Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence

Revising Horn Theories

James P. Delgrande
School of Computing Science
Simon Fraser University
Burnaby, B.C.
Canada V5A 1S6
jim@cs.sfu.ca

Pavlos Peppas
Dept of Business Administration
University of Patras
Patras 265 00
Greece
pavlos@upatras.gr

Abstract

This paper investigates belief revision where the underlying logic is that governing Horn clauses. It proves to be the case that classical (AGM) belief revision doesn't immediately generalise to the Horn case. In particular, a standard construction based on a total preorder over possible worlds may violate the accepted (AGM) postulates. Conversely, Horn revision functions in the obvious extension to the AGM approach are not captured by total preorders over possible worlds. We address these difficulties by first restricting the semantic construction to "well behaved" orderings; and second, by augmenting the revision postulates by an additional postulate. This additional postulate is redundant in the AGM approach but not in the Horn case. In a representation result we show that these two approaches coincide. Arguably this work is interesting for several reasons. It extends AGM revision to inferentially-weaker Horn theories; hence it sheds light on the theoretical underpinnings of belief change, as well as generalising the AGM paradigm. Thus, this work is relevant to revision in areas that employ Horn clauses, such as deductive databases and logic programming, as well as areas in which inference is weaker than classical logic, such as in description logic.

1 Introduction

The area of belief change studies how an agent may modify its beliefs given new information about its environment. The best-known approach in this area is the AGM paradigm [Alchourrón *et al.*, 1985; Gärdenfors, 1988], named after the original developers. This work focussed on belief *revision*, in which new information is incorporated into an agent's belief corpus, as well as belief *contraction*, in which an agent may reduce its set of beliefs. The AGM approach addresses belief change at an abstract level, in which an agent's beliefs are characterised by *belief sets* or deductively closed sets of sentences, and where the underlying logic includes classical propositional logic. In the basic approach to revision, a set of rationality postulates is given which arguably any revision function should satisfy. As well, a semantic construction of revision functions has been given, in terms of a total preorder over possible worlds, called a *faithful ranking*. These syntactic and semantic approaches are then shown to capture the same set of revision functions.

In this paper we consider AGM-style belief revision in the language of *Horn clauses*, where a Horn clause can be expressed as a rule of the form $a_1 \wedge a_2 \wedge \cdots \wedge a_n \rightarrow a$ for $n \geq 0$, and where a, a_i $(1 \leq i \leq n)$ are atoms. (Thus, expressed in conjunctive normal form, a Horn clause will have at most one positive literal.) In our approach an agent's beliefs are represented by a Horn clause knowledge base, and the input is a Horn formula, consisting of a conjunction of Horn clauses. It proves to be the case that AGM-style belief revision doesn't transfer directly to Horn knowledge bases. On the one hand, in the Horn case the AGM postulate set is unsound with respect to a faithful ranking over possible worlds. On the other hand, given a Horn revision function that satisfies the AGM postulates, there may be no corresponding faithful ranking that captures the revision function or, alternately, there may be several faithful rankings that capture the function.

Nonetheless, we show that the AGM approach may be extended to the Horn case naturally and satisfactorily. On the semantic side, we impose a "well-behaved" condition on faithful rankings, expressing that a ranking must be coherent with respect to Horn revision. (In propositional logic, the corresponding condition proves to be trivial.) On the syntactic, postulational, side, we add a postulate to the standard suite of AGM postulates. Interestingly, in the AGM approach this additional postulate is redundant, in that it follows as a theorem from the other AGM postulates. In the Horn case, in which inference is weaker than in the classical case, this postulate is independent of the others. Given these adjustments to the AGM approach, we then prove a representation result, showing that the class of Horn revision functions conforming to the extended postulate set is the same as those capturable by "well-behaved" faithful rankings.

This topic is interesting for several reasons. It sheds light on the theory of belief change, in that it weakens the assumption that the underlying logic contains propositional logic. In so doing, it shows that the AGM approach is more generally applicable than perhaps originally believed. That is, our results provide a *broadening* of the AGM approach to include Horn reasoning, and not just a modification of the AGM approach to accommodate Horn reasoning. Horn clauses are a very useful restriction of classical logic, and have found widespread application in artificial intelligence and database theory. As well, results here may also be relevant to belief change in description logics, a topic that has also received recent attention.

The next section introduces belief change, Horn clause reasoning, and work that has been carried out in the intersection of these areas. This is followed by a discussion of issues that arise in Horn clause belief revision (Sections 4 and 5). The following section develops the approach, and in particular presents the representation result for Horn formula revision. The paper concludes with a discussion of future work and a brief conclusion. Proofs are omitted due to space considerations, but are available in the full paper.

2 Formal Preliminaries

We introduce here the terminology that we will use in the rest of the paper. $\mathcal{P} = \{a, b, c, \dots\}$ is a finite set of propositional variables. $\mathcal{L}_H$ denotes the Horn formula language over $\mathcal{P} \cup \{\bot\}$. That is, $\mathcal{L}_H$ is the least set given by:

1. $a_1 \wedge a_2 \wedge \cdots \wedge a_n \rightarrow a$ is a *Horn clause*, where $n \geq 0$ and $a, a_i \in \mathcal{P} \cup \{\bot\}$ for $1 \leq i \leq n$.
   If $n = 0$ then $\rightarrow a$ is also written a, and is a *fact*.
2. Every Horn clause is a *Horn formula*.
3. If ϕ and ψ are Horn formulas then so is $\phi \wedge \psi$.

In our approach, we deal exclusively with Horn formulas, and so *formula* will refer to a Horn formula; the only exception is when we discuss background work, in which case the context is clear. Formulas are denoted by lower case Greek letters; sets of formulas are denoted by upper case Greek letters.

An *interpretation* m is a subset of $\mathcal{P}$, where $a \in m$ is *true* and $a \notin m$ is *false*. Occasionally we will explicitly list negated atoms in an interpretation; for example for $\mathcal{P} = \{p, q\}$ the interpretation $\{p\}$ will sometimes be written $\{p, \neg q\}$ or more briefly $p\bar{q}$. The symbol $\bot$ is always assigned *false*. $\mathcal{M}$ is the set of *interpretations* or *(possible) worlds*. Sentences of $\mathcal{L}_H$ are *true* or *false* in an interpretation according to the standard rules in propositional logic. Truth of ϕ in m is denoted $m \models \phi$. As well, for $W \subseteq \mathcal{M}$, $W \models \phi$ iff for every $m \in W$, $m \models \phi$. For formula ϕ, $|\phi|$ is the set of models of ϕ. For set of worlds W, $t_H(W)$ denotes the set of formulas satisfied by all worlds in W, i.e. $t_H(W) = \{\phi \in \mathcal{L}_H \mid m \models \phi \text{ for every } m \in W\}$. Note that this means that $t_H(\emptyset) = \mathcal{L}_H$.

$\Gamma \vdash \phi$ iff ϕ is derivable from set of formulas Γ [Delgrande, 2008] (again, members of Γ and ϕ are Horn, and $\vdash$ is defined in terms of Horn formulas). $\psi \vdash \phi$ is an abbreviation for $\{\psi\} \vdash \phi$, and $\psi \equiv \phi$ is logical equivalence, i.e. $\psi \vdash \phi$ and $\phi \vdash \psi$. This extends in the obvious fashion to sets of formulas. For a set of formulas Γ, the closure of Γ under Horn derivability is denoted $\mathcal{C}n_H(\Gamma)$. A (Horn) theory H is a set of formulas such that $H = \mathcal{C}n_H(H)$, also referred to as a *belief set*. $\mathcal{H}$ is the set of Horn theories. For theory H and formula ϕ, $H + \phi = \mathcal{C}n_H(H \cup \{\phi\})$ is the *expansion* of H. $H_\bot = \mathcal{L}_H$ is the inconsistent belief set.

Models of Horn formulas are distinguished by the fact that they are closed under intersection of positive atoms in an interpretation. That is:

If $m_1, m_2 \in |\phi|$ then $m_1 \cap m_2 \in |\phi|$.

We note that the converse is also true; i.e., if a set of models W is closed under intersection of positive atoms in an interpretation, then there is a Horn formula ϕ such that $|\phi| = W$.

$Cl_\cap(W)$ is the intersection closure of a set of interpretations W, i.e. $Cl_\cap(W)$ is the least set such that $W \subseteq Cl_\cap(W)$ and $m_1, m_2 \in Cl_\cap(W)$ implies that $m_1 \cap m_2 \in Cl_\cap(W)$.

A (partial) preorder $\preceq$ is a reflexive, transitive binary relation. A total preorder is a partial preorder such that either $m_1 \preceq m_2$ or $m_2 \preceq m_1$ for every m_1, m_2. The strict part of preorder $\preceq$ is denoted by $\prec$, that is, $m_1 \prec m_2$ just if $m_1 \preceq m_2$ and $m_2 \not\preceq m_1$. As usual, $m_1 \approx m_2$ abbreviates $m_1 \preceq m_2$ and $m_2 \preceq m_1$. Finally, for a set of interpretations $W \subseteq \mathcal{M}$, by $\min(W, \preceq)$ we denote the set $\min(W, \preceq)$ = $\{m_1 \in W\text{: for all } m_2 \in W, \text{ if } m_2 \preceq m_1 \text{ then } m_1 \preceq m_2\}$.

3 Background

3.1 Belief Revision

In the AGM approach to belief change, beliefs of an agent are modelled by a deductively closed set of formulas, or *belief set*. Thus a belief set is a set of formulas K such that $K = \mathcal{C}n(K)$.[1] It is assumed that the underlying logic subsumes classical propositional logic. Belief revision is modeled as a function from a belief set K and a formula ϕ to a belief set K' such that ϕ is believed in K', i.e. $\phi \in K'$. Since ϕ may be inconsistent with K, and since it is desirable to maintain consistency whenever possible (i.e. whenever ϕ is consistent) then some formulas may need to dropped from K before ϕ can be consistently added. Formally, a revision operator $*$ maps a belief set K and formula ϕ to a revised belief set $K * \phi$. The AGM postulates for revision can be formulated as follows, where $\equiv_{PC}$ and $+_{PC}$ stand for logical equivalence and expansion, respectively, in classical propositional logic.

(K*1) $K * \phi = \mathcal{C}n(K * \phi)$

(K*2) $\phi \in K * \phi$

(K*3) $K * \phi \subseteq K +_{PC} \phi$

(K*4) If $\neg\phi \notin K$ then $K +_{PC} \phi \subseteq K * \phi$

(K*5) $K * \phi$ is inconsistent only if ϕ is inconsistent

(K*6) If $\phi \equiv_{PC} \psi$ then $K * \phi = K * \psi$

(K*7) $K * (\phi \wedge \psi) \subseteq K * \phi +_{PC} \psi$

(K*8) If $\neg\psi \notin K * \phi$ then $K * \phi +_{PC} \psi \subseteq K * (\phi \wedge \psi)$

Thus, the result of revising K by ϕ yields a belief set in which ϕ is believed ((K*1), (K*2)); whenever the result is consistent, the revised belief set consists of the expansion of K by ϕ ((K*3), (K*4)); the only time that K is inconsistent is when ϕ is inconsistent ((K*5)); and revision is independent of the syntactic form of the formula for revision ((K*6)). The last two postulates deal with the relation between revising by a conjunction and expansion: whenever

[1] $\mathcal{C}n(K)$ denotes the closure of K under classical logical implication.

consistent, revision by a conjunction corresponds to revision by one conjunct and expansion by the other. Motivation for these postulates can be found in [Gärdenfors, 1988; Peppas, 2008]. The intent of these postulates is that they should hold for *any* rational belief revision function.

Katsuno and Mendelzon [1991] have shown that a necessary and sufficient condition for constructing an AGM revision operator is that there is a function that associates a total preorder on the set of possible worlds with any belief set K, as follows:[2]

Definition 1 *A* faithful assignment *is a function that maps each belief set K to a total preorder $\preceq_K$ on $\mathcal{M}$ such that for any possible worlds w_1, w_2:*

1. *If $m_1, m_2 \in |K|$ then $m_1 \approx_K m_2$*
2. *If $m_1 \in |K|$ and $m_2 \notin |K|$, then $m_1 \prec_K m_2$.*

The resulting preorder is referred to as the *faithful ranking* associated with K. Intuitively, $w_1 \preceq_K w_2$ if w_1 is at least as plausible as w_2. Katsuno and Mendelzon then provide the following representation result, where $t(W)$ is the set of formulas of classical logic true in W:

Theorem 1 ([Katsuno and Mendelzon, 1991]) *A revision operator * satisfies postulates (K*1)–(K*8) iff there exists a faithful assignment that maps each belief set K to a total preorder $\preceq_K$ such that*

$$K * \phi \;=\; t(\min(|\phi|, \preceq_K)).$$

Thus the revision of K by ϕ is characterised by those models of ϕ that are most plausible according to the agent. Given that we are working with a finite language, this construction is in fact equivalent to the earlier *systems of spheres* approach of [Grove, 1988]. It is easier to present our results in terms of faithful assignments, and so we do so here.

Another form of belief change in the AGM approach is belief *contraction*, in which an agent's beliefs strictly decrease. Thus in the contraction $K - \phi$, one has $\phi \notin K - \phi \subseteq K$, while $\neg\phi$ is not necessarily believed. There are two primary means of constructing contraction functions. Using *remainder sets*, a contraction $K - \phi$ is defined in terms of maximal subsets of K that fail to imply ϕ. Via *epistemic entrenchment*, an ordering is defined on sentences of K, and a contraction $K - \phi$ is (roughly) defined in terms of the most entrenched set of sentences that does not imply ϕ. Of interest, and pertinent to the approach at hand, these various constructions are all in a certain sense interdefinable, as are revision and contraction functions. Hence, given a contraction function $-$, one may define a revision function by the so-called Levi identity:

$$K * \phi = (K - \neg\phi) +_{PC} \phi. \tag{1}$$

See [Gärdenfors, 1988; Peppas, 2008] for details.

[2] In fact, Katsuno and Mendelzon deal with formulas instead of belief sets. Since we deal with finite languages only, the difference is immaterial. We use belief sets in order to adhere more closely to the original AGM approach.

3.2 Related Work

Earlier work on belief change involving Horn formulas dealt with the Horn fragment of a propositional theory, rather than Horn clause belief change as a distinct phenomenon. For example, the complexity of specific approaches to revising knowledge bases is addressed in [Eiter and Gottlob, 1992], including the case where the knowledge base and formula for revision are Horn formulas. [Liberatore, 2000] considers the problem of compact representation for revision in the Horn case. Given a knowledge base K and formula ϕ, both Horn, the main problem considered is whether a revised knowledge base can be expressed by a propositional formula whose size is polynomial with respect to the sizes of K and ϕ.

[Langlois *et al.*, 2008] approaches the study of revising Horn formulas by characterising the existence of a complement of a Horn consequence; such a complement corresponds to the result of a contraction operator. This work may be seen as a specific instance of a general framework developed in [Flouris *et al.*, 2004].

The main difference between our work and the above approaches to Horn revision, is that in our approach revision functions *always* produce Horn theories (they are postulated to do so). This of course adds an extra burden to the revision process since it now needs to comply with both the *principle of minimal change* (see [Gärdenfors, 1988; Peppas, 2008]), and the requirement to produce Horn theories (which in this context can be seen as an instance of the *principle of categorical matching*). Our results show that, with some adjustments to the original AGM framework, this double objective can indeed be achieved.

Similar results have already been acheived for Horn contraction. [Delgrande, 2008] addresses maxichoice belief contraction in Horn clause theories, where contraction is defined in terms of remainder sets. Further developments in Horn contraction can be found in [Booth *et al.*, 2009], [Delgrande and Wassermann, 2010], and [Zhuang and Pagnucco, 2010].

4 Horn Revision: Preliminary Considerations

The postulates and semantic construction of Section 3.1 are easily adapted to Horn theories. For the postulates, we have the following, expressed in terms of Horn theories.

An *AGM (Horn) revision* function $*$ is a function from $\mathcal{H} \times \mathcal{L}_H$ to $\mathcal{H}$ satisfying the following postulates.

(H*1) $H * \phi = \mathcal{C}n_H(H * \phi)$.

(H*2) $\phi \in H * \phi$.

(H*3) $H * \phi \subseteq H + \phi$.

(H*4) If $\bot \notin H + \phi$ then $H + \phi \subseteq H * \phi$.

(H*5) If ϕ is consistent then $\bot \notin H * \phi$.

(H*6) If $\psi \equiv \phi$ then $H * \psi = H * \phi$.

(H*7) $H * (\psi \wedge \phi) \subseteq (H * \psi) + \phi$.

(H*8) If $\bot \notin (H * \psi) + \phi$ then $(H * \psi) + \phi \subseteq H * (\psi \wedge \phi)$.

As well, faithful assignments can be defined for the Horn case, basically by changing notation:

Definition 2 *A* faithful assignment *is a function that maps each Horn theory H to a total preorder $\preceq_H$ on $\mathcal{M}$ such that for any possible worlds m_1, m_2:*

1. *If $m_1, m_2 \in |H|$ then $m_1 \approx_H m_2$*
2. *If $m_1 \in |H|$ and $m_2 \notin |H|$, then $m_1 \prec_H m_2$.*

The resulting preorder is referred to as the *faithful ranking* associated with H. Finally, one can define a function $*$ in terms of a faithful ranking by:

$$H * \phi \;=\; t_H(\min(|\phi|, \preceq_H)). \tag{2}$$

The use of $*$ in (2) is suggestive; ideally one would next establish a correspondence between functions that satisfy the postulates and those that can be defined via Definition 2. However, there are significant difficulties in immediately establishing such a representation result. We review these problems next, and then present our solution in the following section.

5 Problems with Naïve AGM Horn revision

1. Interdefinability results do not hold in Horn belief change. As mentioned in Section 3.1, in the AGM approach revision may be defined in terms of contraction via the Levi Identity (1). However, previous work [Delgrande and Wassermann, 2010] suggests that Horn contraction is unsuitable for specifying a revision operator. As well, if one considers the Levi Identity, revision by a Horn formula ϕ is defined in terms of the contraction by $\neg\phi$. Since ϕ is a conjunction of Horn clauses, $\neg\phi$ in general will not be Horn, and so the Levi identity would seem to be inapplicable for Horn theories.

These points are not definitive (there is, after all, no formal result stating an impossibility of interdefinability of Horn contraction and revision), but they do suggest the overall difficulty in obtaining such a result. Consequently, we focus on a direct definition of Horn revision, in terms of ranking functions in the next section. Having developed such an approach, we then suggest that the relation between Horn contraction and revision is a suitable and interesting topic for future research.

2. Distinct rankings may yield the same revision function. Consider the Horn language defined by $\mathcal{P} = \{p, q\}$, and the following three total preorders:

$$\begin{array}{lll} pq & \prec & \overline{pq} \prec p\overline{q} \prec \overline{p}q \end{array} \tag{3}$$
$$\begin{array}{lll} pq & \prec & \overline{pq} \prec \overline{p}q \prec p\overline{q} \end{array} \tag{4}$$
$$\begin{array}{lll} pq & \prec & \overline{pq} \prec \overline{p}q \approx p\overline{q} \end{array} \tag{5}$$

It can be verified that if one defines revision via Definition 2, the three total preorders yield the same revision function. In particular, there is no way in which the relative ranking of worlds given by $p\overline{q}$ and $\overline{p}q$ can be distinguished.

The difficulty is that, in AGM revision, given a revision function that satisfies the AGM revision postulates, for interpretations m_1 and m_2 one can define $m_1 \preceq_K m_2$ just if there is ϕ such that $m_1 \in |K * \phi|$ and $m_2 \in |\phi|$. The problem in the example is that for $m_1 = \{\overline{p}q\}$ and $m_2 = \{p\overline{q}\}$, then for any ϕ with $m_1, m_2 \in |\phi|$, one also has $m_3 = \overline{pq} \in |\phi|$. In the example, for $H = \mathcal{C}n_H(p \wedge q)$ and $\phi = p \wedge q \rightarrow \bot$, $H * \phi = \mathcal{C}n_H(\neg p \wedge \neg q)$. Consequently the relative ranking of m_1 and m_2 cannot be specified.

3. Postulates may not be satisfied in a faithful ranking. Consider the Horn language with atoms $\mathcal{P} = \{p, q, r\}$ and the ranking:

$$pqr \prec \overline{p}q\overline{r} \approx p\overline{qr} \prec \overline{pq}r \prec \overline{pqr} \prec \text{all other worlds} \tag{6}$$

The agent's belief set H is given by $\mathcal{C}n_H(p \wedge q \wedge r)$. Let μ be $p \wedge q \rightarrow \bot$ and ϕ be $\neg p \wedge \neg q$. Defining $*$ as in Definition 2, it can be verified that:

$$\begin{array}{rcl} H * \mu & = & \mathcal{C}n_H((p \wedge q \rightarrow \bot) \wedge \neg r)). \\ (H * \mu) + \phi & = & \mathcal{C}n_H(\neg p \wedge \neg q \wedge \neg r). \\ H * (\mu \wedge \phi) & = & \mathcal{C}n_H(\neg p \wedge \neg q \wedge r). \end{array}$$

Thus $(H * \mu) + \phi$ and $H * (\mu \wedge \phi)$ are not equivalent and violate both (H*7) and (H*8).

Informally, the culprit is the set of worlds $\{\overline{p}q\overline{r}, p\overline{qr}\}$. This set (as with the previous problem) is not expressible by a Horn formula, since it is not closed under intersection of (positive) atoms. It can be observed that the "missing" interpretation is given by $\overline{pqr}$, where in addition we have $\overline{p}q\overline{r} \approx p\overline{qr} \prec \overline{pqr}$. The problem arises then because one may revise by a Horn formula (viz. $\mu = p \wedge q \rightarrow \bot$) that yields the set of minimal models $\{\overline{p}q\overline{r}, p\overline{qr}\}$, but in producing the corresponding Horn theory $t_H(\{\overline{p}q\overline{r}, p\overline{qr}\})$ = $\mathcal{C}n_H((p \wedge q \rightarrow \bot) \wedge \neg r))$, a new *non-minimal* model $\overline{pqr}$ creeps in.[3]

4. There is a Horn AGM revision function satisfying (H*1)-(H*8) that cannot be modelled by a preorder on worlds. Consider the following pseudo-preorder on worlds:

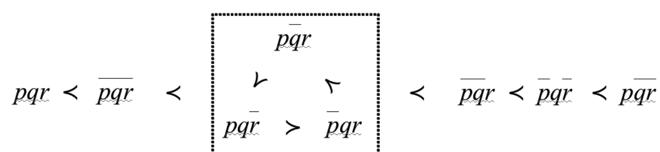

Figure 1

That is, the most preferred world is pqr, followed by $\overline{pqr}$, followed by $pq\overline{r}, \overline{p}qr, p\overline{q}r$ which form a *cycle* (i.e. $pq\overline{r} \prec \overline{p}qr \prec p\overline{q}r \prec pq\overline{r}$). Finally comes the sequence of worlds $\overline{pq}r \prec \overline{p}q\overline{r} \prec p\overline{qr}$.

Based on the pseudo-preorder $\preceq$ of Figure 1, we can define a Horn revision function according to Definition 2. The only time this revision function can run into trouble (i.e. produce an inconsistent result for a consistent ϕ) is if ϕ is inconsistent with the first two worlds $pqr, \overline{pqr}$ and consistent with all three world in the cycle; i.e. $\{pqr, \overline{pqr}\} \cap |\phi| = \emptyset$ and $\{pq\overline{r}, \overline{p}qr, p\overline{q}r\} \subseteq |\phi|$. However, no such Horn formula ϕ

[3]As a sublety, this doesn't imply that a maximum set of equivalently-ranked worlds must be definable by a Horn formula. For example (5) contains a set of worlds that isn't definable by a Horn formula. However, in our approach this ranking will prove to be an acceptable ordering on worlds.

exists. To see this notice that any Horn formula μ such that $\{pq\overline{r}, \overline{p}qr, p\overline{q}r\} \subseteq [\mu]$, has $\overline{pqr}$ as a model (since the models of Horn formulas are closed under intersections of the positive atoms in an interpretation).

It proves to be the case that * as defined in Figure 1 satisfies all AGM postulates for Horn revision. However, it can also be shown that no (partial or total) preorder on worlds can model *. The details are omitted due to space constraints but are included in the full paper.

6 Horn Revision: The Approach

As is clear from the previous discussion, there are substantial differences between classical AGM revision and Horn revision. Since the AGM postulates and the Katsuno-Mendelzon construction are essentially the same for Horn theories, these differences come about from the weakened expressibility of Horn clause theories.

Consider again the issues discussed in the previous section. The first issue isn't a problem per se. Rather, it suggests that, when we look at Horn theories, belief change operators are not interdefinable, or at best are not readily interdefinable.

The second issue also isn't a problem as such. Instead, it indicates that a ranking may be underconstrained by a revision function. In our approach, we address this indirectly by defining a canonical ranking, in which if there is no reason to distinguish the rank of worlds then they are assigned the same rank.

The third issue, that a ranking may violate the postulates (H*7) and (H*8), is indeed a problem. As discussed, the difficulty essentially is that some orderings are unsuitable with respect to Horn revision. The solution then is to add a constraint to faithful orderings such that these "unsuitable orderings" are ruled out. This is covered by the notion of *Horn compliance*, defined below.

As the fourth problem demonstrates, the Horn AGM postulates may fail to rule out undesirable relations on sets of worlds, which is to say, the postulate set is too weak to eliminate certain undesirable non-preorders. This then requires adding to the postulate set to (semantically) further constrain the set of allowable orderings.

6.1 A Representation Result

In accordance with the previous discussion, on the one hand we add a condition to restrict rankings on worlds; on the other we add a postulate to the set of Horn AGM postulates.

On the semantic side, we restrict rankings to those that yield coherent results with respect to Horn revision. That is, we want to allow only those orderings where revision by a Horn formula will yield a set of worlds corresponding to a Horn formula.

Call a set of worlds W *Horn elementary* iff it is definable via a Horn formula, i.e. if there is a Horn formula ϕ such that $W = |\phi|$. So W is Horn elementary iff $W = Cl_\cap(W)$. A preorder $\preceq_H$ is *Horn compliant* iff for every formula $\phi \in \mathcal{L}_H$, $\min(|\phi|, \preceq_H)$ is Horn elementary.

For example, the preorder in (5) is Horn compliant. Note that, while the set $\{\overline{p}q, p\overline{q}\}$ is not Horn elementary, there is no Horn formula ϕ over $\mathcal{P} = \{p, q\}$ such that $\min(|\phi|, \preceq_H) = \{\overline{p}q, p\overline{q}\}$. On the other hand, the ordering in (6) is not Horn compliant since $\min(|p \wedge q \rightarrow \bot|, \preceq_H) = \{\overline{p}q\overline{r}, p\overline{qr}\}$, and $\{\overline{p}q\overline{r}, p\overline{qr}\}$ is not Horn elementary.

With respect to postulates, we introduce the following schema:

(Acyc) If for $0 \leq i < n$ we have $(H * \mu_{i+1}) + \mu_i \not\vdash \bot$, and $(H * \mu_0) + \mu_n \not\vdash \bot$, then $(H * \mu_n) + \mu_0 \not\vdash \bot$.

We note the following minor results:

Proposition 1 *(Acyc) is a logical consequence of the AGM postulates (i.e. where the underlying logic contains classical propositional logic).*

Proposition 2 *(Acyc) is independent of the Horn AGM postulates.*

Figure 1 suffices to show the independence of (Acyc) from the Horn AGM postulates: defining revision according to Definition 2 yields a revision function that satisfies the AGM postulates for Horn revision; however it can be verified that it violates (Acyc).

Informally, (Acyc) rules out cycles (of any length n) as found for example in Figure 1. To see this, consider the instance of (Acyc) for $n = 2$:

If $(H * \mu_1) + \mu_0 \not\vdash \bot$ and $(H * \mu_2) + \mu_1 \not\vdash \bot$ and $(H * \mu_0) + \mu_2 \not\vdash \bot$ then $(H * \mu_2) + \mu_0 \not\vdash \bot$.

If revision is defined as in Definition 2, then $(H*\mu_1)+\mu_0 \not\vdash \bot$ iff $\min(|\mu_1|, \preceq_H) \cap |\mu_0| \neq \emptyset$. This last relation implies that, for $w_1 \in \min(|\mu_1|, \preceq_H)$ and $w_0 \in \min(|\mu_0|, \preceq_H)$ it must be that $w_0 \preceq_H w_1$. Consequently, the postulate can then be read as requiring that if $w_0 \preceq_H w_1 \preceq_H w_2 \preceq_H w_0$ then $w_0 \preceq_H w_2$, and with some further deliberation, $w_0 \approx_H w_1 \approx_H w_2$ (thus ruling out the pseudo-preorder of Figure 1).

The notion of Horn compliance on the one hand, and the postulate (Acyc) on the other, prove to be sufficient to extend the AGM approach to capture revision in Horn theories. We obtain the following results:

Theorem 2 *Let H be a Horn belief set and $\preceq_H$ a Horn compliant faithful ranking associated with H. Define an operator $* : \mathcal{H} \times \mathcal{L}_H \mapsto \mathcal{H}$ by $H * \phi = t_H(\min(|\phi|, \preceq_H))$. Then $*$ satisfies postulates (H*1)–(H*8) and (Acyc).*

Proof Summary: The proof is much the same as for classical AGM revision. Horn compliance is required to show that (H*7) and (H*8) hold. (Acyc) basically follows the informal rationale of the postulate given above. □

Theorem 3 *Let $* : \mathcal{H} \times \mathcal{L}_H \mapsto \mathcal{H}$ be a function satisfying postulates (H*1)–(H*8) and (Acyc). Then for fixed theory H, there is a faithful ranking $\preceq_H$ on $\mathcal{M}$ such that $\preceq_H$ is Horn compliant and $H * \phi = t_H(\min(|\phi|, \preceq_H))$.*

Proof Summary: For interpretations m_1, m_2, define $f(m_1, m_2)$ to be the logically strongest Horn formula with models m_1, m_2. It may be that $\{m_1, m_2\} \subset |f(m_1, m_2)|$. Nonetheless we can define $m_1 \preceq_H^\circ m_2$ iff $m_1 \in |H * f(m_1, m_2)|$. The relation $\preceq_H^\circ$ is not transitive, but (Acyc) guarantees that $\preceq_H^\circ$ can be extended safely to a total preorder $\preceq_H$ that is shown to be Horn compliant and that exactly captures the original revision function. □

We conclude this section with some comments on Horn compliance. Firstly notice that, in certain cases, Horn compliance excludes some popular preorders like, for example, those proposed in [Dalal, 1988]. To see this, assume that $\mathcal{P} = \{p, q\}$, and $H = \mathcal{C}n_H(\{p, q\})$. Then according to Dalal, $pq \prec_H p\overline{q} \approx_H \overline{p}q \prec_H \overline{pq}$, which however is not a Horn compliant preorder.

This raises the question of whether Horn compliant preorders are at all possible. Luckily the answer is affirmative. For any Horn theory H, one can construct a Horn compliant preorder $\preceq_H$ as follows: start with the H-worlds, and then attach any linear order on the non-H-worlds. The resulting preorder can be easily shown to be Horn-compliant (of course there are also other ways to build Horn-compliant preorders – see the full paper for more details).

7 Discussion

It can be noted that while the above results are expressed in terms of Horn theories, they represent an *extension* rather than a *modification* of the AGM approach. That is, we could redefine (AGM-style) revision in the context of a logic that contains Horn derivability. Postulates would consist of (H*1)-(H*8) and (Acyc) as before, while the construction would be in terms of Horn compliant faithful rankings. This would subsume classical AGM revision: classical propositional logic obviously is stronger than Horn logic. In classical propositional logic the notion corresponding to Horn compliance is trivial, since (over a finite alphabet) for any formula ϕ of propositional logic, $\min(|\phi|, \preceq)$ is definable via a formula of propositional logic. On the postulational side, as indicated, (Acyc) is derivable from the other postulates if one has classical propositional logic.

These results also suggest several (in our opinion) very interesting directions for future work. First, we argued that AGM revision can be extended by weakening the underlying logic to that of Horn logic. This raises the issue of whether the overall framework can be broadened to subsume other weakened inference relations, while maintaining the overall AGM character, as reflected in the standard AGM postulates.

The area of belief change in Horn theories is in the process of being mapped out. Other research has characterised Horn contraction, while the present paper has addressed revision. However, there has been no work that we are aware of in linking the areas of Horn contraction and revision. Moreover, the constructions in Horn contraction have focussed on the standard contraction constructions of remainder sets and epistemic entrenchment, while the present work has used the standard revision construction of a faithful ranking. Hence there is also a disconnect in the underlying formal characterisations. Consequently, research on linking Horn contraction and revision would help shed further light on the foundations of belief change.

Last, there is burgeoning interest in addressing belief change in description logics (see [Qi and Yang, 2008] for instance) or in analogous areas such as ontology evolution. Given that a Horn clause may also find interpretation as a subsumption, by mapping a rule $p \wedge q \rightarrow r$ to a subsumption of the form $P \sqcap Q \sqsubseteq R$, the present approach may also shed light on approaches to revision in description logics.

References

[Alchourrón *et al.*, 1985] C.E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Booth *et al.*, 2009] R. Booth, T. Meyer, and I. J. Varzinczak. Next steps in propositional Horn contraction. In *Proceedings of IJCAI-09*, CA, 2009.

[Dalal, 1988] M. Dalal. Investigations into theory of knowledge base revision. In *Proceedings of AAAI-88*, pages 449–479, 1988.

[Delgrande and Wassermann, 2010] J. P. Delgrande and R. Wassermann. Horn clause contraction functions: Belief set and belief base approaches. In *Proceedings of KR-10*, pages 143–152, Toronto, 2010.

[Delgrande, 2008] J.P. Delgrande. Horn clause belief change: Contraction functions. In *Proceedings of KR-08*, pages 156–165. AAAI Press, 2008.

[Eiter and Gottlob, 1992] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57(2-3):227–270, 1992.

[Flouris *et al.*, 2004] G. Flouris, D. Plexousakis, and G. Antoniou. Generalizing the AGM postulates: Preliminary results and applications. In *Proceedings of NMR-04*, pages 171–179, Canada, 2004.

[Gärdenfors, 1988] P. Gärdenfors. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Cambridge, MA, 1988.

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

[Katsuno and Mendelzon, 1991] H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

[Langlois *et al.*, 2008] M. Langlois, R.H. Sloan, B. Szörényi, and G. Turán. Horn complements: Towards Horn-to-Horn belief revision. In *Proceedings of AAAI-08*, Chicago, Il, 2008.

[Liberatore, 2000] Paolo Liberatore. Compilability and compact representations of revision of Horn knowledge bases. *ACM Transactions on Computational Logic*, 1(1):131–161, 2000.

[Peppas, 2008] P. Peppas. Belief revision. In *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, 2008.

[Qi and Yang, 2008] G. Qi and F. Yang. A survey of revision approaches in description logics. In *Proceedings of RR-08*, pages 74–88. Springer-Verlag, 2008.

[Zhuang and Pagnucco, 2010] Z. Q. Zhuang and M. Pagnucco. Horn contraction via epistemic entrenchment. In *JELIA-10*, pages 339–351. Springer Verlag, 2010.