Stringent non-monotonic inference: An alternative approach to expectation-based inference relations

Abhaya Nayak
Department of Computing
Macquarie University
NSW 2109, Australia
abhaya@ics.mq.edu.au

Samir Chopra
Knowledge Systems Group
University of New South Wales
NSW 2052, Australia
schopra@cse.unsw.edu.au

1 Introduction

One of the major factors that influences the reasoning of an agent is its currently held beliefs, and the strengths thereof. Consider how an agent should reason were it to receive a piece of veridical information, namely, "Tweety does not fly". If this piece of information is consistent with everything that the agent currently believes, then it should simply incorporate this new piece of information into its belief corpus (and its inference engine - presumably classical - will take the responsibility of generating the inferences from this enhanced belief corpus). If the piece of information in question is inconsistent with its current beliefs as a whole, so the received theory goes, the agent should identify a set of beliefs to jettison – and ideally this set should be chosen so as to minimise the ensuing loss of information. While there is overall agreement on this proposal, the question of what constitutes "minimal loss of information" has remained a live research issue.

This idea of minimising loss of information in face of recalcitrant evidence is intimately connected with patterns of nonmonotonic inference. In fact, nonmonotonic inference and the theory of belief change (that typically employs minimal loss of information) have been viewed as different perspectives on the same underlying logic [MG90, GM94]: the nonmonotonic (plausible) inference $\alpha \mid\!\sim \beta$ (β is a plausible conclusion from α) can be taken to mean that accepting α as a new piece of information results in our adopting β as a belief [GM94]. Accordingly, the principle of minimal change has played a crucial role in approaches to nonmonotonic inference.

In this paper we examine a particular approach to unpack the intent behind the notion of "minimal loss of information" that has been adopted and made popular by Gärdenfors and Makinson (GM) [GM94]. This approach is based on an ordering over the "expectations" of an reasoning agent. In this approach, a reasoner infers β from α exactly when β is "more expected in light of" α than its negation $\neg\beta$. We argue that this approach to plausible inference is rather permissive in sanctioning inferences that are counterintuitive. We instead suggest that a reasoner should infer β from α only when β is *sufficiently* more expected in light of α than its negation $\neg\beta$. (Explication of this notion of "sufficiency" is postponed until Section 3 until the formal machinery is set up.)

Having suggested this new notion of "minimal loss of information", we investigate the formal properties of the corresponding inference operation in Section 3.1, and that of the corresponding belief change model in Section 4. In the process we notice three interesting points of departure from the GM approach.

1. The notion of "minimal loss of information" used in the AGM approach involves only the strength of an agent's currently held beliefs. The notion we suggest involves a more encompassing measure – the degree of beliefs is relevant in determining what inferences a piece of evidence does license, but so is the plausibility of other entertainable hypotheses that are not exactly accepted by the agent. Our proposal is vindicated by the practice in the Bayesian tradition which concerns itself with the non-trivial comparison of plausible hypotheses.

2. The principle of minimal loss of information is often equated with the Principle of Minimal Change. This latter principle is often captured by the Principle of Choosing the Best, where the underlying preference relation reflects the plausibility of epistemic alternatives (propositions, worlds, scenarios, etc.). Intuitively, Choosing the Best epistemic alternatives entails choosing the most plausible alternatives thus minimising the loss of informa-

tion. However, it turns out that the model of belief change that ensues from the approach advocated in this paper corresponds to the dual principle of "Rejecting the Worst".

3. An interesting spin off of aspect (1) of our proposal is that it allows for abductive inferences. Of course, like the traditional AGM approach, our approach also provides for revision based inferences. Hence, our proposal provides a more comprehensive approach to inference than pure revision based inferences or pure abductive inferences in that in accommodates both of them as special cases of a general inference procedure. Furthermore, a tangible consequence of factor (2) above is that it provides room for *iterated abductive inference*.

2 Informational Economy and Non-monotonic Inference

The Principle of Informational Economy says that information is a valuable commodity, and in any reasoning task that involves trade-off between information and veracity, loss of information should be minimised. There is general agreement on this principle, although there is room for disagreement on what constitutes minimal loss of information. In the context of belief revision equating minimisation of loss of information with the minimisation of the number of beliefs that are lost leads to what is known as Maxichoice Revision which has rather unpleasant consequences[1]. In order to avoid this, it has been suggested that instead of retaining as many of the old beliefs as possible, an agent should only retain those beliefs that are most appropriate – believe in those propositions that are most expected – "in light" of the new information. To this end, additional structure over the set of possible beliefs is required. One such structuring is provided by the *expectations orderings* over the agent's language. Accordingly we will assume a finite propositional language L equipped with the standard Boolean connectives $(\neg, \wedge, \vee, \rightarrow)$, a classical consequence relation Cn (monotone, compact and idempotent) and the usual constants $\top, \bot$. Greek lowercase letters $\alpha, \beta \ldots$ will denote arbitrary formulae; Γ a set of sentences and Roman letters $p, q \ldots$ propositional atoms. The derivability relation $\vdash$ is defined as usual: $\Gamma \vdash \alpha$ iff $\alpha \in Cn(\Gamma)$. The set of all formulas in the language will be denoted $\mathcal{L}$.

2.1 Expectation Orderings

Following [GM94] we assume an expectations ordering $\leq$ over $\mathcal{L}$ satisfies the following properties:

- If $\alpha \leq \beta$ and $\beta \leq \gamma$, then $\alpha \leq \gamma$ (Transitivity)
- If $\alpha \vdash \beta$, then $\alpha \leq \beta$. (Dominance)
- Either $\alpha \leq \alpha \wedge \beta$ or $\beta \leq \alpha \wedge \beta$ for all sentences α, β. (Conjunctiveness)

The conditions above imply in addition *reflexivity* ($\alpha \leq \alpha$ for every sentence α) and *connectivity* (either $\alpha \leq \beta$ or $\beta \leq \alpha$ for every sentence α and β). Hence an expectations ordering $\leq$ is a total preorder over the sentences in $\mathcal{L}$. $\alpha \leq \beta$ is intended to mean that "α is at most as expected as β". As usual, $\alpha < \beta$ is shorthand for "$\alpha \leq \beta$ and $\beta \not\leq \alpha$".

The agent's expectations are all defeasible. As the following example shows, this defeasibility of expectations is carried over to non-monotonic inference.

Example 1 *Consider a language L with the predicates S_x: x is a Swedish citizen, E_x: x has English parents, and P_x: x is a Protestant. We would expect the Swedes but not the English to be Protestants. Hence the set of expectations includes $S_x \rightarrow P_x$ and $S_x \wedge E_x \rightarrow \neg P_x$ for all individuals x. Obviously, what logically follows from expectations is also an expectation. Hence this set also contains $S_x \rightarrow \neg E_x$ and $S_x \wedge E_x \rightarrow S_x$. Now imagine that the agent learns S_a – that some individual a is Swedish. This information will obviously lead her to conclude P_a whereby $S_a \mid\sim P_a$. However, if she were to learn $S_a \wedge E_a$, she would need to resolve some conflicts before making any rational inference. In particular, while one expectation ($S_x \rightarrow P_x$) licences the inference of P_a, the other expectation ($S_x \wedge E_x \rightarrow \neg P_x$) licences the inference of $\neg P_a$. And these two potential inferences are in conflict with each other. The Principle of Specificity suggests that our agent drop the expectations $S_x \rightarrow P_x$ and $S_x \rightarrow \neg E_x$. Consequently the agent can non-monotonically infer $\neg P_a$ from $S_a \wedge E_a$.*

[1] Essentially, an agent becomes extremely opinionated after performng any non-trivial maxi-choice revision - for every sentence, the agent either believes or disbelieves it.

Therefore, not all of an agent's expectations can be used in inferences based on information that is inconsistent with the set of expectations. In this example, the Principle of Specificity was appealed to to determine which of the alternative strategies to remove expectations. Furthermore, 'as many of the older expectations as possible' that are compatible with the new information ($S_a \wedge E_a$) were retained in the resultant contracted set. In general, when constructing the contracted set of expectations, the agent is to give up just those expectations so that the contracted set of expectations and its logical consequences contain as many as possible of the sentences in the set of expectations that are compatible with the new information. Thus we see the principle of minimal change at work.

Our appeal to the Principle of Specificity in the above example may appear to be rather *ad hoc*. This principle may not always be relevant, or appropriate. In the context of expectations, it is assumed that all such contextual information is encoded in the expectation ordering itself. Hence in order to resolve any conflict, instead of appealing to extraneous principles such as specificity, we can exploit the information implicit in the expectation ordering itself. Gärdenfors and Makinson offer the following suggestion: α non-monotonically entails β just in case β logically follows from α together with all those propositions that are 'expected in light of' α (a similar application is made in [Rot91]). This informal notion of *'expected in light of'* α is made more concrete via the following condition (here, $\mid\sim$ is termed a *comparative expectation inference relation*):

Definition 1 $(C_{\mid\sim}) : \alpha \mid\sim \gamma$ *iff* $\gamma \in Cn(\{\alpha\} \cup \{\beta : \neg\alpha < \beta\}$.

Intuitively, for a proposition to be expected with respect to the new evidence, all that is required is that the proposition in question should be more expected than the negation of the evidence. But why is it sufficient on part of a proposition β to be "expected in light of α" to be simply more expected than $\neg\alpha$ rather than α itself? The answer lies in a technicality; but this approach can be alternatively motivated by reformulating the inference operation:

Definition 2 $(C_{\mid\sim}) : \alpha \mid\sim \gamma$ *iff* $\alpha \to \neg\gamma < \alpha \to \gamma$.

If we paraphrase $x \to y$ as the information in y in light of (with respect to) x, then the above says that γ should be inferred from α just in case γ is more expected than $\neg\gamma$ in light of α . For illustration, we consider a variation of the earlier example.

Example 2 *Let S, P and E denote the atoms of L, meaning S is Swedish, P is a Protestant and E is English. The expectation ordering is $\bot < \neg S < S \to (P \wedge \neg E) < (S \wedge E) \to \neg P < \top$ with the understanding that*

- *Tautologies are most expected while non-consequences of $\neg S$ are least expected.*
- *Consequences of $(S \wedge E) \to \neg P < \top$ that are not tautologous are the next most expected.*
- *The next set of sentences in the expectation hierarchy are the consequences of $S \to (P \wedge \neg E)$ that are not also consequences of $(S \wedge E) \to \neg P$.*
- *Consequences of $\neg S$ that do not follow from $S \to (P \wedge \neg E)$ are the next to least expected.*

Note how the effect of the specificity principle is built into this expectation ordering. Now consider the consequences of S. By Definition 1, a sentence follows from S iff it is a classical consequence of $\{S, S \to (P \wedge \neg E), (S \wedge E) \to \neg P\}$. Accordingly, $S \mid\sim P$. By Definition 2, $S \wedge E \not\mid\sim P$ instead, since, while $S \wedge E \to P$ is a classical consequence of $S \to (P \wedge \neg E)$, it does not classically follow from $S \wedge E \to \neg P$ whereby $S \wedge E \to P < S \wedge E \to \neg P$.

Thus we see that an expectation ordering, properly used, can give us the result that we intuitively expect. One might wonder if an inference operation based on such an expectation ordering has the "right" formal properties. The interesting properties of this inference relation are presented in [KLM90] and discussion of relationships holding among them in terms of inter-derivability may be found in [GM94]. We present the properties below:

Definition 3 *The basic conditions for a non-monotonic inference relation are as follows:*

- *If $\alpha \vdash \gamma$, then $\alpha \mid\sim \gamma$* *(Supraclassicality)*
- *If $\vdash \alpha \leftrightarrow \beta$ and $\alpha \mid\sim \gamma$, then $\beta \mid\sim \gamma$* *(Left Logical Equivalence)*
- *If $\beta \to \gamma$ and $\alpha \mid\sim \beta$, then $\alpha \mid\sim \gamma$* *(Right Weakening)*

- *If* $\alpha \mid\!\sim \beta$ *and* $\alpha \mid\!\sim \gamma$, *then* $\alpha \mid\!\sim \beta \wedge \gamma$ *(And)*
- $\alpha \mid\!\sim \alpha$ *(Reflexivity)*
- *If* $\alpha \mid\!\sim \beta_i \forall \beta_i \in B$ *and* $B \vdash \gamma$, *then* $\alpha \mid\!\sim \gamma$ *(Closure)*
- *If* $\alpha \mid\!\sim \beta$, *then* $\mid\!\sim \alpha \rightarrow \beta$ *(Weak Conditionalization)*
- *If* $\alpha \wedge \beta \mid\!\sim \gamma$, *then* $\alpha \mid\!\sim \beta \rightarrow \gamma$ $(Conditionalization)$
- *If* $/\!\!\mid\!\sim \neg\alpha$, *and* $\mid\!\sim \alpha \rightarrow \beta$, *then* $\alpha \mid\!\sim \beta$ *(Weak Rational Monotony)*
- *If* $\alpha \mid\!\sim \perp$, *then* $\alpha \vdash \perp$ *(Consistency Preservation)*

A representation theorem is then provided:

Theorem 1 *[GM94]*

- *Let* $\leq$ *be an expectations ordering over* L. *Then the inference relation that it determines by Definition 1 satisfies the set of properties above.*
- *Let* $\mid\!\sim$ *be any inference relation on* L *that satisfies the properties above. Then* $\mid\!\sim$ *is a comparative expectation inference relation i.e., there is an expectations ordering* $\leq$ *over* L *such that* $\mid\!\sim = \mid\!\sim_{\leq}$.

As Gärdenfors and Makinson show, *epistemic entrenchment* (EE) relations and expectations orderings generate precisely the same class of inference relations, which implies that belief change based on expectation orderings adheres to the same principle of minimal change that underlies the explicit constructions for belief sets based on EE orderings[2].

Non-monotonic inference is intimately associated with belief revision; it can be shown that the comparative expectation inference relation satisfies precisely those properties that correspond to the AGM axioms of belief revision (K^*1)-(K^*8). The associated revision operator is defined in terms of expectations orderings as follows:

Definition 4 $\beta \in K^*_\alpha$ *iff* $\alpha \rightarrow \neg\beta \leq \alpha \rightarrow \beta$ *where* $\leq$ *is an epistemic entrenchment ordering over* L.

Here, $\beta \in K^*_\alpha$ is seen as a non-monotonic inference from α to β given the set K as the set of background expectations. This connection effectively closes the circle on the connections between expectations and epsitemic entrenchment orderings, between non-monotonic inference relations and revision operators: we obtain via expectation orderings precisely the same construction for belief change operators that we would using epistemic entrenchments (and equivalently using a system of spheres semantics). In sum, the GM condition $C_{\mid\!\sim}$ is a condition which captures the notion of minimal change via a condition on the expectations ordering over L. Another way to read the condition that GM place on the non-monotonic inference relation is to view it as a way of selecting the 'best' propositions in light of the new information.

2.2 A Critical Appraisal

Having defined minimal change above, it is natural to ask, what sort of minimal belief change scenario has been modeled? The guiding principle behind the condition $C \mid\!\sim$ above seems to be "for each proposition (expectation) β, check if β is more expected in "light of α" than its *negation*, $\neg\beta$, if so, accept it. However, this principle seems to be lenient and lacking in real prescriptive content. This is so because all that is being tested is whether a proposition β is *merely more expected than its negation in light of* α. This means that if there is a proposition that we are agnostic about, but simply slightly less so than its negation in light of α, then we should accept it as a full belief. This seems to be a very weak condition. The following example illustrates this point.

[2] Epistemic entrenchment relations include the following two conditions:

- $\beta \leq \alpha, \forall\beta$ only if $\vdash \alpha$ $(Maximality)$
- If $K \not\vdash \perp$, then $\alpha \notin K$ iff $\alpha \leq \beta, \forall\beta \in \mathcal{L}$ $(Minimality)$

Example 3 Tweety, the non-flying bird: *Consider the atoms* $B(t)$: Tweety is a bird*,* $F(t)$: Tweety flies*, Let* α *denote the sentence* $B(t)$ *while both* β *and* γ *denote the sentence* $\neg F(t)$*. Intuitively,* $\neg\alpha < \beta$*. Now, since* $\gamma \vdash \alpha \rightarrow \gamma$ *it follows that* $\beta \vdash \alpha \rightarrow \gamma$*. Then* $\gamma \in Cn(\alpha) \cup \{\beta\}$*. But* $\{\beta\} \subseteq \{\beta' : \neg\alpha \leq \beta'\}$*. whence the GM condition would sanction the inference* $\alpha \mid\sim \gamma$ *i.e.* $B(t) \mid\sim \neg F(t)$*, in other words, it would sanction the inference of "Tweety does not fly" from "Tweety is a bird" which runs counter to our intuition. The proposition that Tweety is not a bird is less expected than the proposition that Tweety flies, the agent believes that Tweety is a bird, and reasonably expects Tweety to fly. In the absence of any information to the contrary, the agent would be expected to non-monotonically infer that Tweety flies. Instead, since the negation of a well expected proposition is less expected than some other unreasonable proposition, the GM condition sanctions precisely the inference of that unreasonable proposition.*

This example shows that the inference operation as proposed by Gärdenfors and Makinson does not live upto our expectation since it allows inferences not sanctioned by our intuitions. There is the added concern that this operation does not distinguish in the way it treats evidence that coheres with our current information and evidence that conflicts with our current expectations. We would expect that an agent's choice of premises for an inference operation is dependent on whether the information is consistent or inconsistent with background knowledge. The proposal we present in the next section – the primary contribution of this paper – addresses these issues.

3 An Alternative Approach to Expectations

Our alternative approach to the definition of non-monotonic inference relations via expectations orderings is embedded in a general theory of epistemic updating. The AGM theory [AGM85, GM88] of belief change takes into account the degree of currently held beliefs–in assimilating new information, a rational agent should ensure that *if some currently held beliefs must be given up, then, given the option, less firmly held beliefs may be given up in favour of more firmly held beliefs.* Possibility theory [DP92], relies on the degree of *disbelief*–in assimilating new information, a rational agent may be forced to suspend disbelief in some sentences that are currently disbelieved (i.e., their negations are believed). If so, the agent should see to it that *given the option, the suspension of disbelief is carried out with respect to less strongly denounced propositions instead of more strongly denounced propositions*. Both approaches — belief change and possibility theory — are inter-translatable since the firmness of the belief in a sentence may be viewed as the strength of denouncement with respect to its negation.

Since each sentence is either believed or disbelieved or neither, given an agent's belief state, sentences of a language may be partitioned into *beliefs* (sentences that the agents takes to be true in her model of the world), *disbeliefs* (sentences that the agent takes to be false in her model of the world) and *plausibilities* (sentences that the agent is agnostic about). The measure used in belief change theory i.e., epistemic entrenchment, is primarily defined over the beliefs. The agent's beliefs are in one category and can be nontrivially compared, while the agent's disbeliefs and unentertained hypotheses are in another category – *all of them* minimally entrenched. The measure used in possibility theory is primarily defined over the disbeliefs. Both these measures effectively refuse to compare different plausibilities. However, another well-established tradition in state-updating, Bayesian theory, is primarily based on the probability of the plausibilities[3].

The Bayesian approach and the belief change (or, for that matter, possibility theory) complement each other in providing us a model for the general task of accepting some new evidence (and subsequence inference thereafter). Belief change and possibility theory primarily provide a model for accepting new information that conflicts with the current knowledge (the problem of *belief revision*). The account given of accepting new information that is not in conflict with the current knowledge is viewed as a special case not to be taken seriously. The Bayesian tradition provides us a model of how to accept evidence that is consistent with the current knowledge. This problem has come to be known as *expansion*. Bayesian doctrine is straighforward about its treatment of evidence that conflicts with the current knowledge – the doctrine is simply not designed to handle such evidence.

A *truly general* account of non-monotonic inference should provide a non-trivial account of handling two types of evidence – disbeliefs and plausibilities – in the sense that it should be based on a measure that allows non-trivial comparison among beliefs (or disbeliefs) *and* among plausibilities. This condition on the non-monotonic inference relation should have a mixed flavor, reflecting the different responses of the agent to information that is consistent or inconsistent with its beliefs. The inference relation should allow the non-trivial comparison of beliefs (*mutatis*

[3] In the Bayesian framework, each belief receives probability 1 and each disbelief gets probability 0. So there is no non-trivial comparison among beliefs (or disbeliefs). Only the comparison among the plausibilities is nontrivial.

mutandis disbeliefs) on the basis of their strength or firmness, and of hypotheses that have neither been accepted nor rejected on the basis of their plausibility, Such an inference relation would truly represent a genuine alternative to current notions of non-monotonic inference in its combination of both Bayesian and non-Bayesian approaches in a single framework.

Our take on the current problem is via an alternate ranking of the expectations ordering on the agent's language (Figure 1 below). Note that now there are three categories of sentences. The maximally expected (entrenched) propositions are tautologies as usual. Below these are the agent's beliefs. Then there are serious possibilities (a la Levi) that are not currently believed but are certainly entertainable/plausible. Their negations are thereby not entertainable. There are, however, those propositions that we are strictly agnostic about – we don't dream of entertaining them or their negations. Note the crucial differences with AGM style EE's: there are propositions that are not part of our belief set and yet, they are not homogenous with respect to their ordering.

Example 4 *I reasonably expect my colleague to be in the lab on this floor. If asked, I would hesitate before saying that I believed he was in there. However, given his normal work patterns, I would think that he was in there without actually believing that he was. When it comes to religious belief, however, I am truly agnostic. I don't know what evidence would count for the existence of a Supreme Being, and I don't know what evidence would conclusively prove the non-existence of a Supreme Being. Given the difficulty in framing the question, I am agnostic about both propositions.*

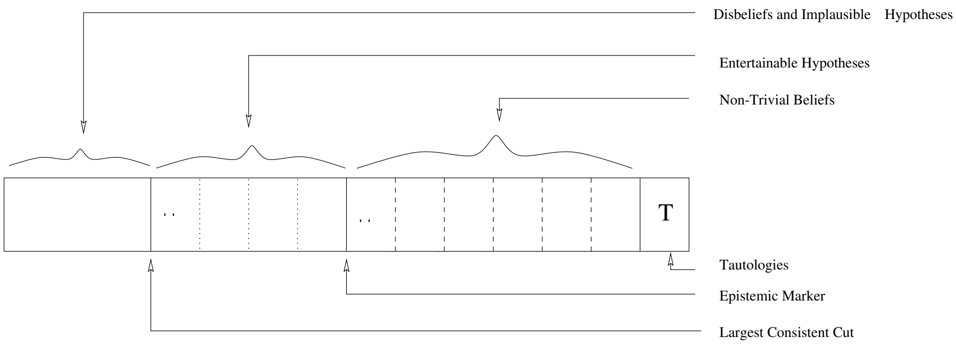

Figure 1: Alternative ranking over expectation ordering

To obtain a definition of our intended non-monotonic inference relation in terms of the expectation ordering $\leq$, we make use of the notion of a "privileged sentence" or "epistemic marker" in the agent's belief set. This epistemic marker gives us a yardstick for measuring the relative entrenchment of beliefs in the agent's belief set. We begin by reviewing the definition of a *cut* over an expectation ordering:

Definition 5 *A set $\Gamma \subset \mathcal{L}$ is a cut over an expectations ordering $\leq$ iff if $\alpha \in \Gamma$ and $\alpha \leq \beta$ then $\beta \in \Gamma$.*

The agent's belief set K corresponds to a particular cut Γ_K over the expectation ordering $\leq$ in $\mathcal{L}$ (since if Γ is a cut, $\Gamma = Cn(\Gamma)$. The "epistemic marker" is a minimal element γ (not necessarily unique) in this cut such that $\Gamma_K = \{\beta | \gamma \leq \beta\}$. Alternatively, the epistemic marker is the weakest belief in the agent's belief set. The second cut in the expectations ordering is the last "safe" cut without inviting inconsistency. Establishing the presence of the epistemic marker in the agent's belief set now lets us define an associated non-monotonic inference relation (we term this inference relation a *stringent expectation inference relation* for reasons that will become clear).

Definition 6 $\mid\!\sim$ *is a stringent expectation inference relation iff there is an expectation ordering $\leq$ over L such that the following condition holds:*
$(C_{\mid\sim})$:$\alpha \mid\!\sim \beta$ *w.r.t to an epistemic marker γ iff:*

1. $\alpha \to \neg\beta < \alpha \to \beta$ *and:*
   (a) *When $\gamma \leq \neg\alpha$, if $\alpha \to \beta < \alpha \to \beta'$ then $\top \leq \alpha \to \beta'$, for every sentence β' in L.*
   (b) *When $\neg\alpha < \gamma$ if $\alpha \to \beta < \alpha \to \beta'$, then $\gamma \leq \alpha \to \beta'$, for every sentence β' in L.*

The definition above is is easily seen to be more stringent in terms of which expectations are considered "sufficiently well expected in light of α" The expectations that we select by the above definition are of two kinds. Part (1) of the definition corresponds to the case when the new information α is inconsistent with background knowledge and says that the agent should accept those sentences that are "most" expected in the light of α. Part (2) corresponds to the case when the new information α is consistent with background knowledge and says that the agent should accept those that are "sufficiently" expected in light of α. Part (1a) and (2a) are equivalent to the GM conditions; it is a necessary condition for the relative information in β to be more expected than that in its rival $\neg\beta$ if β is to be a consequence of α. The net effect of this condition is that not both β and $\neg\beta$ are allowed to follow from α together. We call a conditional $\alpha \rightarrow \beta$ *eligible* if it satisfies the condition (a). Condition (1b) adds the extra condition that the relative information in β wrt α be nontrivially most expected. Another way to motivate (1) is as follows: given the premise α, only conditionals of the form $\alpha \rightarrow \ldots$ are relevant, since only they can be used to infer new information. Now, if a conditional $\alpha \rightarrow \beta$ is tautologous, it will be used for this purpose regardless. Among the nontautologous but relevant conditionals, only those will be used for the inferential purpose that are maximally expected, and satisfy the eligibility criteria. In the other case (2),the situation is analogous, but slightly different. Note that in this case the premise α is consistent with the current knowledge. So, if $\mid\!\sim$ is to be ampliative we need to use relevant conditionals above and beyond those that are currently accepted. These extra ones are those that are maximally expected among the relevant eligible conditionals not already accepted as a part of one's body of beliefs.

We believe our proposal to be closer to the spirit of the original project of devising a non-monotonic inference relation in terms of expectations which was to use premises drawn from not just beliefs but also the "expectations that guide our beliefs without quite being part of them". In fact though, what happens in the GM model is that those expectations that are not considered in the choice of premises for an inference are simply the non-minimal elements of an epistemic entrenchment ordering. But these are disbliefs in any case. So, in order to really do justice to the idea that these are defeasible expectations our ranking must give us a means for truly distinguishing between beliefs and entertainable hypotheses.

Earlier we complained that the GM approach allows inferences that are not sanctioned by our pretheoretic intuition, namely that from the premise *Tweety is a bird* it allows us to infer that *Tweety does not fly*. The example below shows that under the inference based on our stringent expectation ordering, this inference is not allowed no matter whether the evidence *Tweety is a bird* conflicts with our current knowledge or coheres with it.

Example 5 Tweety the flying bird: the inconsistent case *By condition $C_{\mid\!\sim}$ above, $\alpha \mid\!\sim \beta$ iff the relative info in β wrt α is maximally expected (except the trivial consequences of α). $C_{\mid\!\sim}$ will be satisfied trivially if $\alpha \rightarrow \beta \not< \gamma, \forall\gamma$. For our purpose, we need to find a β' such that $\alpha \rightarrow \beta'$, but $\alpha \not\vdash \beta'$. We now set β' to $\alpha \rightarrow \neg\beta$. Since $\alpha \rightarrow \beta < \alpha \rightarrow \neg\beta$ and $\alpha \not\vdash \alpha \rightarrow \neg\beta$ we have that $\alpha \not\mid\!\sim \beta$ as desired.*
The consistent case: *We want to find a β' such that $\alpha \rightarrow \beta < \beta'$, but $x \not< \alpha \rightarrow \beta'$. The previous (revision) case does not work since we set β' to $\alpha \rightarrow \neg\beta$ so $\alpha \rightarrow \beta' = \neg\alpha \vee \neg\alpha \vee \neg\beta = \alpha \rightarrow \neg\beta$ which is presumably a relatively strong belief. We set $\beta' =$* Tweety flies *$= F(t)$. Then, it is plausible that $B(t) \rightarrow \neg F(t) < F(t)$. So, now $B(t) \rightarrow \neg F(t) < M(t)$ and therefore $x \not< B(t) \rightarrow M(t)$.*

However, we need to convince ourselves that we are not throwing the baby out with the bath water – that we *really* sanction desirable inferences. This contention is supported by the following example.

Example 6 Fredo, the venomous. *This example is designed to show that we are not blocking everything. (We are tacitly assuming that the domain in this example is the animal kingdom.) Suppose that we don't know if Fredo is a frog, and we do not know if Fredo is venomous. However we believe that it is rather plain (not colourful). Consider what we can infer from the premise $colourful(fredo)$. The nontrivial, relevant conditionals are: $colourful(fredo) \rightarrow frog(fredo)$, $colourful(fredo) \rightarrow \neg frog(fredo)$, $colourful(fredo) \rightarrow venomous(fredo)$ and $colourful(fredo) \rightarrow \neg venomous(fredo)$. Of these, we believe the conditional $colourful(fredo) \rightarrow venomous(fredo)$*[4] *(and possibly $colourful(fredo) \rightarrow \neg frog(fredo)$ as well) is eligible. Among these, $colourful(fredo) \rightarrow venomous(fredo)$ is definitely a maximal member, so we can infer that Fredo is venomous from the input that Fredo is colourful. Slightly altering this example, we can demonstrate that $colourful(fredo) \mid\!\sim venomous(fredo)$ even when $colourful(fredo)$ is consistent with our current knowledge.*

[4]There is the theory that venomous speces don't have the need to blend with the background for their survival, so they tend to be very colourful and attractive.

3.1 Properties of Stringent Expectation Inference Relations

The properties of this inference relation are given by the following results.

Theorem 2 *Let $\leq$ be an expectation ordering over L. Then the inference relation $\mid\!\sim_{\leq}$ that it determines by $(C_{\mid\!\sim})$ satisfies the set of properties given in Definition 3 (except for Weak Conditionalization and Conditionalization) and Definition 8. 8*

Theorem 3 *Let $\mid\!\sim$ be an inference relation that satisfies the properties given in Definition 3 (except for Weak conditionalization and Conditionalization) and Definition 8. Then $\mid\!\sim$ is a stringent expectation inference relation i.e., there is an expectations ordering $\leq$ over L such that $\mid\!\sim = \mid\!\sim_{\leq}$ (as above).*

Note that our inference relation $\mid\!\sim$ does not satisfy the conditionalization property. However, several weakenings of the Conditionalization postulate obtain. We list these below, but first, we establish some interpretations of the non-monotonic inference relation just defined. We read $\mid\!\sim \alpha$ as "α is evident" (or, $\neg\alpha$ is not entertainable), $\not\mid\!\sim \alpha$ as "α is not evident" (or, $\neg\alpha$ is entertainable), $\not\mid\!\sim \neg\alpha$ as "$\neg\alpha$ is not evident" (or, α is entertainable), $\mid\!\sim \neg\alpha$ as "$\neg\alpha$ is evident" (or α is not entertainable) and $\alpha \mid\!\sim \beta$: α supports β. These interpretations of the inference relation make clear the relationship between the agent's set of expectations and the inferences sanctioned by them. We also define the notions of *knowledge dependent* and *knowledge independent* inferences.

Definition 7

1. *Let $\alpha \mid\!\sim \beta$. We will say that this inference of β from α is knowledge independent iff $\alpha \vdash \beta$. Otherwise, the inference of β from α is knowledge dependent.*
2. *Let α be such that no matter what β is, $\alpha \mid\!\sim \beta$ only if $\alpha \vdash \beta$. (Thus $\alpha \mid\!\sim \beta$ iff $\alpha \vdash \beta$.) In this case, since the background knowledge does not add anything to the inferencial potency of α, we will call α a knowledge independent hypothesis.*

We now list the additional properties, along with intuitive readings.

Definition 8 *The following are additional properties of the inference relation $\mid\!\sim$:*

- *(6A$\circ$) $\mid\!\sim \alpha \leftrightarrow \beta$, $\not\mid\!\sim \neg\alpha \Rightarrow \forall\gamma, \alpha \mid\!\sim \gamma$ iff $\beta \mid\!\sim \gamma$*

  If α is an entertainable hypothesis, and both α, β are simultaneously entertainable or not relative to background knowledge, then accepting them has exactly the same effect on the agent's knowledge (in terms of the inferences that their inclusion permits).

- *(7.1R$\circ$) If $\mid\!\sim \neg\alpha$ and $\exists\gamma$ such that $\alpha \mid\!\sim \gamma$, and $\alpha \wedge \beta \not\vdash \gamma$ then $\alpha \wedge \beta \mid\!\sim \gamma' \Rightarrow \alpha \mid\!\sim \beta \rightarrow \gamma'$.*

  If α is not entertainable, but supports some γ that cannot be classically deduced from current knowledge even when strengthened by $\alpha \wedge \beta$, then every γ' supported by $\alpha \wedge \beta$ is also conditionally supported by α given β.

- *(7.1A$\circ$) If $\exists\gamma$ such that $\alpha \mid\!\sim \gamma$, and $\not\mid\!\sim \alpha \wedge \beta \rightarrow \gamma$, then $\alpha \wedge \beta \mid\!\sim \gamma' \Rightarrow \alpha \mid\!\sim \beta \rightarrow \gamma', \forall\gamma'$.*

  If α supports some γ that cannot be classically deduced from K even if strengthened by α and β, then every γ' that is supported by $\alpha \wedge \beta$ is also conditionally supported by α given β.

- *(7.2R$\circ$) If $\mid\!\sim \neg\alpha$ and $\beta \mid\!\sim \gamma \Rightarrow \beta \vdash \gamma$ then $\alpha \wedge \beta \mid\!\sim \gamma' \Rightarrow \alpha \mid\!\sim \beta \rightarrow \gamma'$.*

  Let β be a knowledge independent hypothesis and α not entertainable. Then $\alpha \wedge \beta$ supports any sentence γ' only if α conditionally supports γ' (wrt β).

- *(7.2A$\circ$) If $\not\mid\!\sim \neg\alpha$ and $\beta \mid\!\sim \gamma \Rightarrow \mid\!\sim \beta \rightarrow \gamma$ then $\alpha \wedge \beta \mid\!\sim \gamma' \Rightarrow \alpha \mid\!\sim \beta \rightarrow \gamma'$.*

  If α is entertainable, and if β supports an inference γ such that γ is conditionally supported by β given background knowledge, then every γ' that is supported by $\alpha \wedge \beta$ is also conditionally supported by α given β.

- *(7.3R$\circ$) If $\alpha \mid\!\sim \gamma$ and $\beta \mid\!\sim \alpha \Rightarrow \alpha \vdash \gamma$ then $\alpha \wedge \beta \mid\!\sim \gamma' \Rightarrow \alpha \mid\!\sim \beta \rightarrow \gamma'$.*

  If α's supporting γ, and β's supporting γ entails that α is a knowledge independent hypothesis, then every γ' that is supported by $\alpha \wedge \beta$ is also conditionally supported by α given β.

- *(7.3A$\circ$) If $\not\mid\!\sim \neg\alpha$ and $\alpha \mid\!\sim \gamma$, $\mid\!\sim \beta \to \gamma \Rightarrow\ \mid\!\sim \alpha \to \gamma$ then $\alpha \wedge \beta \mid\!\sim \gamma' \Rightarrow \alpha \mid\!\sim \beta \to \gamma'$.*

  If α is entertainable, and supports γ, and if the supporting of the conditional inference $\beta \to \gamma$ entails that the conditional inference $\alpha \to \gamma$ is supported relative to background knowledge, then every γ' that is supported by $\alpha \wedge \beta$ is also conditionally supported by α given β.

- *(8$\circ$) $\alpha \not\mid\!\sim \neg\beta \Rightarrow$ if $\alpha \mid\!\sim \beta \to \gamma$, then $\alpha \wedge \beta \mid\!\sim \gamma$* *(Rational Monotony)*

- *(9R$\circ$) $\mid\!\sim \neg\alpha, \alpha \mid\!\sim \neg\beta$ but $\alpha \not\vdash \neg\beta \Rightarrow \alpha \wedge \beta \mid\!\sim \gamma$ iff $\alpha \wedge \beta \vdash \gamma$*

  If α is not entertainable, and the inference of $\neg\beta$ from background knowledge supported by α is a knowledge dependent inference, then $\alpha \wedge \beta$ is a knowledge independent hypothesis.

- *(9A$\circ$) $\not\mid\!\sim \alpha \to \neg\beta$, $\alpha \mid\!\sim \neg\beta \Rightarrow$ if $\alpha \wedge \beta \mid\!\sim \gamma$ then $\mid\!\sim \alpha \wedge \beta \to \gamma$*

  If $\alpha \to \neg\beta$ is not evident, but $\neg\beta$ is supported by background knowledge strengthened by α, then every γ supported by $\alpha \wedge \beta$ is also conditionally supported by $\alpha \wedge \beta$ relative to background knowledge.

3.2 Classification in KLM Taxonomy

Wheer does the inference relation just devised fit into the hierarchy of systems of non-monotonic inference?

4 An Alternative Operator For Belief Change

Given the natural relationship between non-monotonic inference relations and belief revision operators (a statement of the form $\beta \in K * \alpha$ is seen as non-monotonic inference from α to β given K as a set of *background* expectations (explicit translations of the AGM axioms into corresponding properties for non-monotonic inference are provided in [MG90]), we expect a corresponding belief revision operator to our inference relation , one that accounts for the varying responses of the agent to new information depending on its consistency with older information. The revision operator so devised works as both an abduction operator and AGM revision operator (not precisely though). Note the naming conventions followed: the numeric part of the name in general signifies which AGM postulate it is an analogue of and the (optional) alphabetic part signifies whether this property concerns the abductive behaviour or the revision behaviour of the operator $\circ$; $(7.1A\circ)$ corresponds to the AGM postulate $(7*)$ and concerns the abductive behaviour of $\circ$.

$(1\odot)$	$K^{\odot}_{\alpha}$ is a theory
$(2\odot)$	$\alpha \in K^{\odot}_{\alpha}$
$(4\odot)$	If $K \not\vdash \neg\alpha$ then $Cn(K \cup \{\alpha\}) \subseteq K^{\odot}_{\alpha}$
$(5\odot)$	$K^{\odot}_{\alpha} = K_{\perp}$ iff $\vdash \neg\alpha$
$(6R\odot)$	If $\vdash \alpha \leftrightarrow \beta$, then $K^{\odot}_{\alpha} = K^{\odot}_{\beta}$
$(6A\odot)$	If $K \vdash \alpha \leftrightarrow \beta$, then $K^{\odot}_{\alpha} = K^{\odot}_{\beta}$, given that $K \not\vdash \neg\alpha$
$(7.1R\odot)$	If $K^{\odot}_{\alpha} \not\subseteq Cn(\alpha \wedge \beta)$ then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K^{\odot}_{\alpha} \cup \{\beta\})$ given $K \vdash \neg\alpha$
$(7.1A\odot)$	If $K^{\odot}_{\alpha} \not\subseteq Cn(K \cup \{\alpha, \beta\})$ then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K^{\odot}_{\alpha} \cup \{\beta\})$
$(7.2R\odot)$	If $K^{\odot}_{\beta} = Cn(\beta)$ then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K^{\odot}_{\alpha} \cup \{\beta\})$ given $K \vdash \neg\alpha$
$(7.2A\odot)$	If $K^{\odot}_{\beta} = Cn(K \cup \{\beta\})$ then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K^{\odot}_{\alpha} \cup \{\beta\})$ given $K \not\vdash \neg\alpha$
$(7.3R\odot)$	If $K^{\odot}_{\alpha} \cap Cn(\beta) \subseteq Cn(\alpha)$ then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K^{\odot}_{\alpha} \cup \{\beta\})$
$(7.3A\odot)$	If $K^{\odot}_{\alpha} \cap Cn(K \cup \{\beta\}) \subseteq Cn(K \cup \{\alpha\})$ then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K^{\odot}_{\alpha} \cup \{\beta\})$ given $K\neg \vdash \neg\alpha$
$(8\odot)$	If $K^{\odot}_{\alpha} \not\vdash \neg\beta$ then $Cn(K^{\odot}_{\alpha} \cup \{\beta\}) \subseteq K^{\odot}_{\alpha\wedge\beta}$
$(9R\odot)$	If $K \vdash \neg\alpha,\ K^{\odot}_{\alpha} \vdash \neg\beta$ but $\alpha \not\vdash \neg\beta$ then $K^{\odot}_{\alpha\wedge\beta} = Cn(\alpha \wedge \beta)$.

$(9A\odot)$ If $K \cup \{\alpha\} \not\vdash \neg\beta$ but $K^{\odot}_{\alpha} \vdash \neg\beta$
then $K^{\odot}_{\alpha\wedge\beta} \subseteq Cn(K \cup \{\alpha, \beta\})$.

It is helpful to view $K^{\odot}_{\alpha}$ as the set of sentences that the evidence α can explain given the background knowledge K. Properties $(1\odot$–$6R\odot)$ are basic postulates of the AGM revision operation [AGM85]. Postulate $(6A\odot)$ says that if two pieces of evidence contain the same information relative to, and they do not conflict with, the current knowledge, then accepting them has the same effect on the current knowledge (this is a stronger postulate than $(6R\odot)$). Postulates $(7.1R\odot$–$7.3A\odot)$ are several variations (weakenings) of the AGM postulate $(7*)$. We defer extended discussions of these properties to the full version this paper but point out that their intuitive readings are provided in Definition 8.

Theorem 4 *Let $\odot$ be a revision operator that satisfies $(K^{\odot}1)$ - $(K^{\odot}9)$. Then for any K there is an expectations ordering $<$ satisfying $(C_{<,\odot})$.*

Theorem 5 *Let $<$ be an expectations ordering on L. Then there exists a revision operator satisfying $(C_{<,\odot})$ and axioms $(K^{\odot}1)$ - $(K^{\odot}9)$.*

5 Conclusion

In this paper, we argued that although it is taken for granted that there can be no meaningful comparison among tenable hypotheses that have neither been accepted nor rejected, a case can be made for nontrivial comparison among them on the basis of their plausibility. Equipped with a measure that can compare such hypotheses as well as beliefs and disbeliefs, we modelled a comprehensive an alternative non-monotonic inference relations and belief change operator. Much work remains to be done in exploring the properties and in devising a comprehensive theory of updating that includes the best of belief change, possibility theory and Bayesian approaches.

References

[AGM85] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[DP92] Didier Dubois and Henri Prade. Belief change and possibility theory. In Peter Gärdenfors, editor, *Belief Revision*, pages 142–182. Cambridge University Press, 1992.

[GM88] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspect of Reasoning About Knowledge*, pages 83–96, 1988.

[GM94] Peter Gärdenfors and David Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65:197–245, 1994.

[KLM90] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence*, 44:167–207, 1990.

[MG90] David Makinson and Peter Gärdenfors. Relations between the logic of theory change and nonmonotonic logic. In A. Fuhrmann and M. Morreau, editors, *The Logic Of Theory Change*, pages 185–205. Lecture Notes in Artificial Intelligence 465, Springer Verlag, 1990.

[Rot91] Hans Rott. Two methods of constructing contractions and revisions of knowledge systems. *Journal of Philosophical Logic*, 20:149–173, 1991.