [\sigma]$

(**I4**) If $\alpha \in [\sigma]$, then $[\sigma \circ \rho] = [\sigma \circ \alpha \circ \rho]$

(**I5**) If $\beta \models \alpha$, then $[\sigma \circ \alpha \circ \beta \circ \rho] = [\sigma \circ \beta \circ \rho]$

(**I6**) If $\neg\beta \notin [\sigma \circ \alpha]$, then $[\sigma \circ \alpha \circ \beta \circ \rho]$ = $[\sigma \circ \alpha \circ \alpha \wedge \beta \circ \rho]$

(**I7**) $[\sigma \circ \neg\beta \circ \beta] \subseteq Cn([\sigma], \beta)$

The following are translations of the Lehmann axioms with reference to a ranked information state RI; E_{RI} is as defined in Definition 4.8. Furthermore, we define $RI \circ \alpha$ as $RI \cup \{\alpha, r\}$.

(L1) E_{RI} is consistent

(L2) $\alpha \in E_{RI \cup \{(\alpha,s)\}}$ if $r > \max\{s \mid (\alpha, s) \in RI\}$ or $E_{RI} \vDash \alpha$

(L3) If $\beta \in E_{RI \cup \{(\alpha,r)\}}$ then $\alpha \to \beta \in E_{RI}$

(L4) If $\alpha \in E_{RI}$ then $E_{RI \cup RI'} = E_{RI \cup \{(\alpha,r)\} \cup RI'}$

(L5) If $\beta \vDash \alpha$ then $E_{RI \cup \{(\alpha,r),(\beta,r),(\gamma,r)\}} = E_{RI \cup \{(\beta,r),(\gamma,r)\}}$

(L6) If $\neg\beta \notin E_{RI \cup \{(\alpha,r)\}}$ then $E_{RI \cup \{(\alpha,r),(\beta,r)\} \cup RI'} = E_{RI \cup \{(\alpha,r),(\alpha \wedge \beta,r)\} \cup RI'}$

(L7) $E_{RI \cup \{(\neg\beta,r),(\beta,r)\}} \subseteq Cn(E_{RI}, \beta)$

PROPOSITION 5.1. *(L1), (L2), (L3) and (L7) hold for a given ranked information state RI; (L4), (L5), (L6) do not.*

Proof: (L1) is trivial. For (L2) let $RI' = RI \cup \{(\alpha, r)\}$. Now suppose that $r > \max\{s \mid (\beta, s) \in RI\}$. Then $M_{\preceq'^r}(\top) = M_{\preceq'_{B^r}}(\top) = M(\alpha)$ and so $M_{\preceq'^0}(\top) \subseteq M(\alpha)$. Also, by the choice of r, it follows that if $\overline{\delta} \in RD' \cup WD'$ then $M(\neg\delta) \cap M_{\preceq'^0}(\top) \neq \emptyset$. From this it follows that $M_{\preceq'^0}(\neg\delta) \subseteq M_{\preceq'^0}(\top)\ \forall\overline{\delta} \in RD' \cup WD'$, from which the result follows. For the case where $\alpha \in E_{RI}$, let $RI' = RI \cup \{(\alpha, r)\}$. So $M_{\preceq^0}(\top) \subseteq M(\alpha)$ and from this it follows easily that $M_{\preceq'^0}(\top) \subseteq M(\alpha)$. Now we first show that $RD' \cup WD' \subseteq RD \cup WD$. Pick any $(\overline{\delta}, s) \in RI$ such that $\delta \notin (RD \cup WD)$. Then $\overline{\delta}$ is inconsistent with $K(\preceq^i)$ for some $i > s$. And since $K(\preceq^i) \subseteq K(\preceq'^i)$, $\overline{\delta}$

is also inconsistent with $K(\preceq'^i)$. So $\delta \notin (RD' \cup WD')$. Now, by supposition, $M_{\preceq^0}(\neg\delta) \subseteq M(\alpha)$ for every $\delta \in WD \cup RD$. But then it also follows that $M_{\preceq'^0}(\neg\delta) \subseteq M(\alpha)$ for every $\delta \in WD \cup RD$ and therefore $M_{\preceq'^0}(\neg\delta) \subseteq M(\alpha)$ for every $\delta \in WD' \cup RD'$, from which the result follows. For (L3) let $RI' = RI \cup \{(\alpha, r)\}$. Now pick any $u \in M(K(RI)) \cap M(\alpha)$. If $u \in M_{\preceq^0}(\top)$ then it follows immediately that $u \in M_{\preceq'^0}(\top)$ and so $u \in M(\beta)$. Otherwise $u \in M_{\preceq^0}(\neg\delta)$ for some $\delta \in RD \cup WD$. Since $u \in M(\alpha)$ it then follows that $u \in M_{\preceq'^0}(\neg\delta)$ and also that $\delta \in WD' \cup RD'$. And therefore $u \in M(\beta)$. For (L4), let $RI = \{(p \wedge q, 0)\}$, let $\alpha = p$ and $r = 0$ and let $RI' = \{(\neg q, 0)\}$. For (L5), let $RI = \emptyset$, $\alpha = p$, $\beta = p \wedge q$, $\gamma = \neg q$ and $r = 0$. For (L6), let $RI = \emptyset$, $\alpha = p$, $\beta = q$, $RI' = \{(\neg(p \wedge q), 0)\}$ and $r = 0$. For (L7), let $RI' = RI \cup \{(\neg\beta, r), (\beta, r)\}$. It is easily seen that $\preceq_{B^r} = \preceq_{B^r \cup \{\neg\beta, \beta\}}$, and therefore $\preceq^0 = \preceq'^0$. Also then, $WD \cup RD = WD' \cup RD'$, and so $K(RI) \equiv K(RI')$, from which the result follows. □

Observe that in (L2) it is not sufficient for α to be consistent with E_{RI} since there may be disbeliefs in RI amounting to a refusal to commit to α, with a rank higher than the rank of any of the beliefs in RI e.g., $RI = \{(\overline{p}, 1)\}$ and $\alpha = p$. Also, note that the (L4), (L5), (L6) counterexamples use only beliefs, not disbeliefs.

5.2. Conformance with the Darwiche-Pearl Axioms

In (Darwiche and Pearl, 1994) a framework for iterated belief revision is presented via the four axioms below ($\circ$ is the update operator, α, μ represent new epistemic inputs, ψ represents a belief state).

(C1) If $\alpha \models \mu$, then $(\psi \circ \mu) \circ \alpha \equiv \psi \circ \alpha$.

(C2) If $\alpha \models \neg\mu$, then $(\psi \circ \mu) \circ \alpha \equiv \psi \circ \alpha$.

(C3) If $\psi \circ \alpha \models \mu$, then $(\psi \circ \mu) \circ \alpha \models \mu$.

(C4) If $\psi \circ \alpha \not\models \neg\mu$, then $(\psi \circ \mu) \circ \alpha \models \mu$.

The following are translations of the DP axioms (as before $RI \circ \alpha = RI \cup \{\alpha, r\}$).

(C1) If $\alpha \vDash \mu$ then $E_{RI \cup \{(\mu,r),(\alpha,r)\}} = E_{RI \cup \{(\alpha,r)\}}$

(C2) If $\alpha \vDash \neg\mu$ then $E_{RI \cup \{(\mu,r),(\alpha,r)\}} = E_{RI \cup \{(\alpha,r)\}}$

(C3) If $E_{RI \cup \{(\alpha,r)\}} \vDash \mu$ then $E_{RI \cup \{(\mu,r),(\alpha,r)\}} \vDash \mu$

(C4) If $E_{RI \cup \{(\alpha,r)\}} \nvDash \neg\mu$ then $E_{RI \cup \{(\mu,r),(\alpha,r)\}} \vDash \mu$

PROPOSITION 5.2. *(C1) holds if $r > \max\{s \mid (\beta, s) \in RI\}$. (C2) doesn't hold. (C3) holds, and (C4) holds if $D_{RI} = \emptyset$ i.e., if RI contains no disbeliefs.*

Proof: For (C1), suppose that $r > \max\{s \mid (\beta, s) \in RI\}$ and let $RI' = RI \cup \{(\alpha, r)\}$ and $RI'' = RI \cup \{(\mu, r), (\alpha, r)\}$. It is easily seen that $K(\preceq'_{B^r}) \equiv K(\preceq''_{B^r})$ and therefore $K(\preceq'^0) \equiv K(\preceq''^0)$. Also, by the choice of r, it follows that if $\overline{\delta} \in RD' \cup WD'$ then $M(\neg\delta) \cap M_{\preceq'^0}(\top) \neq \emptyset$ and if $\overline{\delta} \in RD'' \cup WD''$ then $M(\neg\delta) \cap M_{\preceq''^0}(\top) \neq \emptyset$. From this it follows that $M_{\preceq'^0}(\neg\delta) \subseteq M_{\preceq'^0}(\top)\ \forall\overline{\delta} \in RD' \cup WD'$ and $M_{\preceq''^0}(\neg\delta) \subseteq M_{\preceq''^0}(\top)\ \forall\overline{\delta} \in RD'' \cup WD''$. And then the result follows. For (C2), let $RI = \emptyset$, $\alpha = p$, $\mu = \neg p$. (C3) follows from the postulate (L2), which was proved in proposition 5.1. For (C4), let $RI' = RI \cup (\alpha, r)$ and $RI'' = RI' \cup (\mu, r)$ and suppose that $D_{RI} = \emptyset$ and $E_{RI'} \not\models \neg\mu$. Thus, $M_{\preceq'^0}(\top) \cap M(\mu) \neq \emptyset$. Since $M_{\preceq''^0}(\top) \subseteq M_{\preceq'^0}(\top)$ it then follows that $M_{\preceq''^0}(\top) \subseteq M(\mu)$. And since $D_{RI'} = \emptyset$, we have that $M(K(RI'')) = M_{\preceq''^0}(\top)$, from which the result follows. □

To show that (C1) does not hold in general, consider the case where $RI = \{(\neg p, 1)\}$, $\alpha = p$, $\mu = p \vee q$ and $r = 0$. To see that (C4) does not hold when disbeliefs occur in RI, consider the example where $RI = \{(\overline{p \wedge q}, 1)\}$, $\alpha = p, \mu = p \wedge q$ and let $r = 0$. The condition for (C4) to hold shows that the explicit introduction of disbeliefs really does lead to a more expressive language since now, we can explicitly relate the set of disbeliefs to properties of iterated revision on our epistemic states.

Observe that (C3) and (C4) look different from the DP versions of (C3) and (C4) which would look like this:

($C3'$) If $E_{RI \cup \{(\alpha,r)\}} \models \mu$ then $E_{RI \cup \{(\mu,r),(\alpha,r)\}} \models \mu$

($C4'$) If $E_{RI \cup \{(\alpha,r)\}} \not\models \neg\mu$ then $E_{RI \cup \{(\mu,r),(\alpha,r)\}} \models \mu$

$C3'$ follows from (C3) and $C4'$ follows from (C4), but not conversely.

6. Conclusion and future work

In this paper we have provided a different perspective on the model for non-prioritized belief change based on default theories (Ghose and Goebel, 1998). 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. **Finally, in bringing all these features together formally, our account of non-prioritized**

ranked belief change was shown to incorporate AGM belief change as a special case of our more general model.

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 modal logic frameworks would 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 (Dubois et al., 1994), in which necessity-valued statements may be seen as beliefs, and possibility-valued statements as disbeliefs.

What sort of rational agent has been modeled in this study? The agent is viewed as one that keeps track of the information that it has received. In situations that call for it to make an epistemic commitment, it does so by evaluating the information it has received in terms of its epistemic value. The reasoner resolves inconsistencies at the time of epistemic commitment; this resolution is aided by the ordering that it has placed on the information received. We leave it to the reader to judge whether such a picture of a reasoning agent is in accordance with his or her conceptions of rationality.

References

Alchourr´on, C. E., P. Gärdenfors, and D. Makinson: 1985, 'On the logic of theory change: Partial meet functions for contraction and revision'. *Journal of Symbolic Logic* **50**, 510–530.

Benferhat, S., C. Cayrol, D. Dubois, J. Lang, and H. Prade: 1993, 'Inconsistency Management and Prioritized Syntax-Based Entailment'. In: R. Bajcsy (ed.): *IJCAI-93. Proceedings of the 13th International Joint Conference on Artificial Intelligence held in Chambery, France, August 28 to September 3, 1993*. San Mateo, CA, pp. 640–645, Morgan Kaufmann.

Brewka, G.: 1991, 'Belief Revision in a Framework for Default Reasoning'. In: A. Fuhrmann and N. Morreau (eds.): *The Logic of Theory Change*, Lecture Notes in Artifi cial Intelligence, Number 465. Berlin: Springer, pp. 206–222.

Chopra, S., K. Georgatos, and R. Parikh: 2000, 'Relevance Sensitive Nonmonotonic Inference on Belief Sequences'. In: *Proceedings of NMR'2000*. Full version appears in *Journal of Applied Non-Classical Logics*, 2001.

Cross, C. B. and R. H. Thomason: 1992, 'Conditionals and knowledge-base-update'. In: P. Gärdenfors (ed.): *Belief Revision*, Vol. 29 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge: Cambridge University Press, pp. 247–275.

Darwiche, A. and J. Pearl: 1994, 'On the logic of iterated belief revision'. In: *Proceedings of Theoretical Aspects of Rationality and Knowledge*. pp. 5–23.

Dubois, D., J. Lang, and H. Prade: 1994, 'Possibilistic logic'. In: *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. 3. pp. 439–513.

Fagin, R., J. Halpern, Y. Moses, and M. Y. Vardi: 1995, *Reasoning about Knowledge*. MIT Press.

Fuhrmann, A.: 1991, 'Theory contraction through base contraction'. *Journal of Philosophical Logic* **20**, 175–203.

Galliers, J. R.: 1992, 'Autonomous belief revision and communication'. In: P. Gärdenfors (ed.): *Belief Revision*, Theoretical Computer Science. Cambridge University Press, pp. 220–246.

Ghose, A. and R. Goebel: 1998, 'Belief states as default theories: Studies in non-prioritised belief change'. In: H. Prade (ed.): *ECAI 98. 13th European Conference on Artificial Intelligence*. New York, pp. 8–12, Johm Wiley & Sons, Ltd.

Grove, A.: 1988, 'Two modellings for theory change'. *Journal of Philosophical Logic* **17**, 157–170.

Hansson, S.: 1997, 'Semi-revision'. *Journal of Applied Non-Classical Logics* **7**(1–2), 151–175.

Hansson, S.-O.: 1993, 'Changes of disjunctively closed bases'. *Journal of Logic, Language and Information* **2**(4), 255–284.

Hansson, S. O.: 1999, *A textbook of belief dynamics*. Kluwer Academic Publishers.

Hughes, G. and M. Cresswell: 1972, *An introduction to Modal Logic*. Methuen.

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

Kraus, S., D. Lehmann, and M. Magidor: 1990, 'Nonmonotonic Reasoning, Preferential Models and Cumulative Logics'. *Artificial Intelligence* **44**, 167–207.

Lehmann, D.: 1995, 'Belief revision, revised'. In: *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*. pp. 1534–1540.

Levi, I.: 1977, 'Subjunctives, dispositions and changes'. *Synthese* **34**, 423–455.

Meyer, T.: 1999, 'Basic Infobase Change'. In: N. Foo (ed.): *Advanced Topics in Artificial Intelligence*, Vol. 1747 of *Lecture Notes In Artificial Intelligence*. Berlin, pp. 156–167, Springer-Verlag.

Meyer, T. A., W. A. Labuschagne, and J. Heidema: 2000, 'Infobase Change: A First Approximation'. *Journal of Logic, Language and Information* **9**(3), 353–377.

Nebel, B.: 1992, 'Syntax based approaches to belief revision'. In: P. Gärdenfors (ed.): *Belief Revision*, Theoretical Computer Science. Cambridge University Press, pp. 52–88.

Poole, D.: 1988, 'A logical framework for default reasoning'. *Artificial Intelligence* **36**, 27–47.

Reiter, R.: 1980, 'A Logic for Default Reasoning'. *Artificial Intelligence* **13**, 81–132.

Rott, H.: 1996, *Change, Choice and Inference*. Oxford University Press.

Spohn, W.: 1988, 'Ordinal Conditional Functions: A Dynamic Theory of Epistemic States'. In: W. L. Harper and B. Skyrms (eds.): *Causation in Decision: Belief, Change and Statistics: Proceedings of the Irvine Conference on Probability and Causation: Volume II*, Vol. 42 of *The University of Western Ontario Series in Philosophy of Science*. Dordrecht, pp. 105–134, Kluwer Academic Publishers.