ECAI 2008
M. Ghallab et al. (Eds.)
IOS Press, 2008

doi:10.3233/978-1-58603-891-5-85

Conflicts between Relevance-Sensitive and Iterated Belief Revision

Pavlos Peppas and **Anastasios Michael Fotinopoulos** and **Stella Seremetaki**[1]

Abstract. The original AGM paradigm focuses only on one-step belief revision and leaves open the problem of revising a belief state with whole sequences of evidence. Darwiche and Pearl later addressed this problem by introducing extra (intuitive) postulates as a supplement to the AGM ones. A second shortcoming of the AGM paradigm, seemingly unrelated to iterated revision, is that it is too liberal in its treatment of the notion of relevance. Once again this problem was addressed with the introduction of an extra (also very intuitive) postulate by Parikh. The main result of this paper is that Parikh postulate for relevance-sensitive belief revision is inconsistent with each of the Darwiche and Pearl postulates for iterated belief revision.

1 INTRODUCTION

The original AGM paradigm for belief revision, [1, 3, 11], focuses only on one-step transitions leaving open the problem of how to revise a belief state with a whole sequence of evidence. This problem was later addressed by Darwiche and Pearl who formulated four intuitive new postulates (known as the DP postulates) to regulate iterated revisions. Possible world semantics were introduced to characterize the new postulates, and with some adjustments (see section 4) the DP postulates were shown to be compatible with the original AGM ones.[2]

Although Darwiche and Pearl's work has received some criticism, it remains very influential in the literature of iterated belief revision and has served as a basis for further developments in the area [7, 5].

A shortcoming of a different nature of the original AGM paradigm is that it neglects the important role of *relevance* in belief revision. As noted by Parikh, [8], when a belief state ψ is revised by new information μ, only the part of ψ that is related to μ should be effected; the rest of ψ should remain the same. Parikh proceeded to formulate a postulate, called (P), that captures this intuition (albeit in limited cases). Postulate (P) was later shown to be consistent with the AGM postulates and possible-world semantics were introduced to characterize it, [10].

The main contribution of this paper is to show that, in the presence of the AGM postulates, Parikh postulate for relevance-sensitive belief revision is inconsistent with *each* of the (seemingly unrelated) Darwiche and Pearl postulates for iterated belief revision.

This of course is quite disturbing. Both the concept of relevance and the process of iteration are key notions in Belief Revision and we can do away with neither. Moreover, the encoding of these notions proposed by Parikh, Darwiche, and Pearl appears quite natural and it is not obvious how one should massage the postulates in order to reconcile them. Further to this point, subsequent postulates introduced to remedy problems with the (DP) ones, [7, 5], are also shown to be incompatible with postulate (P) (see section 6). On the positive side, these incompatibility results reveal a hitherto unknown connection between relevance and iteration that deepens our understanding of the belief revision process.

The paper is structured as follows. The next section introduces some notation and terminology. Following that we briefly review (Katsuno and Mendelzon's re-formalization of) the AGM postulates, Darwiche and Pearl's approach for iterated revisions, and Parikh's proposal for relevance-sensitive belief revision (sections 3, 4, and 5). Section 6 contains our main incompatibility results. Finally in section 7 we make some concluding remarks.

2 PRELIMINARIES

Throughout this paper we shall be working with a finitary propositional language L. We shall denote the (finite) set of all propositional variables of L by A. For a set of sentences Γ of L, we denote by $Cn(\Gamma)$ the set of all logical consequences of Γ, i.e. $Cn(\Gamma) = \{\varphi \in L\text{: } \Gamma \vdash \varphi\}$. A theory K of L is any set of sentences of L closed under $\vdash$, i.e. $K = Cn(K)$. We shall denote the set of all theories of L by $\mathcal{T}_{\mathcal{L}}$. A theory K of L is complete iff for all sentences $\varphi \in L$, $\varphi \in K$ or $\neg\varphi \in K$. As it is customary in Belief Revision, herein we shall identify consistent complete theories with possible worlds. We shall denote the set of all consistent complete theories of L by $\mathcal{M}_{\mathcal{L}}$. If for a sentence φ, $Cn(\varphi)$ is complete, we shall also call φ complete. For a set of sentences Γ of L, $[\Gamma]$ denotes the set of all consistent complete theories of L that contain Γ. Often we shall use the notation $[\varphi]$ for a sentence $\varphi \in L$, as an abbreviation of $[\{\varphi\}]$. When two sentences φ and χ are logically equivalent we shall often write $\varphi \equiv \chi$ as an abbreviation of $\vdash \varphi \leftrightarrow \chi$. Finally, the symbols $\top$ and $\bot$ will be used to denote an arbitrary (but fixed) tautology and contradiction of L respectively.

3 THE KM POSTULATES

In the AGM paradigm belief sets are represented as logical theories, new evidence as sentences of L, and the process of belief revision is modeled as a function mapping a theory K and a sentence μ to a new theory $K * \mu$. Moreover, eight postulates for $*$ are proposed, known as the AGM postulates, that aim to capture the notion of *rationality* in belief revision.

Katsuno and Mendelzon in [6] slightly reshaped the AGM constituents to make the formalization more amendable to implementa-

[1] Dept of Business Administration, University of Patras, Patras 26500, Greece, emails: pavlos@upatras.gr, anfotino@upatras.gr, serstel@master.math.upatras.gr
[2] To be precise, the DP postulates were shown to be compatible with the re-formalization of the AGM postulates introduced by Katsuno and Mendelzon in [6].

tion. In particular, the object language L is set to be a finitary propositional one,[3] and belief sets are defined as *finite* sets of sentences of L. Since a belief set contains only finitely many elements, one can in fact represent it as a single sentence ψ; namely the conjunction of all its elements. This is the representation eventually adopted in [6] and the one we shall use herein. To emphasize the finiteness of the new representation we shall call ψ a *belief base* and reserve the term *belief set* for the closure of ψ, i.e. the theory $Cn(\psi)$.

With the above reformulation, a revision function becomes a function $*$ mapping a sentence ψ and a sentence μ to a new sentence $\psi * \mu$; i.e. $* : L \times L \mapsto L$. Moreover in the new formalization the AGM postulates are equivalent to the following six, known as the KM postulates:

(KM1) $\psi * \mu \vdash \mu$.
(KM2) If $\psi \wedge \mu$ is satisfiable then $\psi * \mu \equiv \psi \wedge \mu$.
(KM3) If μ is satisfiable then $\psi * \mu$ is also satisfiable.
(KM4) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$ then $\psi_1 * \mu_1 \equiv \psi_2 * \mu_2$.
(KM5) $(\psi * \mu) \wedge \varphi \vdash \psi * (\mu \wedge \varphi)$.
(KM6) If $(\psi * \mu) \wedge \varphi$ is satisfiable then $\psi * (\mu \wedge \varphi) \vdash (\psi * \mu) \wedge \varphi$.

4 ITERATED BELIEF REVISION

One thing to notice about the KM postulates is that they all refer to single-step revisions; no constraints are placed on how the revision policy at the initial belief base ψ may relate to the revision policies at it descendants (i.e. at the belief bases resulting from ψ via a sequence of revisions).

Darwiche and Pearl's solution to this problem came in the form of four additional postulates, known as the DP postulates, listed below, [2]:

(DP1) If $\varphi \vdash \mu$ then $(\psi * \mu) * \varphi = \psi * \varphi$.
(DP2) If $\varphi \vdash \neg\mu$ then $(\psi * \mu) * \varphi = \psi * \varphi$.
(DP3) If $\psi * \varphi \vdash \mu$ then $(\psi * \mu) * \varphi \vdash \mu$.
(DP4) If $\psi * \varphi \not\vdash \neg\mu$ then $(\psi * \mu) * \varphi \not\vdash \neg\mu$.

The DP postulates are very intuitive and their intended interpretation, is loosely speaking as follows (see [2, 7] for details). Postulate (DP1) says that if the subsequent evidence φ is logically stronger than the initial evidence μ then φ overrides whatever changes μ may have made. (DP2) says that if two contradictory pieces of evidence arrive sequentially one after the other, it is the later that will prevail. (DP3) says that if revising ψ by φ causes μ to be accepted in the new belief base, then revising first by μ and then by φ can not possibly block the acceptance of μ. Finally, (DP4) captures the intuition that "no evidence can contribute to its own demise" [2]; if the revision of ψ by φ does not cause the acceptance of $\neg\mu$, then surely this should still be the case if ψ is first revised by μ before revised by ψ.

An initial problem with the DP postulates (more precisely with (DP2)) was that they were inconsistent with the KM postulates. However this inconsistency was not deeply rooted and was subsequently resolved. There are in fact (at least) two ways of removing it. The first, proposed by Darwiche and Pearl themselves, [2], involves substituting belief bases with *belief states*, and adjusting the KM postulates accordingly. The second, proposed by Nayak, Pagnucco, and Peppas, [7], keeps belief bases as the primary objects of change, but modifies the underlying assumptions about the nature of $*$.

[3] In the original AGM paradigm, the object language L is not necessarily finitary nor propositional. The details of L are left open and only a small number of structural constraints are assumed of L and its associated entailment relation $\vdash$ (see [3, 11]).

In particular, notice that, with the exception of (KM4), the KM postulates apply only to a *single* initial belief base ψ; no reference to other belief bases are made. Even postulate (KM4) can be weakened to comply with this policy:

(KM4)' If $\mu_1 \equiv \mu_2$ then $\psi * \mu_1 \equiv \psi * \mu_2$.

With the new version of (KM4), the KM postulates allow us to define a revision function as a *unary* function $* : L \mapsto L$, mapping the new evidence μ to a new belief base $*_\psi(\mu)$, given the initial belief base ψ as *background*. This is the first modification made by Nayak et al. The second is to make revision functions *dynamic*. That is, revision functions may change as new evidence arrives. With this relaxation it is possible for example to have one revision function $*_\psi$ associated initially with ψ, and a totally different one after the revision of ψ by a sequence of evidence that have made the full circle and have converted ψ back to itself.[4] Notice that the weakening of (KM4) to (KM4)' is consistent with such dynamic behavior.

As shown by Nayak et al., these two modifications suffice to reconcile the DP postulates with the KM ones, and it is these modifications we shall adopt for the rest of the paper.[5] Hence for the rest of the paper, unless specifically mentioned otherwise, we shall use the term "KM postulates" to refer to (KM1)-(KM6) with (KM4) replaced by (KM4)', and we shall assume that revision function are unary and dynamic.

We close this section with a remark on notation. Although we assume that revision functions are *unary* (relative to some background belief base ψ), for ease of presentation we shall keep the original notation and denote the revision of ψ by μ as $\psi * \mu$ rather than $*_\psi(\mu)$.

5 RELEVANCE-SENSITIVE BELIEF REVISION

Leaving temporarily aside the issue of iterated belief revision, we shall now turn back to one-step revisions to review the role of *relevance* in this process.

As already mentioned in the introduction, Parikh in [8] pointed out that the AGM/KM postulates fail to capture the intuition that during the revision of a belief base ψ by μ, only the part of ψ that is related to μ should be effected, while everything else should stay the same.

Of course determining the part of ψ that is relevant to some new evidence μ is not a simple matter. There is however at least one special case where the role of relevance can be adequately formalized; namely, when it is possible to break down ψ into two (syntactically) independent parts such that only the fist of the two parts is syntactically related to the new evidence μ.

More precisely, for a sentence φ of L, we shall denote by A_φ the smallest set of propositional variables, through which a sentence that is logically equivalent to φ can be expressed. For example, if φ is the sentence $(p \vee q \vee z) \wedge (p \vee q \vee \neg z)$, then $A_\varphi = \{p, q\}$, since φ is logically equivalent to $p \vee q$, and no sentence with fewer propositional variables is logically equivalent to φ. We shall denote by L_φ the propositional sublanguage built from A_φ via the usual boolean connectives. By $\overline{L_\varphi}$ we shall denote the sublanguage built from the *complement* of A_φ, i.e. from $A - A_\varphi$. Parikh proposed the following postulate to capture the role of relevance in belief revision (at least for the special case mentioned above):[6]

[4] In such cases, although the sequence of evidence does not effect the beliefs of the agent, it does however change the way the agent reacts to new input.

[5] It should be noted though that our results still hold even if Darwiche and Pearl's proposal of switching to belief states was adopted.

[6] The formulation of (P) in [8] is slightly different from the one presented below since Parikh was working with theories rather than belief bases. The two version are of course equivalent.

(P) If $\psi \equiv \chi \wedge \varphi$ where χ, φ are sentences of disjoint sublanguages L_χ, L_φ respectively, and $L_\mu \subseteq L_\chi$, then $\psi * \mu \equiv (\chi \circ \mu) \wedge \varphi$, where $\circ$ is a revision function of the sublanguage L_χ.

According to postulate (P), whenever it is possible to break down the initial belief base ψ into two independent parts χ and φ, and moreover it so happens that the new evidence μ can be expressed entirely in the language of the first part, then during the revision of ψ by μ, it is only the first part that is effected; the unrelated part φ crosses over to the new belief base verbatim.

Notice that of the nature of the "local" revision operator $\circ$ and its relationship with the "global" revision operator $*$ is not clearly stated in (P). Peppas et al., [10], therefore proposed a re-formulation of axiom (P) in terms of two new conditions (R1) and (R2) that do not refer to a "local" revision operator $\circ$. Only the first of these two conditions will be needed herein:

(R1) If $\psi \equiv \chi \wedge \varphi$, $L_\chi \cap L_\varphi = \emptyset$, and $L_\mu \subseteq L_\chi$, then $Cn(\psi) \cap \overline{L_\chi} = Cn(\psi * \mu) \cap \overline{L_\chi}$.

At first (R1) looks almost identical to (P) but it is in fact strictly weaker than it (see [10] for details). It is essentially condition (R1) that we will be using to derive our incompatibility results.

6 INCOMPATIBILITY RESULTS

As already announced in the introduction, we shall now prove that in the presence of the KM postulates, (R1) – and therefore (P) – is inconsistent with each of the postulates (DP1)-(DP4). The proof relies on the semantics characterization of these postulates so we shall briefly review it before presenting our results.

We start with Grove's seminal representation result [4] and its subsequent reformulation by Katsuno and Mendelzon [6].

Let ψ be a belief base and $\leq_\psi$ a total preorder in $\mathcal{M}_\mathcal{L}$. We denote the strict part of $\leq_\psi$ by $<_\psi$. We shall say that $\leq_\psi$ is *faithful* iff the minimal elements of $\leq_\psi$ are all the ψ-worlds:[7]

(SM1) If $r \in [\psi]$ then $r \leq_\psi r'$ for all $r' \in \mathcal{M}_\mathcal{L}$.
(SM2) If $r \in [\psi]$ and $r' \notin [\psi]$ then $r <_\psi r'$.

Given a belief base ψ and a faithful preorder $\leq_\psi$ associated with it, one can define a revision function $* : L \mapsto L$ as follows:

(S*) $\psi * \mu = \gamma(min([\mu], \leq_\psi)$.

In the above definition $min([\mu], \leq_\psi)$ is the set of minimal μ-world with respect to $\leq_\psi$, while γ is a function that maps a set of possible worlds S to a sentence $\gamma(S)$ such that $[\gamma(S)] = S$.

The preorder $\leq_\psi$ essentially represents comparative plausibility: the closer a world is to the initial worlds $[\psi]$, the more plausible it is. Then according to (S*), the revision of ψ by μ is defined as the belief base corresponding to the most plausible μ-worlds.

In [4, 6] it was shown that the function induced from (S*) satisfies the KM postulates and conversely, every revision function $*$ that satisfies the KM postulates can be constructed from a faithful preorder by means of (S*).[8]

[7] In [6] a third constraint was required for faithfulness, namely that logically equivalent sentences are assigned the same preorders. This is no longer necessary given the new version of (KM4).

[8] We note that the weakening of (KM4) to (KM4)' does not effect these results since it is accommodated by a corresponding weakening of the notion of faithfulness.

This correspondence between revision functions and faithful preorders can be preserved even if extra postulates for belief revision are introduced, as long as appropriate constraints are also imposed on the preorders. In particular, Darwiche and Pearl proved that the following four constraints (SI1)-(SI4) on faithful preorders correspond respectively to the four postulates (DP1)-(DP4).

(SI1) If $r, r' \in [\mu]$ then $r \leq_\psi r'$ iff $r \leq_{\psi*\mu} r'$.
(SI2) If $r, r' \in [\neg\mu]$ then $r \leq_\psi r'$ iff $r \leq_{\psi*\mu} r'$.
(SI3) If $r \in [\mu]$ and $r' \in [\neg\mu]$ then $r <_\psi r'$ entails $r <_{\psi*\mu} r'$.
(SI4) If $r \in [\mu]$ and $r' \in [\neg\mu]$ then $r \leq_\psi r'$ entails $r \leq_{\psi*\mu} r'$.

Notice that all of the above constraints make associations between the preorder $\leq_\psi$ related to the initial belief base ψ and the preorder $\leq_{\psi*\mu}$ related to the belief base that results from the revision of ψ by μ.

The semantic constraint(s) corresponding to postulate (P) have also been fully investigated in [10].[9] Herein however we shall focus only on condition (R1); in fact we shall be even more restrictive and consider only the semantic counterpart of (R1) in the special case of consistent and *complete* belief bases:

(PS) If *Diff*$(\psi, r) \subset$ *Diff*(ψ, r') then $r <_\psi r'$.

In the above condition, for any two worlds w, w', *Diff*(w, w') represents the set of propositional variables that have different truth values in the two worlds; in symbols, *Diff*$(w, w') = \{p \in A : w \vdash p$ and $w' \nvdash p\} \cup \{p \in A : w' \vdash p$ and $w \nvdash p\}$. Whenever a sentence ψ is consistent and complete, we use *Diff*(ψ, w') as an abbreviation of *Diff*$(Cn(\psi), w')$.

Intuitively, (PS) says that the plausibility of a world r depends on the propositional variables in which it differs from the initial (complete) belief base ψ: the more the propositional variables in *Diff*(ψ, r) the less plausible r is. In [10] it was shown that, for the special case of consistent and complete belief bases, (PS) is the semantic counterpart of (R1); i.e. given a consistent and complete belief base ψ and a faithful preorder $\leq_\psi$, the revision function $*$ produced from $\leq_\psi$ via (S*) satisfies (R1) iff $\leq_\psi$ satisfies (PS).

Although it is possible to obtain a fully-fledged semantic characterization of postulate (P) by generalizing (PS) accordingly (see [10]) the above restricted version suffices to establish the promised results:

Theorem 1 *In the presence of the KM postulates, postulate (P) is inconsistent with each of the postulates (DP1)-(DP4).*

Proof. Since (P) entails (R1) it suffices to show that (R1) is inconsistent with each of (DP1)-(DP4).

Assume that the object language L is built from the propositional variable p, q, and z. Moreover let ψ be the complete sentence $p \wedge q \wedge z$ and let $\leq_\psi$ be the following preorder in $\mathcal{M}_\mathcal{L}$:

$$pqz <_\psi pq\overline{z} <_\psi p\overline{q}z <_\psi \overline{p}qz <_\psi p\overline{q}\overline{z} <_\psi \overline{p}q\overline{z} <_\psi \overline{p}\overline{q}z <_\psi \overline{p}\overline{q}\overline{z}$$

In the above definition of $\leq_\psi$, and in order to increase readability, we have used sequences of literals to represent possible worlds (namely the literals satisfied by a world), and we have represented the negation of a propositional variable v by $\overline{v}$.

Notice that $\leq_\psi$ satisfies (PS). In what follows we shall construct sentences μ_1, μ_2, μ_3, and μ_4, such that no preorder satisfying (PS) and related to $\psi * \mu_1$ (respectively to $\psi * \mu_2$, $\psi * \mu_3$, $\psi * \mu_4$) can also

[9] In the same paper axiom (P) was shown to be consistent with all the AGM/KM postulates.

satisfy (SI1) (respectively (SI2), (SI3), (SI4)). Given the correspondence between (R1) and (PS) on one hand, and the correspondence between (DP1)-(DP4) and (SI1)-(SI4) on the other, this will suffice to prove the theorem.

Inconsistency of (PS) and (SI1):
Let μ_1 be the sentence $\overline{q} \vee \overline{z}$. According to the definition of $\leq_\psi$, there is only one minimal μ_1-world, namely $pq\overline{z}$, and therefore by (S*), $\psi * \mu_1 \equiv p \wedge q \wedge \overline{z}$. Consider now the possible worlds $w = p\overline{q}z$ and w' = $p\overline{q}\overline{z}$. Clearly, $\mathit{Diff}(\psi * \mu_1, w') = \{q\} \subset \{q, z\} = \mathit{Diff}(\psi * \mu_1, w)$. Consequently, no matter what the new preorder $\leq_{\psi * \mu_1}$ is, as long as it satisfies (PS) it holds that $w' <_{\psi * \mu_1} w$. On the other hand, since $w, w' \in [\mu_1]$ and $w \leq_\psi w'$, (SI1) entails that $w \leq_{\psi * \mu_1} w'$. Contradiction.

Inconsistency of (PS) and (SI2):
Let μ_2 be the sentence $p \wedge q \wedge \overline{z}$. Once again there is only one minimal μ_2-world, namely $pq\overline{z}$, and therefore $\psi * \mu_2 \equiv p \wedge q \wedge \overline{z}$. Let w and w' be the possible worlds $p\overline{q}z$ and $p\overline{q}\overline{z}$ respectively. It is not hard to verify that $\mathit{Diff}(\psi * \mu_2, w') \subset \mathit{Diff}(\psi * \mu_2, w)$ and therefore (PS) entails $w' <_{\psi * \mu_1} w$. On the other hand, given that $w, w' \in [\neg\mu_2]$ and $w \leq_\psi w'$, (SI2) entails that $w \leq_{\psi * \mu_1} w'$, leading us to a contradiction.

Inconsistency of (PS) and (SI3):
Let μ_3 be the sentence $(p \wedge q \wedge \overline{z}) \vee (p \wedge \overline{q} \wedge z)$. Given the above definition of $\leq_\psi$ it is not hard to verify that $\psi * \mu_3 \equiv p \wedge q \wedge \overline{z}$. Once again, define w and w' to be the possible worlds $p\overline{q}z$ and $p\overline{q}\overline{z}$ respectively. Clearly $\mathit{Diff}(\psi * \mu_3, w') \subset \mathit{Diff}(\psi * \mu_3, w)$ and therefore (PS) entails $w' <_{\psi * \mu_3} w$. On the other hand notice that $w \in [\mu_3]$, $w' \in [\neg\mu_3]$ and $w <_\psi w'$. Hence (SI3) entails that $w <_{\psi * \mu_3} w'$. Contradiction.

Inconsistency of (PS) and (SI4):
Let $\mu_4 = \mu_3 = (p \wedge q \wedge \overline{z}) \vee (p \wedge \overline{q} \wedge z)$ and assume that w, w' are as previously defined. Then clearly, like before, (PS) entails $w' <_{\psi * \mu_4} w$. On the other hand, since $w \in [\mu_4]$, $w' \in [\neg\mu_4]$ and $w \leq_\psi w'$, (SI4) entails that $w \leq_{\psi * \mu_4} w'$. A contradiction once again. □

The above inconsistencies extend to subsequent developments of the Darwiche and Pearl approach as well. Herein we consider two such extensions in the form of two new postulates introduced to rectify anomalies with the original DP approach.

The first postulate, called the *Conjunction Postulate*, was introduced by Nayak, Pagnucco, and Peppas in [7]:

(CNJ) If $\mu \wedge \varphi \not\vdash \perp$, then $\psi * \mu * \varphi \equiv \psi * (\mu \wedge \varphi)$.

As shown in [7], in the presence of the KM postulates, (CNJ) entails (DP1), (DP3), and (DP4). Hence the following result is a direct consequence of Theorem 1:

Corollary 1 *In the presence of the KM postulates, postulate (P) is inconsistent with postulate (CNJ).*

The last postulate we shall consider herein is the *Independence Postulate* proposed by Jin and Thielscher in [5]:

(Ind) If $\neg\mu \notin \psi * \varphi$ then $\mu \in \psi * \mu * \varphi$.

Jin and Thielscher proved that, although weaker than (CNJ), (Ind) still entails (DP3) and (DP4). Hence, from Theorem 1, it follows:

Corollary 2 *In the presence of the KM postulates, postulate (P) is inconsistent with postulate (Ind).*

7 CONCLUSION

In this paper we have proved the inconsistency of Parikh's postulate (P) for relevance-sensitive belief revision, with each of the four postulates (DP1)-(DP4) proposed by Darwiche and Pearl for iterated belief revision. This result suggests that a major refurbishment may be due in our formal models for belief revision. Both relevance and iteration are central to the process of belief revision and neither of them can be sacrificed. Moreover, the formalizations of these notions by postulates (P) and (DP1)-(DP2) respectively seem quite intuitive and it is not clear what amendments should be made to reconcile them.

On a more positive note, the inconsistencies proved herein reveal a hitherto unknown connection between relevance and iteration, which will eventually lead to a deeper understanding of the intricacies of belief revision.

REFERENCES

[1] C. Alchourron and P. Gardenfors and D. Makinson, "On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision", *Journal of Symbolic Logic*, vol 50, pp 510-530, 1985.

[2] A. Darwiche and J. Pearl, "On the Logic of Iterated Belief Revision", *Artificial Intelligence*, vol. 89, pp 1-29, 1997.

[3] P. Gardenfors, *Knowledge in Flux*. MIT Press, 1988.

[4] A. Grove, "Two modellings for theory change", *Journal of Philosophical Logic* vol. 17, pp. 157-170, 1988.

[5] Y. Jin and M. Thielscher, "Iterated Revision, Revised", *Proceedings of the 19th International Joint Conference in Artificial Intelligence*, pp 478-483, Edinburgh, 2005.

[6] H. Katsuno and A. Mendelzon, "Propositional Knowledge Base Revision and Minimal Change", in *Artificial Intelligence*, vol. 52, pp 263-294, 1991.

[7] A. Nayak, M. Pagnucco, and P. Peppas, "Dynamic Belief Revision Operators", *Artificial Intelligence*, vol. 146, pp 193-228, 2003.

[8] R. Parikh, "Beliefs, Belief Revision, and Splitting Languages", in J. Lawrence Moss, M. de Rijke, (eds)., *Logic, Language, and Computation*, vol 2, pp 266–268, CSLI Lecture Notes No. 96,. CSLI Publications, 1999.

[9] P. Peppas, N. Foo, and A. Nayak, "Measuring Similarity in Belief Revision", *Journal of Logic and Computation*, vol. 10(4), 2000.

[10] P. Peppas, S. Chopra, and N. Foo, "Distance Semantics for Relevance-Sensitive Belief Revision", *Proceedings of the 9th International Conference on the Principles of Knowledge Representation and Reasoning (KR2004)*, Whistler, Canada, June 2004.

[11] P. Peppas, "Belief Revision", in F. van Harmelen, V. Lifschitz, and B. Porter (eds), *Handbook in Knowledge Representation*, Elsevier Publications, 2007.

[12] M. Winslett, "Reasoning about Action using a Possible Models Approach", in *Proceedings of the 7th National (USA) Conference on Artificial Intelligence (AAAI'88)*, pp. 89-93, 1988.