A

General Belief Revision

JAMES P. DELGRANDE, Simon Fraser University
PAVLOS PEPPAS, University of Patras and University of Technology Sydney
STEFAN WOLTRAN, Technische Universität Wien

In Artificial Intelligence, a key question concerns how an agent may rationally revise its beliefs in light of new information. The standard (AGM) approach to belief revision assumes that the underlying logic contains classical propositional logic. This is a significant limitation, since many representation schemes in AI don't subsume propositional logic. In this paper we consider the question of what the minimal requirements are on a logic, such that the AGM approach to revision may be formulated. We show that AGM-style revision can be obtained even when extremely little is assumed of the underlying language and its semantics; in fact, one requires little more than a language with sentences that are satisfied at models, or possible worlds. The classical AGM postulates are expressed in this framework and a representation result is established between the postulate set and certain preorders on possible worlds. To obtain the representation result, we add a new postulate to the AGM postulates, and we add a constraint to preorders on worlds. Crucially, both of these additions are redundant in the original AGM framework, and so we *extend*, rather than *modify*, the AGM approach. As well, iterated revision is addressed and the Darwiche/Pearl postulates are shown to be compatible with our approach. Various examples are given to illustrate the approach, including Horn clause revision, revision in extended logic programs, and belief revision in a very basic logic called *literal revision*.

CCS Concepts: •**Computing methodologies** → **Nonmonotonic, default reasoning and belief revision;** *Reasoning about belief and knowledge;*•**Theory of computation** → *Logic;*

General Terms: Theory

Additional Key Words and Phrases: knowledge representation, belief revision, nonclassical logic

ACM Reference Format:
J. ACM V, N, Article A (January YYYY), 35 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

In all but the simplest of circumstances and environments, an agent will have to alter its beliefs to take into account new information. Such new information may fill in gaps in the agent's beliefs, or it may correct an agent's incorrectly-held belief. So, very broadly, in this process of *belief revision* an agent will receive information about the domain; this information may or may not conflict with the agent's beliefs; but one way or another this new information is to be incorporated into the agent's beliefs.

However, this process of incorporating new beliefs into an agent's belief corpus is not arbitrary, but rather is bound by various commonsense principles. For example

The first author gratefully acknowledges financial support from the Natural Sciences and Engineering Research Council of Canada. The third author was supported by the Austrian Science Fund (FWF): P25521-N23, Y698-N23.
Authors' address: James Delgrande, School of Computing Science, Simon Fraser University, Burnaby, B.C., Canada, V5A 1S6, e-mail: jim@cs.sfu.ca. Pavlos Peppas, Department of Business Administration, University of Patras, Patras 26500, Greece, and Centre for AI, School of Software, University of Technology Sydney, NSW 2007, Australia, email: pavlos@upatras.gr. Stefan Woltran, Fakultät für Informatik, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria, e-mail: woltran@dbai.tuwien.ac.at.

DOI: http://dx.doi.org/10.1145/0000000.0000000

assume that the new information is given by a formula ϕ. Then if the goal of revision is to incorporate this information, following the process of revision, ϕ should appear among the agent's beliefs. One possibility would be to simply add ϕ to the agent's beliefs; in such a case ϕ would indeed be in the resulting belief set. However, ϕ might conflict with the agent's prior beliefs, and if this was the case, the agent would fall into inconsistency. So another reasonable principle is that an agent's beliefs should be consistent after revision by a formula ϕ (unless ϕ itself is inconsistent). This in turn requires that an agent may also have to remove some beliefs in a revision. One possibility in this case would be to remove *all* of the agent's prior beliefs. However this is clearly far too drastic, and so one would want to stipulate that in some fashion the agent retains as many of its old beliefs as consistently possible.

The upshot is that belief revision (and more broadly, *belief change* as a whole) is an area with difficult and subtle problems. Research in this area can be regarded as beginning with the seminal work of Alchourrón, Gärdenfors, and Makinson [Alchourrón et al. 1985] (see also [Gärdenfors 1988]), resulting in what has come to be known as the *AGM approach*. In this framework, the focus was on the belief change operators of *revision*, in which an agent alters its beliefs to incorporate a new formula, and *contraction*, in which an agent reduces its stock of beliefs so that a given formula is not believed. In this approach, as we review in the next section, *postulates* are provided, which express principles that arguably should govern any rational change operator, as well as formal *constructions* that express how one may build a specific change operator. These two notions are tied together by providing *representation results* that prove that the class of change functions captured by a set of postulates is exactly that given in the corresponding construction. The resulting framework has, since its inception, been the central pillar and focus of research in belief change [Peppas 2008].

A key assumption of the AGM approach, and the point that will concern us here, is that the logic underlying the agent's knowledge base at least contains classical propositional logic. On the one hand this seems to be a quite reasonable assumption; after all, classical propositional logic is often seen as being very basic and lacking in expressivity. However, on the other hand, this apparent simplicity is deceptive. The best propositional reasoners take exponential time in the worst case, and general consensus is that this won't change (given that the satisfiability problem of propositional logic is NP complete). As well, full classical negation and disjunction are sometimes seen as being undesirable, particularly when moving toward a first-order formalism. Yet other approaches employ nonclassical notions of, for example, negation, and resist an easy comparison with classical propositional logic. What this means is that in Artificial Intelligence in general, and Knowledge Representation in particular, there has been extensive work on representation formalisms that don't subsume classical propositional logic, including work in Horn clause reasoners, description logics, extended logic programs, and others. And so what this also means is that the AGM framework for belief change is inapplicable in these approaches.

This has led to the study of AGM-style belief change with respect to systems that do not subsume propositional logic. The focus of much of this work has been on belief change in Horn theories, including belief contraction and belief revision. In particular, [Delgrande and Peppas 2015] reconstructs full AGM-style belief revision in the context of propositional Horn theories. As a result, while the AGM approach assumes that the underlying logic subsumes classical propositional logic, it is clear that this is not a *necessary* condition.

In the present paper, we consider the question of just what are the minimal restrictions that need to be placed on a logic in order to be able to define AGM-style revision in that logic. It proves to be the case that very little needs to be assumed in order to provide a sufficient setting for defining revision. Essentially we assume that we have a *language* (although we assume nothing about the structure of the language), and that we have a set of *models*, and a function that specifies, for each formula, the set of models that *satisfy* the formula.

While we work within a very general setting, we show that nonetheless a fundamental semantic characterisation of belief revision based on the notion of a *faithful ranking* [Katsuno and Mendelzon 1991] can be suitably defined in our approach. However, in the general case, an additional constraint that we call *regularity* is imposed on faithful rankings. Notably this condition is redundant when the underlying logic subsumes classical propositional logic. As well, we provide a set of postulates that corresponds to the standard AGM revision postulates. Similar to faithful rankings, an additional postulate, that we call (Acyc), is required. Again, this postulate is redundant in the case of propositional logic. These two characterisations are proven to be equivalent via a representation result that shows that the class of general revision functions conforming to the augmented postulate set is the same as those expressible by regular faithful rankings. We also consider iterated belief revision, showing that the Darwiche-Pearl postulates [Darwiche and Pearl 1997] are consistent with the general approach to revision.

Subsequently, various specific instances of the approach are discussed. Classical propositional logic and Horn clause logic are first viewed, briefly, as instances of this approach. Following this we review belief revision in various classes of *extended logic programs*. Last, we develop an approach called *literal revision* where the underlying formal system is perhaps the simplest approach that could reasonably be called a logic. In this last system, an agent's belief set is equivalent to a set of propositional literals, and the task is to consistently revise by a formula expressed as a conjunction of literals. Since the defined system satisfies our set of assumptions, it follows that a full revision function can be defined, even in such an impoverished system.

These results are interesting for several reasons:

— Foremost, the AGM framework is extended to include any system that might reasonably be called a logic.[1] As described above, systems that do not subsume classical propositional logic are playing an increasing role in knowledge representation. Notable areas of interest include, among others, description logics [Baader et al. 2007] and the answer set approach to logic programming [Gelfond and Lifschitz 1988; Brewka et al. 2011]. The present approach then implicitly defines AGM-style belief revision within such approaches.

— Consequently, our results provide a guide to the formulation of specific revision operators in non-classical logics, including description logics, modal logics, many-valued logics, etc.

— In addition, our results provide a significant short cut in developing representation results: For any logic, once the language, model theory, and a notion of regularity are suitably defined, our representation result applies to that logic.

— Last, the approach sheds light on the foundations of belief change. On the one hand, it demonstrates that the AGM framework, as least as regards revision, is much more widely applicable than previously believed. On the other hand, our results indicate that when the underlying logic is weaker than classical propositional logic, revision and contraction become distinct, independent change operations.

The next section provides background and motivation: the AGM approach is briefly reviewed and, following this, issues that may arise in inferentially-weak systems are discussed. Section 3 covers previous work concerning belief change in such systems.

[1] Or at least any logic defined in terms of a model theory and satisfaction relation. We return to this point in Section 7 once the approach has been presented.

Section 4 defines the formal framework, expresses the AGM approach in this framework, and provides a representation result. Section 5 addresses iterated revision; while the next section describes various instantiations of the approach. Section 7 discusses issues raised by the approach, and the final section gives a brief conclusion.

2. BELIEF CHANGE

2.1. The AGM Approach

The AGM approach to belief change [Alchourrón et al. 1985; Gärdenfors 1988; Peppas 2008] studies change operators at the *knowledge level*, independent of syntactic issues such as how information is to be represented in a knowledge base. It is assumed that the underlying logic contains classical propositional logic. An agent's beliefs are modelled by a deductively closed set of formulas, called a *theory* or a *belief set*. Thus a belief set is a set of formulas K such that $K = Cn(K)$, where $Cn(K)$ denotes the closure of K under a consequence operator that subsumes classical logical consequence. Belief revision is modelled as a function from a belief set K and a formula ϕ to a belief set K' such that ϕ is believed in K', that is, $\phi \in K'$. If ϕ is consistent with K (that is to say, $\neg\phi \notin K$), then it is simply added to K and the revision is given by $Cn(K \cup \{\phi\})$. This "adding" of a formula to a belief set is usefully considered as a distinct operation, called *expansion*; it is defined by:

$$K + \phi = Cn(K \cup \{\phi\}).$$

The interesting case in revising by a formula ϕ is when ϕ is inconsistent with the agent's belief set K. Since ϕ is to be believed in the revised knowledge base, this means that (assuming that ϕ is consistent), some formulas must be dropped from K before ϕ can be consistently added. In general, there will be many ways in which K can be reduced so that ϕ can be consistently added — for example, one alternative is to drop all formulas in K.

Clearly such a revision function would in general be too drastic. This leads to the consideration that a revision function isn't arbitrary, but rather is assumed to be guided by various *rationality* criteria. A key assumption is that of *informational economy*, that when revising beliefs, we want to retain as much as possible of our prior beliefs. As a consequence, a rational belief revision operator is one in which (among other things) a belief set K undergoes *minimal change* in order to incorporate a formula for revision. Of course, a notion such as *minimal change*, at least as an English phrase, is informal, and so part of the task of specifying a revision function, only partly addressed by the AGM approach, is to formally specify what is meant by such change.

The AGM framework does not determine a unique revision function, but rather it specifies constraints that a rational change function should satisfy; beyond these constraints the approach offers no advice as to how a specific operator should be constructed. The overall methodology for studying belief change is to approach a change operator from two directions: On the one hand, a set of *postulates* can be given to characterise those properties that any rational change operator should satisfy. On the other hand, a *construction* can be given to formally characterise the class of instances of that operator. Then, ideally, the two approaches are shown to coïncide via a *representation result*, showing that the approaches capture the same class of operators.

The AGM postulates for revision can be expressed as follows. Below, $\equiv_{PC}$ and $+_{PC}$ stand for logical equivalence and expansion, respectively, in classical propositional logic.

(K*1) $K * \phi = Cn(K * \phi)$
(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 $(K * \phi) +_{PC} \psi \subseteq K * (\phi \wedge \psi)$

The first six postulates are called the *basic postulates*, while the last two are called the *extended postulates*. The first two postulates assert that the result of revising K by ϕ yields a belief set (K*1) in which ϕ is believed (K*2). (K*3) and (K*4) assert that if a formula for revision is consistent with a belief set K, then revision consists of the expansion of K by ϕ. (K*5) says that unless ϕ is inconsistent, $K * \phi$ is consistent. (If ϕ is inconsistent, then (K*2) requires the result to be inconsistent.) (K*6) asserts that revision is independent of the syntactic form of the formula for revision. The last two postulates deal with the relation between revising by a conjunction and expansion: whenever consistent, revision by a conjunction corresponds to revision by one conjunct and expansion by the other. Postulates (K*3) and (K*4), and (K*7) and (K*8), can be seen as expressing that in a revision as little information is removed from K as is consistently possible. Further motivation for these postulates can be found in [Gärdenfors 1988; Peppas 2008]. We shall call any function $*$ that satisfies (K*1) – (K*8) an *AGM revision function*.

Adam Grove [Grove 1988] provided a possible worlds characterisation of revision functions, based in turn on David Lewis's *system of spheres* [Lewis 1973]. We shall deviate slightly from Grove's terminology and instead of systems of spheres we shall be working with total preorders over propositional interpretations, or *possible worlds*.[2]

First, recall that a preorder $\preceq$ (here, over possible worlds) is a reflexive, transitive, binary relation on the set of possible worlds $\mathcal{M}$. The relation $\preceq$ is called *total* iff for all $w_1, w_2 \in \mathcal{M}$, either $w_1 \preceq w_2$ or $w_2 \preceq w_1$.

For a subset S of $\mathcal{M}$, we say that a world w is *minimal* in S with respect to $\preceq$ iff $w \in S$ and for all $w' \in S$, $w' \preceq w$ entails $w \preceq w'$. We denote the set of minimal elements of S with respect to $\preceq$ by $\min(S, \preceq)$:

$$min(S, \preceq) = \{w \in S \mid \text{ for all } w' \in S, \text{ if } w' \preceq w \text{ then } w \preceq w'\}.$$

We say that a preorder $\preceq$ over $\mathcal{M}$ is *faithful* with respect to a theory K iff [3]

(F1) $\preceq$ is total
(F2) if $[K] \neq \emptyset$, then $\min(\mathcal{M}, \preceq) = [K]$.

Intuitively, $w_1 \preceq w_2$ if w_1 is at least as plausible as w_2. Grove also provides an additional condition, corresponding to the *limit assumption* of [Lewis 1973]:

(S4) for any consistent sentence ϕ, $\min([\phi], \preceq) \neq \emptyset$.

This condition is only needed in the infinite case; since we will assume that the set of possible worlds is finite, we will not require such a condition. Grove then provides the following representation result (modulo the different terminology), where $t(S)$ is the set of formulas of classical logic true in the set of possible worlds S:

THEOREM 2.1 ([GROVE 1988]). *Let K be a theory and * a revision function. Then $*$ satisfies postulates (K*1) – (K*8) at K iff there exists a preorder $\preceq$ over $\mathcal{M}$ that is*

[2]The two constructs are equivalent for the purpose of constructing a revision function. This was already noted by Katsuno and Mendelzon [Katsuno and Mendelzon 1991], who also first suggested expressing a possible worlds characterisation in terms of a total preorder rather than a Lewis-style system of spheres.
[3]We use $[\cdot]$ to represent the set of possible worlds associated to a theory K or a formula ϕ; a formal definition is given in Section 4. For the time being, $[\cdot]$ can be thought of as a set of classical models.

faithful to K and satisfies (S4), and such that

$$K * \phi \;=\; t(\min([\phi], \preceq)). \tag{1}$$

Thus the revision of K by ϕ is characterised by the set of those models of ϕ that are most plausible according to the agent.

Another form of belief change in the AGM approach is called belief *contraction*. Assume that $\phi \in K$ and that ϕ is not a tautology. In contracting the formula ϕ from the belief set K, denoted $K - \phi$, the agent no longer believes ϕ (while not necessarily believing $\neg\phi$). That is, if ϕ is not a tautology, then one requires that $\phi \notin K - \phi$. Informally, contraction is thought of as being a more basic (or fundamental) operation than revision, since in contraction an agent's beliefs can only decrease, while in revision in the interesting case an agent's beliefs change.

However, it proves to be the case that in the standard AGM approach, revision and contraction functions are interdefinable. Given a contraction function $-$, one can define a revision function by the so-called Levi identity:

$$K * \phi = (K - \neg\phi) +_{PC} \phi. \tag{2}$$

Analogously, given a revision function $*$, one can define a contraction function via the Harper identity:

$$K - \phi = K \cap K * \neg\phi. \tag{3}$$

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

So, to conclude, our interests lie with AGM-style belief revision, which we have introduced here, and with the goal of extending it to arbitrary logics. It is worth briefly discussing some notions that we will not be considering. First, although we will allude to belief contraction, it is not our focus, and we do not consider the interesting question of AGM-style belief contraction. Second, an intuition underlying belief revision is that the agent is receiving information about some domain, but where the domain itself is unchanging. An alternative intuition is that an agent receives information about a change in the domain; this leads to a different class of operators, called *belief update* [Katsuno and Mendelzon 1992]. It seems likely that the techniques developed here could be applied without difficulty to belief update, although we do not do so in this paper. Last, we mentioned that the AGM framework is described at the *knowledge level* wherein presumably-irrelevant syntactic concerns are ignored, and wherein an agent's beliefs are given by a belief set. An alternative is to take syntax into account. In this case, distinct but logically-equivalent knowledge bases may behave differently under revision by the same formula. This leads to the notion of *belief base* revision [Hansson 1999], which again we do not consider here.

2.2. AGM Revision and Classical Propositional Logic

In this subsection, we consider the question of why AGM-style revision requires that the underlying logic subsumes classical propositional logic. We do this by informally surveying problems that arise in attempting to define an AGM-style belief change operator in an inferentially-weak system, where by "inferentially weak" we mean not having the expressivity of classical propositional logic. (So this term includes both fragments of classical propositional logic, as well as those nonclassical logics that do not subsume classical propositional logic.) In surveying problems that may arise, we focus on Horn clause theories,[4] and refer to other approaches as appropriate. While we refer

[4]Some of this material is drawn from [Delgrande 2008; Delgrande and Wassermann 2013; Delgrande and Peppas 2015].

to Horn clause theories to illustrate the problems that may arise, it should be clear that such problems may be expected to occur in other weak systems.

We begin with a review of some basic terminology. In classical logic, a *clause* is a disjunction of literals. A *Horn clause* is a clause with at most one positive literal. A *definite clause* is a clause with exactly one positive literal. A Horn clause $\neg a_1 \vee \neg a_2 \vee \cdots \vee \neg a_n \vee a$ can be perspicuously written as an implication involving atoms only: $a_1 \wedge a_2 \wedge \cdots \wedge a_n \Rightarrow a$. A clause with no positive literal can be written $a_1 \wedge a_2 \wedge \cdots \wedge a_n \Rightarrow \perp$. A *Horn formula* is just a conjunction of Horn clauses, while a *Horn theory* is a set of Horn formulas closed under a suitable notion of Horn derivability [Delgrande and Peppas 2015]. An example of a formula that is not expressible in a Horn theory is $p \vee q$.

Models of Horn formulas are distinguished by the fact that they are closed under intersection of positive atoms in an interpretation. That is, if w_1 and w_2 are models of ϕ expressed as a set of atoms then $w_1 \cap w_2$ is also a model of ϕ. The converse is also true; that is, 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 the models of ϕ are W; and if a set of models W is not closed under intersection then it is not representable in a Horn theory. For example consider the formula $p \equiv \neg q$ over the alphabet $\{p, q\}$. The models of $p \equiv \neg q$ are $\{p\}$ and $\{q\}$, and $p \equiv \neg q$ is not expressible using Horn clauses. If, along with $\{p\}$ and $\{q\}$, we include the model $\emptyset$ ($= \{p\} \cap \{q\}$), then the resulting set of models $\{\{p\}, \{q\}, \emptyset\}$ corresponds to the Horn clause $\neg p \vee \neg q$.

Consider now the issue of defining AGM-style revision with respect to Horn clause knowledge bases. To begin, it can be observed that the simpler case, of definite clauses, is trivial. First, any set of definite clauses is consistent (for example, just assign the value *true* to every atom). Hence, to revise a definite clause knowledge base by a definite clause, one just adds the formula for revision to the knowledge base and takes the deductive closure (suitably defined for definite clauses). However, a set of Horn clauses may be inconsistent (for example, p and $p \Rightarrow \perp$ are together inconsistent) and so revision is nontrivial in the Horn case. These observations suggest that, for revision to be meaningful, a logic must have some notion of inconsistency. These observations also suggest that care must be taken when non-classical negation is encountered, as may be found for example in an *extended logic program*; we will encounter an example in the next section.

So with Horn clauses it would seem that we have a system, weaker than classical propositional logic, that might nonetheless have revision defined according to the standard AGM definitions. And indeed one can easily define faithful rankings and a set of AGM-like postulates in terms of Horn clauses: Interpretations of Horn formulas are, after all, just the interpretations of classical propositional logic. And a set of AGM-like postulates, rephrased in terms of a Horn-logic consequence relation, is straightforwardly specifiable. However, if one does this, it proves to be the case that the standard representation results fail. Thus, it is possible to define an operator that satisfies all the revision postulates (restricted to Horn formulas) but for which there is no corresponding faithful preorder.[5] Similarly, one can specify a faithful ranking for which the operator defined by (1) does not satisfy the (Horn AGM) revision postulates; see [Delgrande and Peppas 2015] for a counterexample. Essentially these problems arise from the relative inexpressiveness of Horn theories: full disjunction is missing, as is full negation.[6]

[5]Very informally the problem is the following [Delgrande and Peppas 2015]: In the case of propositional logic, given a revision operator $*$, one determines that the least ψ worlds should be ranked "no lower than" the least ϕ worlds exactly when $K * (\phi \vee \psi) \subseteq K * \phi$. But in the Horn case, if ϕ and ψ are Horn formulas, $\phi \vee \psi$ may nonetheless not be a Horn formula, and so this construction technique is inapplicable.

[6]That is, for a Horn formula ϕ, the negation $\neg\phi$ may not be Horn-definable.

Similar issues may be expected to arise in other inferentially-weak systems. For example, many description logics [Baader et al. 2007] lack full disjunction or negation. In perhaps the simplest description logic $\mathcal{EL}$, there is no notion of inconsistency and so revision is trivial in this case. However, all description logics have a concept $\top$ that is true of all individuals, and most have another concept $\bot$ that is true of none. Given the standard (Tarskian) assumption that there is at least one individual, a description logic knowledge base is inconsistent just if $\top \sqsubseteq \bot$ is entailed; that is, the *top* concept is subsumed by the *bottom* concept.

Thus in inferentially-weak logics, the direct adaptation of the AGM approach to revision may be anticipated to be problematic. A plausible alternative is to first define a suitable contraction function, and then define revision via the Levi Identity (2). However in general this strategy is also problematic. Consider again Horn theories. To begin with, there is more than one way that one may define contraction. Informally, for a contraction $K - \phi$ there are the two notions: $K - \phi$ can be defined as a subset of K that does not entail ϕ, or $K - \phi$ can be defined as a subset of K that is consistent with ϕ; in symbols:

(1) If $\phi \in K$ then one requires $\phi \notin K - \phi$.
(2) If $K \cup \{\phi\}$ is inconsistent then one requires that $(K - \phi) \cup \{\phi\}$ is consistent.

Note that in the second case, if the underlying logic contains propositional logic, we would have $\neg\phi \in K$, and so the AGM contraction would in fact be expressed as $K - \neg\phi$; in (2), we write $K - \phi$ instead, because in an arbitrary logic $\neg\phi$ may not be a formula.

These two conceptions of contraction are easily shown to coincide for propositional logic: For the antecedents, one has

$$\phi \in K \text{ iff } K \cup \{\neg\phi\} \text{ is inconsistent},$$

and for the consequents we have that

$$\phi \notin K - \phi \text{ iff } (K - \phi) \cup \{\neg\phi\} \text{ is consistent}.$$

However, for Horn clause theories these are distinct, simply because if ϕ is a Horn formula, then $\neg\phi$ may not be. (As a simple example, $\neg p \wedge \neg q$ is a Horn formula, whereas $\neg(\neg p \wedge \neg q)$, or $(p \vee q)$, is not.)[7] There has been extensive work in contraction in Horn theories [Delgrande 2008; Booth et al. 2009; Zhuang and Pagnucco 2010; 2011; Booth et al. 2011; Zhuang and Pagnucco 2012; Delgrande and Wassermann 2013]. However, such work either ends up with postulates that differ from the standard AGM set, or else makes use of non-Horn clauses along the route to determining Horn contraction. What this means is that, given the current state of research, it is not clear that an approach to contraction for inferentially-weak logics that follows the AGM approach is possible.

However, there is a more immediate reason why defining revision via contraction may not work, and that is that in general the Levi Identity may fail, and so revision would not then be definable in terms of contraction via this identity. Thus in Horn theories, as well as in weak description logics such as the $\mathcal{EL}$ family, the Levi Identity can't be used since for an arbitrary formula ϕ in one of these approaches, $\neg\phi$ may not be defined.

Informally, these results suggest that in inferentially weak systems, revision and contraction become two distinct operations, in that they are no longer obviously inter-

[7]It is interesting to note that these two formulations for contraction differ in other ways. For example the first makes sense in a system with no notion of inconsistency (such as in definite clauses or the description logic $\mathcal{EL}$) whereas the second does not. Hence the first may potentially be useful in such logics, whereas the second would presumably be inapplicable.

definable. In fact, with contraction, it appears that while some semantic constructions in the AGM approach may be adaptable to weaker logics, others may not be so readily adapted.[8] Moreover, to date the prospects of coming up with a contraction function in such weak systems that satisfies the full AGM postulate set are unclear.

It is of interest that, while a modification of AGM contraction to accommodate other logics is uncertain, our results here show that this is not the case for revision. We show instead that the AGM approach can be adapted to apply in a very wide class of logics. Included in this class is Horn logic, description logics, relevance logics, extended logic programs, and, more broadly, any system that seems to satisfy a very basic notion of *logic*.

3. RELATED WORK

This section reviews related work that has been carried out in belief revision, focussing on what we called "inferentially weak" logics in the previous section. Such work can be considered as belonging to one of two broad groups. The first involves revision in fragments of classical logic, while the second addresses revision in nonclassical logics. We also examine at the end of this section an approach that addresses *distance-based* revision from first principles.

In the first group, perhaps the earliest work in studying revision in a system weaker than classical propositional logic is that described in [Restall and Slaney 1995], where revision in the relevance logic [Anderson and Belnap Jr. 1975] of *first degree entailment*, $\mathbf{E}_{fde}$, is studied. In $\mathbf{E}_{fde}$ a formula may be true or false, as usual, but it may also be both true and false, or neither true nor false. As a result, the so-called *paradoxes of implication*, such as $\phi \land \neg\phi \Rightarrow \psi$, do not hold. Restall and Slaney's work focusses on the semantic constructions, in particular those based on *epistemic entrenchment*, *partial meet*, and *systems of spheres*. In each case it is shown how a construction can be adapted to the 4-valued semantics. In the case of epistemic entrenchment and partial meet, a revision function is obtained from a contraction function via the Levi Identity. Interestingly, in the case of $\mathbf{E}_{fde}$, the Harper Identity fails.

With regards to Horn revision, Zhuang et al. [2013] present a technique for obtaining a Horn revision in terms of contraction. As previously described, the difficulty in defining Horn revision in terms of contraction is that, in employing the Levi Identity, one must deal with the negation of a Horn formula; this, in general, is not Horn. Zhuang et al. circumvent this difficulty by contracting by a sequence of *Horn strengthenings* [Selman and Kautz 1996] of the negation of the formula for revision.

As noted in the previous section, Delgrande and Peppas [2011; 2015] investigate belief revision where the underlying logic is that governing Horn clauses. In this work, the AGM approach is augmented in two ways. First, a further postulate is added to the set of revision postulates. This postulate, in semantic terms, rules out certain undesirable circularities among possible world orderings. Second, a condition is imposed on faithful rankings to exclude certain undesirable orderings (i.e., orderings that would yield a result that is not expressible in Horn logic). Of key importance is the fact that both of these restrictions, while necessary for the Horn case, are redundant in the standard AGM approach. A representation result shows that the class of revision functions captured by these restricted faithful rankings is precisely that given by the (extended) set of Horn revision postulates. Consequently, this work extends AGM revision to the inferentially-weaker Horn case. Moreover it is shown that Horn revision is compatible with other work in revision, including iterated revision and work concerning relevance.

[8] While it is beyond the scope of the present paper, we can note that the approach of *remainder sets* appears to be naturally extendable to the Horn case [Delgrande and Wassermann 2013], while more effort is required to adapt *epistemic entrenchment* to this case [Zhuang and Pagnucco 2010].

The present paper then can be seen as in part extending and generalising these results to arbitrary logics.

More recently, belief revision in other fragments of propositional logic, including Krom and affine formulas, has been addressed in [Creignou et al. 2014]. However, the main focus of that work is not concerned with representation results. Instead, the authors propose to *adapt* known revision operators by means of a certain post-processing and then study the limits of this approach in terms of satisfaction of the postulates. One of the main results of that paper is that in this framework it is not possible to keep Postulate (K*8) satisfied.

[Delgrande et al. 2013] addresses AGM-style revision in logic programs under the answer set semantics [Gelfond and Lifschitz 1988; Brewka et al. 2011]. This approach makes use of a standard, monotonic (albeit non-classical) model theory based on the notion of *SE-models* [Turner 2003].[9] Using techniques from [Delgrande and Peppas 2015], it is shown how classical AGM-style revision can be extended to various classes of logic programs by means of SE-models. That is, the AGM postulates are rephrased to refer to logic programs; a semantic construction for revision operators is given based on orderings over SE models; and then a representation result shows that these approaches coincide. See also [Schwind and Inoue 2013] for a related approach.

Recently, AGM-style revision has also gained interest in the field of abstract argumentation. Here, the outcome of so-called argumentation frameworks [Dung 1995] is revised on the level of *extensions*, see e.g. [Coste-Marquis et al. 2014]. In order to guarantee that the result of a revision remains expressible as an argumentation framework, similar issues as recognized for Horn revision come into play; a recent paper [Diller et al. 2018] shows how AGM revision of argumentation frameworks needs to be defined such that it is guaranteed to work properly within the restricted language of argumentation frameworks. In these papers, the outcomes of argumentation frameworks are treated like models of propositional formulas which obey certain restrictions. Using a completely different recent approach, [Baumann and Brewka 2015] develops a weaker logic in order to study revision of argumentation frameworks; this approach is closer to the SE-model based revision in logic programming (where, likewise, a weaker monotonic logic underlying the nonmonotonic semantics of logic programming is taken as a base logic for formalizing belief change).

Regarding belief revision in general, [Gabbay et al. 2008] tackles a somewhat different problem than that addressed here. For a non-classical logic whose semantics can be axiomatised in first-order logic, they show how a revision operator for classical logic can be used to define a revision operator in the non-classical logic. This is done by translating a belief set and formula expressed in the non-classical logic, along with an axiomatic specification of the logic, into classical logic. The (standard, AGM) revision operator is applied to the resulting theory; and the results are subsequently re-expressed in the original logic. The overall result then is a methodology for "exporting" an AGM revision operator in classical logic to non-classical logics. This methodology has recently been applied also to Horn logic [Brewka et al. 2016].

Ribeiro and Wassermann [2014] consider revision in non-classical logics. Their approach is to begin with the basic AGM postulates, and then consider additional postulates (in place of (K*7) and (K*8)) that would express a notion of minimality. Two constructions are provided, based on the (contraction) constructions of remainder sets and kernels, and representation results are provided making use of an additional postulate

[9]Without going into details, an SE-model of a logic program P is an ordered pair of classical models, satisfying certain constraints and related to the classical models of P.

of *relevance*[10] on the one hand, and *core retainment* on the other. Last, [Wassermann 2011] provides a survey of research on belief change in non-classical logics.

The work of Lehman, Magidor, and Schlechta [2001] (see also [Schlechta 2004] for further details and in a wider context) also deals with foundational issues in belief revision. However, the focus of Lehman, Magidor, and Schlechta (LMS for short) is on distance-based revision: they assume that a fixed, global, distance is given between all pairs of possible worlds. Once such a distance is given, an AGM revision function is implicitly specified for every set of possible worlds; as well, the results of iterated revision are automatically determined by such a distance function. LMS begins by considering the notion of "distance" broadly, as an abstract, algebraic concept. Once such an algebraic theory is developed, it is then applied in the specific instance of distances between sets of models. Two cases are considered. First, the notion of distance is not necessarily symmetric; this is developed with respect to finite sets of models. Second they consider the case where distance is a symmetric relation; this case allows sets of models of infinite cardinality.

There are some similarities between the current approach and that of LMS. Notably, both have a postulate schema, called (Loop) in LMS and (Acyc) here, that rules out cycles of $<$ relations between worlds; these schemas in turn go back to the (Loop) condition in [Kraus et al. 1990]. As well, in both approaches not all sets of models are representable by a set of formulas. In LMS this is a result of allowing infinite theories for symmetric distances; in the present approach, this results from dealing with non-classical logics and fragments of classical logic. There are also major differences: LMS gives a *restriction* of the AGM approach, in that certain revision functions are ruled out in their approaches. The present approach is a *generalisation* of the AGM framework, in that revision is extended to arbitrary logics. The present approach is also more general, since nothing is assumed about the underlying language, whereas LMS makes explicit use of disjunction. On the other hand, we note that the LMS approach permits logics with infinitely many models (though confined to symmetric distances between models), whereas our approach is limited to logics with finitely many possible worlds.

4. THE APPROACH

In this section we present our approach. The first subsection defines the general framework, while the next subsection expresses the AGM postulates in this framework. The last subsection provides a representation result.

4.1. Building the Framework

Our framework is built from three primitive entities:

— A nonempty (possibly countably infinite) *language* $\mathcal{L}$. The elements of $\mathcal{L}$ are called *sentences* or, equally, *formulas*. We shall use the last few letters of the Greek alphabet, like ϕ, χ, ψ, ..., to denote sentences, and the first few upper-case letters of the English alphabet, like A, B, ..., to denote sets of sentences. Nothing is assumed of the internal structure of sentences (not even the Boolean connectives).

— A finite set $\mathcal{M}$ containing at least two members, the elements of which are called *possible worlds* or simply *worlds*. Worlds will be denoted with the last few lower-case letters of the English alphabet, like, r, w, ... Once again, nothing is assumed of the internal structure of worlds.

— A function f from $\mathcal{L}$ to $2^{\mathcal{M}}$. For a sentence $\phi \in \mathcal{L}$, we often write $[\phi]$ as an alternative to $f(\phi)$.

[10] For a thorough discussion of this, and other proposed postulates, we refer the reader to [Hansson 1999].

With the above three primitive entities we gradually develop the full framework. Let w be any world in $\mathcal{M}$, ϕ any sentence in $\mathcal{L}$, and S an arbitrary set of worlds, that is, $S \subseteq \mathcal{M}$. We say that w ***satisfies*** a sentence ϕ, denoted $w \models \phi$, iff $w \in [\phi]$. Similarly, we say that S ***satisfies*** ϕ, denoted $S \models \phi$, iff for every $w \in S$ we have $w \models \phi$. Moreover we define

$$t(S) = \{\phi \in \mathcal{L} \mid S \models \phi\}.$$

It can be noted that, by definition, $\emptyset \models \phi$ for any $\phi \in \mathcal{L}$, and therefore $t(\emptyset) = \mathcal{L}$.

Let $A \subseteq \mathcal{L}$ be an arbitrary set of sentences. We define $[A]$ to be the set of worlds

$$[A] = \{w \in \mathcal{M} \mid \text{ for all } \phi \in A, w \models \phi\}.$$

We shall say that a world w satisfies A, denoted $w \models A$, iff $w \in [A]$. Observe that by definition $[\emptyset] = \mathcal{M}$. We shall say that A is *consistent* iff $[A] \neq \emptyset$. We say that a set of sentences B is *consistent with* A iff $A \cup B$ is consistent. Two sets of sentences $A, B \subseteq \mathcal{L}$ are said to be *equivalent*, denoted $A \equiv B$ iff $[A] = [B]$. For $\phi, \psi \in \mathcal{L}$, we shall often write $\phi \equiv \psi$ as an abbreviation of $\{\phi\} \equiv \{\psi\}$. We define the *closure* of a set of sentences A, denoted $Cn(A)$, to be the set

$$Cn(A) = \{\phi \in \mathcal{L} \mid [A] \subseteq [\phi]\}.$$

A is said to be a *theory* iff $A = Cn(A)$. Finally, for two sets of sentences A, B, by $A + B$ we denote the set

$$A + B = Cn(A \cup B).$$

Up to now we have made no assumptions about the primitive ingredients $\mathcal{L}$, $\mathcal{M}$, and f of our framework. To proceed further however we impose a simple restriction:

(Expr) For any two distinct worlds $w, w' \in \mathcal{M}$, there exists a sentence ϕ such that $w \models \phi$ and $w' \not\models \phi$.

This restriction requires that the language is expressive enough to distinguish between any two possible worlds. (Expr) has two important consequences. First, it ensures that no world satisfies fewer sentences (set-theoretically) than some other world. That is, if every sentence satisfied by possible world w was also satisfied by possible world w', this would mean that there was no sentence ϕ such that $w \models \phi$ and $w' \not\models \phi$. But this contradicts (Expr). Consequently, all worlds are in a sense "maximal". Second, (Expr) rules out *incoherent* worlds, that is, worlds at which no sentence of $\mathcal{L}$ is satisfied. To see this, assume that w is a world at which no sentence is satisfied. We have assumed that $\mathcal{M}$ contains at least two possible worlds. Let w' be a world distinct from w. Then (Expr) says that there is a sentence satisfied by w, contradicting the assumption that w is incoherent.

The following auxiliary result will be useful in the forthcoming discussion:

LEMMA 4.1. *For any possible world* $w \in \mathcal{M}$, $[t(\{w\})] = \{w\}$.

PROOF. Let w be any possible world in $\mathcal{M}$. Clearly, by the definition of t, $w \in [t(\{w\})]$. Hence what is left to show is that $[t(\{w\})] \subseteq \{w\}$. Consider any possible world $w' \in \mathcal{M}$ such that $w' \neq w$ or, for our purposes, $w' \notin \{w\}$. Then by (Expr), there is a $\phi \in \mathcal{L}$ such that $w \models \phi$ and $w' \not\models \phi$. From $w \models \phi$ it follows that $\phi \in t(\{w\})$. Hence from $w' \not\models \phi$ we derive that $w' \notin [t(\{w\})]$. □

The following small results will be used extensively in the forthcoming discussion. They are stated without a proof since they follow immediately from the definitions:

PROPOSITION 4.2. *For any sets of sentences* $A, B \subseteq \mathcal{L}$, *and sets of worlds* $S, Q \subseteq \mathcal{M}$:

(*1*) $[A] = [Cn(A)]$.

(*2*) $[A \cup B] = [A] \cap [B]$.
(*3*) $t([A] \cap [B]) = A + B$.
(*4*) *If* $S \neq \emptyset$ *then* $t(S)$ *is consistent.*
(*5*) $S \subseteq [t(S)]$.
(*6*) *If* $S \subseteq Q$ *then* $t(Q) \subseteq t(S)$.

We observe that $Cn(.)$ is a Tarskian consequence relation [Tarski 1936], that is, it satisfies the following conditions:

PROPOSITION 4.3. *For any sets of sentences* $A, B \subseteq \mathcal{L}$*:*

(*1*) $A \subseteq Cn(A)$. *Inclusion*
(*2*) *If* $A \subseteq B$ *then* $Cn(A) \subseteq Cn(B)$ *Monotony*
(*3*) $Cn(A) = Cn(Cn(A))$ *Idempotence*

PROOF. The proof of the proposition is obvious, with the possible exception of the containment $\supseteq$ in Part 3.

For this part and direction, from set theory we have that for any set of formulas A that $[A] \subseteq \cap\{[\phi] \mid [A] \subseteq [\phi]\}$. By repeated application of Proposition 4.2.2 we have that $\cap\{[\phi] \mid [A] \subseteq [\phi]\} = [\{\phi \mid [A] \subseteq [\phi]\}]$ and so $[A] \subseteq [\{\phi \mid [A] \subseteq [\phi]\}]$. We observe that $\{\phi \mid [A] \subseteq [\phi]\}$ is just $Cn(A)$ and so we get $[A] \subseteq [Cn(A)]$.

$[A] \subseteq [Cn(A)]$ implies that for every formula ϕ that if $[Cn(A)] \subseteq [\phi]$ then $[A] \subseteq [\phi]$. We can observe from the definition of Cn that we have: $[A] \subseteq [\phi]$ iff $\phi \in Cn(A)$. Applying this to the preceding gives that: for every ϕ if $\phi \in Cn(Cn(A))$ then $\phi \in Cn(A)$, or $Cn(Cn(A)) \subseteq Cn(A)$, which was to be shown. □

4.2. The AGM Approach in the Generalized Framework

In the classical AGM framework the epistemic input for revision was assumed to be a single sentence ϕ. Subsequently, the AGM framework was generalised to allow for (possibly infinite) *sets* of sentences as epistemic input [Peppas 1996; 2004; Zhang and Foo 2001]. Since herein we aim for generality, we shall follow the later approach.

Postulates. A revision function $*$ maps a theory K (also called a *belief set*) and a (possibly infinite) set of sentences A to a revised belief set $K * A$. For ease of notation, if $A = \{\phi\}$ for a sentence $\phi \in \mathcal{L}$, we shall often use $K * \phi$ as an abbreviation of $K * \{\phi\}$.

Assume that K is a theory, and A, B are sets of sentences, that is, $A, B \subseteq \mathcal{L}$. The AGM postulates for revision can be reformulated as follows:

(K*1) $K * A = Cn(K * A)$
(K*2) $A \subseteq K * A$
(K*3) $K * A \subseteq K + A$
(K*4) If $K \cup A$ is consistent then $K + A \subseteq K * A$
(K*5) If A is consistent then $K * A$ is consistent
(K*6) If $A \equiv B$ then $K * A = K * B$
(K*7) $K * (A \cup B) \subseteq (K * A) + B$
(K*8) If $(K * A) \cup B$ is consistent, then $(K * A) + B \subseteq K * (A \cup B)$

The postulates (K*1) – (K*8) are the well-known AGM postulates for revision, rephrased to allow sets of sentences for the epistemic input. However at the high level of abstraction at which our framework is developed, a ninth postulate is (sometimes) necessary:

(Acyc) If $A_1, \ldots, A_n$ are sets of sentences such that A_n is consistent with $K * A_1$, and for all $1 \leq i < n$, A_i is consistent with $K * A_{i+1}$, then A_1 is consistent with $K * A_n$

If our abstract framework is instantiated to classical propositional logic, then (Acyc) follows from (K*1) – (K*8) (see [Delgrande and Peppas 2015, Proposition 3]). In general however this is not true.

Preorders on Possible Worlds. For defining preorders on possible worlds, we basically adopt the definitions given earlier. The only notable difference is that, since now there could be sets of worlds that are not expressible by any set of sentences, we need to be a bit more careful with the set of minimal worlds that satisfy the input.

We shall say that a set S of worlds is *elementary* iff there exists a set of sentences $A \subseteq \mathcal{L}$ such that $[A] = S$. The following proposition is immediate, but useful.

PROPOSITION 4.4.

(*1*) *A set S of worlds is elementary iff $S = [t(S)]$*
(*2*) *If w is a possible world then $\{w\}$ is elementary.*

PROOF.

(1) For the right-to-left direction, the proof of the proposition is trivial: if $S = [t(S)]$ then, by definition, S is elementary.
For the opposite direction, assume that S is elementary. Then there exists a set of formulas A such that $S = [A]$. Hence

$$\begin{aligned} Cn(A) &= \{\phi \mid [A] \subseteq [\phi]\} \\ &= \{\phi \mid S \subseteq [\phi]\} \\ &= \{\phi \mid S \models \phi\} \\ &= t(S) \end{aligned}$$

Thus $Cn(A) = t(S)$. Moreover from Proposition 4.2, $[A] = [Cn(A)]$. Therefore, $S = [A] = [Cn(A)] = [t(S)]$ as desired.

(2) Lemma 4.1 states, for possible world w, that $\{w\} = [t(\{w\})]$. Hence $\{w\}$ is elementary from the previous part.

□

Depending on the specifics of $\mathcal{L}$, $\mathcal{M}$, and f, there may, or may not, exist non-elementary sets of worlds. For example, if our framework is instantiated to classical propositional logic with finitely many propositional variables, all sets of worlds are elementary. However, if the framework is instantiated to Horn logic, then non-elementary sets of worlds exist even when there are only finitely many variables.

A preorder $\preceq$ on possible worlds is called *faithful* to a belief set K iff it satisfies the following conditions:

(F1) $\preceq$ is total
(F2) if $[K] \neq \emptyset$, then $\min(\mathcal{M}, \preceq) = [K]$.

In addition, a preorder on possible worlds is called *regular* iff it satisfies:

(F3) for any $A \subseteq \mathcal{L}$, $\min([A], \preceq)$ is elementary.

The first two conditions (F1) – (F2) are the same as those of the classical AGM framework. The third condition was identified in [Peppas 2004], where it is called (SD), as being necessary for possibly infinite epistemic input. Subsequently it was noted to also be required in *finite* Horn theories by [Delgrande and Peppas 2015]. Of course in the context of propositional logic with finitely many variables, (F3) is vacuous since all sets of worlds are elementary.

The function $*$ induced from a preorder $\preceq$ faithful to a theory K is defined as follows:

($\preceq *$) $\quad K * A = t(\min([A], \preceq))$.

The following example illustrates various aspects of regular faithful rankings. Assume that we are working in Horn logic where $\mathcal{P} = \{p, q, r\}$. Then the following is a regular faithful ranking with respect to $K = Cn(\{p, q, r\})$:[11]

$$pqr \prec \bar{p}\bar{q}r \prec \begin{array}{c} \bar{p}qr \\ p\bar{q}r \end{array} \prec \begin{array}{c} pq\bar{r} \\ \bar{p}\bar{q}\bar{r} \\ \bar{p}q\bar{r} \\ p\bar{q}\bar{r} \end{array} \tag{4}$$

As a subtlety, note that even though the set of worlds $\{\bar{p}qr, p\bar{q}r\}$ is not elementary, the preorder is regular. In particular, there is no set of Horn formulas A such that $\min([A], \preceq) = \{\bar{p}qr, p\bar{q}r\}$, and so (F3) is satisfied in this case. Defining the function $*$ via condition ($\preceq *$), we have that

$$K * (\neg p \vee \neg q) \;=\; K * \neg p \;=\; Cn(\neg p \wedge \neg q \wedge r)$$

and

$$K * \neg r \;=\; Cn(\neg r).$$

In some cases, distinct regular faithful preorders may induce the same function $*$. For example it can be verified that the preorder

$$pqr \prec \bar{p}\bar{q}r \prec \bar{p}qr \prec p\bar{q}r \prec \begin{array}{c} pq\bar{r} \\ \bar{p}\bar{q}\bar{r} \\ \bar{p}q\bar{r} \\ p\bar{q}\bar{r} \end{array} \tag{5}$$

induces the same function as (4) at K. This would not be the case if the underlying logic were classical propositional logic, where for example via (4) we would have

$$K * (p \equiv \neg q) \;=\; Cn((p \equiv \neg q) \wedge r)$$

whereas via (5) we would have

$$K * (p \equiv \neg q) \;=\; Cn(\neg p \wedge q \wedge r).$$

So at this point we have two definitions of a function $*$, one in terms of postulates and the other in terms of preorders over possible worlds. In the next subsection we show that these two notions coincide.

4.3. Representation Results

In the standard AGM approach, the preorder $\preceq$ would be faithful but not necessarily regular, and the aim would be to prove that the functions induced from ($\preceq *$) coincided with those satisfying (K*1) – (K*8). In our general framework however this does not hold, and to a large extent this is due to the existence of non-elementary sets of worlds.

We illustrate the anomaly through a counter-example. Suppose that w_0, w_1, w_2, w_3 are distinct possible worlds, and $A_1, A_2, A_3 \subseteq \mathcal{L}$ are sets of sentences such that,

(i) $w_0 \in [t(\{w_1, w_2, w_3\})]$.
(ii) $w_1, w_2 \in [A_1]$ and $w_3 \notin [A_1]$.
(iii) $w_2, w_3 \in [A_2]$ and $w_1 \notin [A_2]$.

[11] For convenience, we occasionally will deviate from our standard set representation of interpretations (where the set contains the atoms assigned to true), and write interpretations as strings of literals, where $\bar{x}$ is just an abbreviation for $\neg x$.

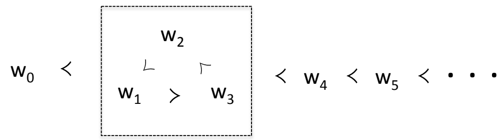

Fig. 1. An example for a pseudo-preorder.

(iv) $w_1, w_3 \in [A_3]$ and $w_2 \notin [A_3]$.

An example of worlds and sets of sentences satisfying conditions (i) – (iv) can be easily constructed, for example, in Horn logic. In particular, assume that $\mathcal{L}$ is built over three propositional variables p, q, r. As usual in Horn logic, we identify possible worlds with the set of variables they satisfy. With this convention, define $w_0 = \emptyset$, $w_1 = \{p, q\}$, $w_2 = \{p, r\}$, and $w_3 = \{q, r\}$. Moreover define $A_1 = \{p\}$, $A_2 = \{r\}$, and $A_3 = \{q\}$. It is not hard to see that all four conditions (i) – (iv) are indeed satisfied.[12]

Now consider the pseudo-preorder over worlds depicted in Figure 1. The minimal world is w_0 followed by a *cycle* of the three worlds $w_1 \prec w_2 \prec w_3 \prec w_1$, followed by a linear order over the remaining worlds.

Clearly, $\prec$ is not transitive and therefore not a preorder. Moreover, as shown next, there is no total preorder $\preceq'$ that is "revision-equivalent" to $\preceq$:

PROPOSITION 4.5. *Let* $w_0, w_1, w_2, w_3 \in \mathcal{M}$ *and* $A_1, A_2, A_3 \subseteq \mathcal{L}$ *be possible worlds and sets of sentences respectively, satisfying conditions (i) – (iv). Moreover let* $\prec$ *be the binary relation defined in Figure 1, and* $\preceq$ *its reflexive closure. Then there is no total preorder* $\preceq'$ *such that* $t(\min([A], \preceq')) = t(\min([A], \preceq))$*, for all* $A \subseteq \mathcal{L}$.

PROOF. Assume towards a contradiction that such a preorder $\preceq'$ does exist. Clearly by condition (ii), $\min([A_1], \preceq) = \{w_1\}$, and consequently, $\min([A_1], \preceq') = \{w_1\}$. This entails that $w_1 \prec' w_2$. In a similar manner, from condition (iii) we derive that $w_2 \prec' w_3$, and from condition (iv) we conclude that $w_3 \prec' w_1$. From the transitivity of $\preceq'$ we then derive that $w_1 \prec' w_1$. Contradiction. $\square$

Despite the above result, it turns out that the function $*$ induced from $\prec$ satisfies all eight postulates (K*1) – (K*8).

PROPOSITION 4.6. *The function* $*$ *induced via* ($\preceq *$) *from the binary relation* $\prec$ *of Figure 1 satisfies (K*1) – (K*8).*

PROOF. Postulates (K*1), (K*2), (K*3), (K*4), and (K*6) follow trivially from ($\preceq *$). For (K*5), let A be any consistent set of sentences. We need to show that $\min([A], \preceq) \neq \emptyset$. If $w_0 \in [A]$ then this is trivially true. Assume therefore that $w_0 \notin [A]$. Next we show that at least one of the worlds w_1, w_2, w_3 is not in $[A]$. Assume on the contrary that $w_1, w_2, w_3 \in [A]$. Then $A \subseteq t(\{w_1, w_2, w_3\})$. Hence, since by construction $w_0 \in [t(\{w_1, w_2, w_3\})]$, it follows that $w_0 \in [A]$, which however contradicts our earlier assumption. Hence we have shown that at least one of w_1, w_2, w_3 is not in $[A]$. From the definition of $\prec$ it then follows that $[A]$ has a minimal element wrt $\preceq$ and therefore $K * A$ is consistent.

[12]Conditions (ii) – (iv) are straightforward to verify. For condition (i), one only needs to recall that for any two worlds w, w' and Horn sentence ϕ, if $w \models \phi$ and $w' \models \phi$, then $w \cap w' \models \phi$.

For (K*7) and (K*8), consider any two sets of sentences A, B of $\mathcal{L}$. Observe that according to Figure 1, $K = t(\{w_0\})$. If B is inconsistent with $K * A$ then (K*7) and (K*8) are trivially true.

Assume therefore that B is consistent with $K * A$; i.e. $[t(\min([A], \preceq)] \cap [B] \neq \emptyset$. Then clearly $[A] \neq \emptyset$. Moreover, as already argued earlier, if $w_1, w_2, w_3 \in [A]$, then $w_0 \in [A]$ and consequently $\min([A], \preceq)$ is singleton (namely $\{w_0\}$). This is also the case, as can be easily verified from Figure 1, if at least one of w_1, w_2, w_3 is missing from $[A]$. Hence in all cases, $\min([A], \preceq)$ is a singleton. From (Expr) we then derive that $[t(\min([A], \preceq)]$ is also a singleton. Consequently from $[t(\min([A], \preceq)] \cap [B] \neq \emptyset$ it follows that the unique minimal A-world also satisfies B. Therefore $\min([A], \preceq) = \min([A \cup B], \preceq)$ and consequently $(K * A) + B = K * A = K * (A \cup B)$. Thus (K*7) and (K*8) are true. □

PROPOSITION 4.7. *The function $*$ induced via ($\preceq *$) from the binary relation $\prec$ of Figure 1 violates (Acyc).*

PROOF. From Conditions (ii) – (iv) and the definition of $\prec$, we have that $[K * A_1] = \{w_1\}$, $[K * A_2] = \{w_2\}$, and $[K * A_3] = \{w_3\}$. Hence A_3 is consistent with $K * A_1$; A_1 is consistent with $K * A_2$; and A_2 is consistent with $K * A_3$. From (Acyc) then we derive that A_1 is consistent with $K * A_3$. Contradiction. □

It is informative to consider an instance of this example in Horn logic: Choose $A_1, A_2, A_3 \subseteq \mathcal{L}$, such that [13]

$$\begin{aligned}
[A_1] &= \{w_1, w_2, w_1 \cap w_2\}, \\
[A_2] &= \{w_2, w_3, w_2 \cap w_3\}, \\
[A_3] &= \{w_1, w_3, w_1 \cap w_3\}.
\end{aligned}$$

Note that we can assume that $w_1 \cap w_2$, $w_2 \cap w_3$, $w_1 \cap w_3$ are all different from w_i ($i \in \{0, 1, 2, 3\}$) and thus are of form w_j for $j > 3$. Moreover, by the definition of $\prec$ it follows that $[K * A_1] = \{w_1\}$, $[K * A_2] = \{w_2\}$, and $[K * A_3] = \{w_3\}$. Hence A_3 is consistent with $K * A_1$, A_1 is consistent with $K * A_2$, and A_2 is consistent with $K * A_3$. From (Acyc) then we derive that A_1 is consistent with $K * A_3$. Contradiction.

To this point we have shown that directly applying the AGM approach to arbitrary logics is problematic. On the one hand, the standard AGM postulates are not strong enough to rule out cycles in an intended corresponding preorder on worlds. On the other hand, a revision function defined in terms of an arbitrary faithful ranking over worlds may violate the AGM postulates. It proves to be the case that by adding the postulate (Acyc) and by restricting faithful rankings to those that are regular, we can obtain a representation result. We first show that any faithful regular preorder satisfies the AGM postulates and (Acyc).

THEOREM 4.8. *Let K be a belief set and $\preceq$ a preorder over $\mathcal{M}$ that is faithful to K and regular. Then the function $*$ induced from ($\preceq *$) satisfies postulates (K*1) – (K*8) and (Acyc).*

PROOF. Postulates (K*1) – (K*4) follow immediately from ($\preceq *$) and the fact that $\preceq$ is faithful to K. For (K*5), let A be any consistent set of sentences in $\mathcal{M}$. Then $[A] \neq \emptyset$ and therefore $\min([A], \preceq) \neq \emptyset$, which again entails that $K * A$ is consistent.

For (K*6), assume that $A, B \subseteq \mathcal{L}$ are such that $A \equiv B$. Then $[A] = [B]$ and consequently, $\min([A], \preceq) = \min([B], \preceq)$. This again entails $K * A = K * B$ as desired.

[13] Recall for instance our earlier example where $\mathcal{L}$ is built over propositional variables p, q, r and we define $w_0 = \emptyset$, $w_1 = \{p, q\}$, $w_2 = \{p, r\}$, and $w_3 = \{q, r\}$, as well as $A_1 = \{p\}$, $A_2 = \{r\}$, and $A_3 = \{q\}$.

For (K*8), consider any two sets $A, B \subseteq \mathcal{L}$ such that B is consistent with $K * A$. Then clearly both A and B are consistent, and moreover we have $[B] \cap [t(\min([A], \preceq))] \neq \emptyset$ by assumption. Since, by (F3), $\min([A], \preceq)$ is elementary, we derive from Proposition 4.4 that $[B] \cap \min([A], \preceq) \neq \emptyset$. This again entails that $\min([A \cup B], \preceq) = [B] \cap \min([A], \preceq)$. Hence $K * (A \cup B) = (K * A) + B$. Thus (K*8) is satisfied.

The argument above also proves that (K*7) holds if B is consistent with $K * A$. If on the other hand B is inconsistent with $K * A$, then $(K * A) + B = \mathcal{L}$, and therefore, clearly, (K*7) is once again satisfied.

Finally for (Acyc), let $A_1, \ldots A_n \subseteq \mathcal{L}$ be sets of sentences such that A_n is consistent with $K * A_1$, and for all $1 \leq i < n$, A_i is consistent with $K * A_{i+1}$.

Since A_1 is consistent with $K * A_2$ it follows that $[A_1] \cap [t(\min([A_2], \preceq))] \neq \emptyset$. Then by (F3) and Proposition 4.4 we derive that $[A_1] \cap \min([A_2], \preceq) \neq \emptyset$. Hence there is an A_1-world, call it w_1', such that $w_1' \preceq r$, for all $r \in [A_2]$. Similarly, from A_2 being consistent with $K * A_3$ we conclude that there is a $w_2' \in [A_2]$ such that $w_2' \preceq r$, for all $r \in [A_3]$. Applying the same argument $(n-1)$-times, we derive that there exist worlds $w_1', \ldots, w_{n-1}'$ such that for all $1 \leq i < n$, $w_i' \preceq r$ for all $r \in [A_{i+1}]$. From the transitivity of $\preceq$ we then derive that $w_1' \preceq r$, for all $r \in [A_n]$. Finally, from A_n being consistent with $K * A_1$ it follows that there is a minimal A_1-world, call it w_1'', that satisfies A_n. Moreover, from $w_1'' \preceq w_1' \preceq r$ (for all $r \in [A_n]$), it follows that w_1'' is also a minimal A_n-world; that is, $w_1'' \in \min([A_n], \preceq)$. Since $\min([A_n], \preceq)$ contains an A_1-world, it follows that A_1 is consistent with $K * A_n$ as desired. □

The next theorem gives the converse result, that, for any revision function satisfying the AGM postulates and (Acyc), there is a corresponding regular faithful ranking on possible worlds. The full proof can be found in the appendix.

THEOREM 4.9. *Let K be a belief set and $*$ a revision function satisfying (K*1) – (K*8) and (Acyc). Then there exists a total preorder $\preceq$ over $\mathcal{M}$ that is faithful to K and regular, such that ($\preceq *$) is satisfied.*

PROOF IDEA. We progressively construct the preorder $\preceq$ alluded to in the statement of the theorem. First we define, using K and $*$, a binary relation $\sqsubseteq$ over $\mathcal{M}$ for which we show that $[K * A] = \min([A], \sqsubseteq)$ for all $A \subseteq \mathcal{L}$. In general, $\sqsubseteq$ is neither transitive nor total (although it is reflexive). The transitive closure of $\sqsubseteq$, denoted $\preceq_0$, is clearly a preorder, but in general it is not total. We therefore construct a series of extensions of $\preceq_0$, denoted $\preceq_1, \preceq_2, \cdots$, that preserve the minimal elements of $[A]$ for all $A \subseteq \mathcal{L}$. The union of this series is denoted $\preceq$ and it is shown to be a total preorder having all the desired properties. □

5. ITERATED REVISION IN THE GENERAL FRAMEWORK

The previous section has shown that the classical AGM approach can be rephrased in a highly general framework. In this section we examine the Darwiche and Pearl [1997] approach to iterated revision with respect to the general approach.

Firstly we note that the Darwiche and Pearl approach (or the *DP approach* for short) differs from the classical AGM approach in an important aspect: revision functions apply to *epistemic states* rather than to *belief sets*.

As already noted, a belief set K (alias theory[14]) represents the beliefs of an agent at a certain point in time. An epistemic state, on the other hand, is a richer structure. It contains, in addition to K, "...the entire information needed for coherent reasoning,

[14]To be precise, [Darwiche and Pearl 1997] follow the conventions of [Katsuno and Mendelzon 1991] and model a belief set as a *sentence* rather than a *theory*. Moreover new input is also modelled as a sentence, rather than a set of sentences. However, in the present context where we assume only finitely many possible worlds, these differences are immaterial.

including, in particular, the very strategy of belief revision which the agent wishes to employ at that given time" [Darwiche and Pearl 1997].

The belief set assigned to an epistemic state S is denoted by $bel(S)$. Crucially, two different epistemic states S, S', can be assigned the same belief set, i.e. $bel(S) = bel(S')$. It is fairly straightforward to switch to epistemic states in our general framework.[15] The main step is to introduce two extra primitives:

- a nonempty set Σ, the elements of which we call *epistemic states*.
- a function $bel : \Sigma \mapsto 2^{\mathcal{L}}$.

Moreover, for Σ and bel to have their intended meaning we need the following two restrictions:

(Σ1) For all $S \in \Sigma$, $bel(S)$ is a *theory* of $\mathcal{L}$.
(Σ2) For every theory K of $\mathcal{L}$, there is an epistemic state $S \in \Sigma$ such that $bel(\Sigma) = K$.

Conditions (Σ1) - (Σ2) are pretty much self-explanatory. (Σ1) restricts the set of sentences assigned to an epistemic state to be a *theory*. (Σ2) says that epistemic states are at least as rich as belief sets; the same assumption is made, implicitly, in [Darwiche and Pearl 1997].

Rephrasing Darwiche and Pearl in our extended general framework, a revision function $*$ is defined as a function mapping an epistemic state S and a set of sentences A to an epistemic state $S * A$. Of course the postulates (K*1) - (K*8), (Acyc) need to be adjusted accordingly. In particular each occurrence of "$K * X$" in the postulates is replaced with "$bel(S * X)$", and each remaining occurrence of "K" in the postulates is replaced with "$bel(S)$". Thus for example, (K*4) becomes:[16]

(K*4) If $bel(S) \cup A$ is consistent then $bel(S) + A \subseteq bel(S * A)$.

Having changed to epistemic states on the postulational side, Darwiche and Pearl make analogous adjustments to the semantic side. Following their lead, we now assign faithful preorders $\preceq$ to epistemic states (rather than to belief sets). Moreover we define $\preceq$ to be a faithful (regular) preorder with respect to an epistemic state S iff it is faithful (regular) with respect to $bel(S)$.

With the above changes, the representation results connecting the postulates with the semantics (Theorems 4.8 and 4.9), also hold for the extended general framework. The reason is that (K*1) - (K*8), (Acyc), regulate only *one-step* transitions; under this restriction the distinction between epistemic states and belief sets is immaterial.

Consider now the postulates introduced in [Darwiche and Pearl 1997], call them the *DP postulates*, to regulate *iterated revision*:[17]

(DP1) If $A \subseteq Cn(B)$, then $bel((S * A) * B) = bel(S * B)$.
(DP2) If $B \cup A$ is inconsistent, then $bel((S * A) * B) = bel(S * B)$.
(DP3) If $A \subseteq bel(S * B)$, then $A \subseteq bel((S * A) * B)$.
(DP4) If $A \cup bel(S * B)$ is consistent, then $A \cup bel((S * A) * B)$ is also consistent.

---

[15]We make this change to epistemic states only for this section where we discuss iterated revision. For the rest of the paper we move back to belief sets.

[16]We note that Darwiche and Pearl work with postulates (R1) - (R6) of [Katsuno and Mendelzon 1991], rather than the original AGM postulates (K*1) - (K*8). As already mentioned the difference is inessential. In fact, by using (K*1) - (K*8) we have a completely uniform mapping of the postulates to the DP approach, instead of dealing with one of the postulates separately, like Darwiche and Pearl need to do with (R4).

[17]As already mentioned, [Darwiche and Pearl 1997] follow the modelling conventions of [Katsuno and Mendelzon 1991] and hence the DP postulates therein are formulated differently. The differences however are inessential.

The theorem below shows that, in our extended general framework, the new postulates (DP1) - (DP4) are consistent with (K*1) - (K*8), (Acyc).

THEOREM 5.1. *Let S be any epistemic state in Σ. Then there exists a revision function $*$ such that at S it satisfies (K*1) - (K*8), (Acyc) as well as (DP1) - (DP4).*

PROOF. See Appendix A.2

Theorem 5.1 shows that our extended general framework is compatible with the DP approach. Yet Darwiche and Pearl did more than just to show consistency between their postulates and the AGM postulates. In particular they provided a semantic characterisation of (DP1) - (DP4) in terms of constraints over faithful preorders.

More precisely, consider an arbitrary revision function $*$ and an epistemic state S. Denote by $\preceq$ the faithful preorder that $*$ assigns to S by means of ($\preceq *$). Moreover, for every set $A \subseteq \mathcal{L}$, denote by $\preceq_A$ the faithful preorder that $*$ assigns to $S * A$. It was shown in [Darwiche and Pearl 1997] that $*$ satisfies (DP1) - (DP4) at S iff the following constraints between $\preceq$ and $\{\preceq_A\}_{A \subseteq \mathcal{L}}$ are satisfied:

(IR1) If $w \models A$ and $w' \models A$ then $w \prec w'$ iff $w \prec_A w'$.
(IR2) If $w \not\models A$ and $w' \not\models A$ then $w \prec w'$ iff $w \prec_A w'$.
(IR3) If $w \models A$ and $w' \not\models A$ then $w \prec w'$ entails $w \prec_A w'$.
(IR4) If $w \models A$ and $w' \not\models A$ then $w \preceq w'$ entails $w \preceq_A w'$.

Unfortunately this nice correspondence between (DP1) - (DP4) and (IR1) - (IR4) breaks down in our extended general framework. We show this with a counter-example. We shall construct a revision function that satisfies (DP1) - (DP4) but violates (IR2).

The counter-example is based on an instantiation of our framework to Horn logic built over only two propositional variable p, q. The possible worlds of this instantiation are $\mathcal{M} = \{pq, p\overline{q}, \overline{p}q, \overline{pq}\}$, where each world satisfies the literals that appear in it.[18]

Let S be an epistemic state such that $bel(S)$ is the closure under Horn logic, denoted $Cn_{\mathcal{H}}$, of the Horn sentence $\neg p \wedge \neg q$; in symbols, $bel(S) = Cn_{\mathcal{H}}(\{\neg p \wedge \neg q\})$. Clearly, for any revision function $*$ satisfying $(K * 1)$ - $(K * 8)$, (Acyc), we have $bel(S * \{pq\})$ = $Cn_{\mathcal{H}}(\{p \wedge q\})$. Now, consider the following two preorders over $\mathcal{M}$:

$$\overline{pq} \prec \begin{matrix} \overline{p}q \\ p\overline{q} \end{matrix} \prec pq$$

$$pq \prec_{pq} \overline{pq} \prec_{pq} \overline{p}q \prec_{pq} p\overline{q}$$

It is not hard to verify that $\preceq$, $\preceq_{pq}$ are regular and faithful preorders with respect to S and $S * \{pq\}$ respectively. Hence there exists a revision function $*$ that satisfies $(K * 1)$ - $(K * 8)$, (Acyc), such that $*$ is associated with $\preceq$ at S, and with $\preceq_{pq}$ at $S * \{pq\}$, via means of ($\preceq *$) . Next we show that $*$ satisfies (DP1) - (DP4) at S for $A = \{p \wedge q\}$.

Consider any set of Horn sentences B. For (DP1), assume that $\{p \wedge q\} \subseteq Cn_{\mathcal{H}}(B)$, and therefore $[B] \subseteq [p \wedge q]$. If B is inconsistent, then (DP1) is trivially true. Assume therefore that $[B] \neq \emptyset$. Then, $[B] = \{pq\}$, and therefore, $min([B], \preceq) = min([B], \preceq_{pq})$, which again entails $bel((S * A) * B) = bel(S * B)$ as desired.

For (DP2), assume that $B \cup \{p \wedge q\}$ is inconsistent. Then $[B] \subseteq \{p\overline{q}, \overline{p}q, \overline{pq}\}$. If B is inconsistent, then (DP2) is trivially true. Moreover, if $[B]$ is a singleton, then clearly

[18]Recall that, for the sake of readability, for a literal x we shall often write $\overline{x}$ as an abbreviation of $\neg x$. Moreover, whenever a sequence of literals appears in the place of a sentence, we take it to be the conjunction of the literals in the sequence; for example "$S * \{p\overline{q}\}$" is an abbreviation for "$S * \{p \wedge \neg q\}$".

$min([B], \preceq) = min([B], \preceq_{pq}) = [B]$, which in turn entails $bel((S * A) * B) = bel(S * B)$. Assume therefore that $[B]$ contains at least two worlds. Next observe that there is no set of Horn sentences X such that $[X] = \{p\overline{q}, \overline{p}q\}$. Consequently the only possible values for $[B]$ are $[B] = \{p\overline{q}, \overline{p}q, \overline{pq}\}$, $[B] = \{p\overline{q}, \overline{pq}\}$, and $[B] = \{\overline{p}q, \overline{pq}\}$. In all three cases, it is easy to see that $min([B], \preceq) = min([B], \preceq_{pq}) = \{\overline{pq}\}$. Hence $bel((S * A) * B) = bel(S * B)$ as desired.

For (DP3), assume that $A \subseteq bel(S * B)$. If B is inconsistent then (DP3) follows immediately. Assume therefore that B is consistent. Then by $(K * 1)$ - $(K * 5)$, from $\{p \wedge q\} \subseteq bel(S * B)$ we derive that $[bel(S * A)] = [bel(S * B)] = \{pq\}$; moreover $B \subseteq bel(S * B)$. Hence, from $(K * 1)$ - $(K * 4)$ it follows that $bel((S * A) * B) = bel(S * A)$, which again entails $A \subseteq bel((S * A) * B)$.

Finally for (DP4), assume that $\{p \wedge q\} \cup bel(S * B)$ is consistent. Then, by (K*2), B is consistent, and $pq \in min([B], \preceq)$. Since pq is also the maximum world with respect to $\preceq$, we derive that $[B] = \{pq\}$. Hence $A \equiv B$, and consequently, $bel(S * A) = bel(S * B) = bel((S * A) * B)$. Therefore $\{p \wedge q\} \cup bel((S * A) * B)$ is consistent, as desired.

We have thus shown that $*$ satisfies the DP postulates at S for $A = \{p \wedge q\}$. On the other hand, $\preceq$ and $\preceq_{pq}$ violate (IR2). In particular, observe that $p\overline{q} \not\models p \wedge q$, $\overline{p}q \not\models p \wedge q$, $p\overline{q} \preceq \overline{p}q$, and yet $\overline{p}q \prec_{pq} p\overline{q}$.

The main reason the correspondence between (DP1) - (DP4) and (IR1) - (IR4) breaks down in our framework is because of the existence of non-elementary sets of worlds. In particular, on the one hand, with "non-elementariness" there could be more than one regular faithful preorder that is associated with the same revision function $*$ and epistemic state S by means of $(\preceq *)$. On the other hand, the extra constraint of regularity for the preorders $\preceq_A$ may not necessarily be compatible with (IR1) - (IR4).

In the rest of the paper we focus once again on *one-step transitions* and hence we switch back to belief sets and to the general framework as it was developed before this section.

6. INSTANCES OF THE APPROACH

In this section we consider various instantiations of the general approach with respect to specific logics. We begin with revision in classical propositional logic, noting that in this case the general approach reduces to the standard AGM approach. Subsequently we review revision in Horn theories, briefly considering as a special case revision in definite clause theories. Third, we discuss revision in extended logic programs. While the model theory looks quite different from that of classical logic, nonetheless it is straightforward to show that our results cover this class of approaches. Last, we examine revision in what is arguably the simplest approach that may be considered to be a non-trivial logic, in what we call *literal revision*.

6.1. Classical Propositional Logic

In propositional logic, our language $\mathcal{L}_P$ is built from a finite set of atoms $\mathcal{P} = \{p, q, \dots\}$ with sentences formed using the usual set of propositional connectives. The set of possible worlds $\mathcal{M}_P$ corresponds to the set of interpretations of $\mathcal{L}_P$, and the function f_P assigning sentences of $\mathcal{L}_P$ to sets of possible worlds is given by the standard satisfaction relation of propositional logic.

In this setting the restriction (Expr) is trivially satisfied. Moreover, in this setting, the postulate (Acyc) is derivable from the AGM postulates (K*1) – (K*8) [Delgrande and Peppas 2015, Proposition 3]. Every set of worlds $S \subseteq \mathcal{M}_P$ is elementary, in that for any $S \subseteq \mathcal{M}_P$ there is a sentence $\phi \in \mathcal{L}_P$ such that $[\phi] = S$. In particular, in the proof of Theorem 4.9 we obtain for any worlds w_1, w_2 that $[B(w_1, w_2)] = \{w_1, w_2\}$. Consequently, the relation $\preceq'$ defined in the proof of Theorem 4.9 corresponds to the definition of $\preceq$ in [Katsuno and Mendelzon 1991], where they show that $\preceq$ defines a total preorder. The overall result is that restricted to finite propositional logic, we just need the standard AGM postulates, all sets of worlds are elementary, and the soundness and completeness results of [Katsuno and Mendelzon 1991] go through. Hence our general approach reduces to the AGM approach (as formulated by Katsuno and Mendelzon) when the underlying logic contains classical propositional logic.

6.2. Horn Logic

We next consider revision in Horn clause theories. Basic definitions and issues were presented in Section 2.2; as well, [Delgrande and Peppas 2015] provides an extensive development of AGM-style revision in Horn theories. Consequently, in this subsection we just examine Horn revision from the perspective of the general approach. However, first we briefly consider a restriction of Horn clauses, to that of *definite clauses*.

A definite clause is a clause (viz. disjunction of literals) that contains exactly one unnegated literal. Hence a definite clause can be written as an implication $a_1 \wedge a_2 \wedge \cdots \wedge a_n \Rightarrow a$ where $n \geq 0$ and each a_i, $1 \leq i < n$, and a are atoms. Thus, without worrying about formalities too much, our language $\mathcal{L}_D$ is the set of conjunctions of definite clauses, based on a finite set of atoms $\mathcal{P}$. The set of possible worlds would again correspond to the set of interpretations on the language. Definite clauses are expressively impoverished, in that any set of definite clauses is satisfiable.[19] What this means for our general approach is that revision is still definable, but it becomes a trivial operation. Thus, for any definite clause belief set K, the notion of a faithful assignment is still meaningful, as is the induced function ($\preceq *$). However, given that any set of definite clauses is satisfiable, this means that for any set of definite clauses A, $[K] \cap [A] \neq \emptyset$ and so we obtain that $K * A = t(\min([A], \preceq)) = Cn(K \cup A)$. Which is a roundabout way of saying that, not unexpectedly, while we obtain AGM-style revision for definite clauses, in fact it reduces to expansion.

Turning to Horn clauses, where a Horn clause is a clause with *at most* one unnegated literal, things become quite a bit more complicated, in fact arguably more complicated than the case of classical propositional logic. As reviewed in Section 2.2, a Horn clause can be written as an implication $a_1 \wedge a_2 \wedge \cdots \wedge a_n \Rightarrow a$, as in the case of definite clauses, but where a may be the falsum $\perp$. In terms of the basic components of our approach, our language $\mathcal{L}_H$ is that of Horn formulas (that is, conjunctions of Horn formulas) over a finite set of atoms. The set of possible worlds again is the set of propositional interpretations. As with propositional logic, our restriction (Expr) is trivially satisfied. It proves to be the case that the postulate (Acyc) is required: with respect to Horn logic, (Acyc) is independent of the postulates (K*1) – (K*8). As well, not every set of worlds is elementary: if a set of worlds is closed under intersection of atoms true in an interpretation, it is elementary; otherwise it is not. So in Horn clause theories, a preorder over interpretations is regular, if for all sets of Horn formulas A, $\min([A], \preceq)$ is closed under intersection. Consequently, we obtain a representation result for Horn clause theories with respect to the general revision postulates on the one hand, and faithful regular preorders over possible worlds on the other.

6.3. Answer Set Programs

Answer set programming (ASP) [Gelfond and Lifschitz 1988; Gebser et al. 2012; Brewka et al. 2011] is a major area of research in knowledge representation and reasoning. On the one hand it has a conceptually simple, declarative, theoretical foundation, while on the other hand efficient implementations are available. We omit a full introduction to ASP here, but refer the reader to the above citations; as well, [Delgrande et al. 2013] is a full development of AGM-style revision in ASP from first principles. So

[19]For example, the interpretation that assigns *true* to every atom satisfies every definite clause.

here we just describe how revision in ASP can be directly expressed using our general approach.

As before, our language is based on a finite set of propositional atoms $\mathcal{P}$. The language, $\mathcal{L}_{LP}$, is that of *generalised logic programs*, where a generalised logic program over $\mathcal{P}$ is a set of rules of the form:

$$a_1; \ldots; a_m;\ \sim b_1; \ldots; \sim b_n\ \leftarrow\ c_1, \ldots, c_j,\ \sim d_1, \ldots, \sim d_k \tag{6}$$

where $a_p, b_q, c_r, d_s \in \mathcal{P}$ and $p, q, r, s \geq 0$. The operators ';' and ',' express disjunctive and conjunctive connectives respectively while the unary operator $\sim$ is default negation or negation-as-failure. Two important subclasses of logic programs are given as follows. A rule r as in (6) is called *disjunctive* if $n = 0$; and *normal* if $m \leq 1$ and $n = 0$. (For a normal rule in which $k = 0$, we are back with a Horn clause.) A program is a *disjunctive logic program* if it consists of disjunctive rules only, and a program is a *normal logic program* if it consists of normal rules only. Any logic program as above induces zero or more *answer sets*, informally classical models of the program that satisfy certain minimality conditions.

Our interests aren't with answer sets here, but rather with the underlying model theory of such programs. This is given by a standard, albeit perhaps intricate, model theory, based on so-called *SE models* [Turner 2003]. The set of SE models is defined to be, for a set of atoms $\mathcal{P}$, the set of all ordered pairs (X, Y) where $X \subseteq Y \subseteq \mathcal{P}$.

This defines the language and set of models. The last component that we need to specify is the mapping f from sentences in the language to possible worlds, in this case, SE models. For this we need some additional terminology. A rule as in (6) can be written

$$H(r)^+; \sim H(r)^-\ \leftarrow\ B(r)^+, \sim B(r)^-$$

where $\sim X = \{\sim a \mid a \in X\}$ and

$$\begin{array}{lcllcl} a_1, \ldots, a_m & = & H(r)^+, & \qquad b_1, \ldots, b_n & = & H(r)^-, \\ c_1, \ldots, c_j & = & B(r)^+, & \qquad d_1, \ldots, d_k & = & B(r)^-. \end{array}$$

The *reduct* of a generalised logic program Π with respect to a set of atoms Y, denoted Π^Y, is the set of rules:

$$\{H(r)^+ \leftarrow B(r)^+ \mid r \in \Pi,\ H(r)^- \subseteq Y,\ B(r)^- \cap Y = \emptyset\}.$$

Note that the reduct consists of negation-free rules only. Informally Y can be thought of as a guess of a model of Π, and the reduct is composed of the rules in Π where the default negations have been "compiled out". An SE model (X, Y) is an *SE model of* Π iff $Y \models \Pi$ and $X \models \Pi^Y$, where $\models$ is the satisfaction relation in classical propositional logic.

So this defines the three major components required in our general approach to revision: the language, set of possible worlds, and satisfaction relation. While it is quite a bit more complex than the previously-described instances of the approach (and indeed won't make a whole lot of intuitive sense to someone not passingly familiar with ASP), it nonetheless fits within our general specification of a "logic".

Continuing, it turns out that the notion of an *elementary* set of worlds is non-trivial in ASP, in that there are sets of SE models S for which there is no program Π where $[\Pi] = S$. For the classes of programs that we are interested in, we have the following constraints on sets of SE models:

A set of SE models S is elementary:[20]

[20]The following conditions are referred to as *well-defined*, *complete*, and *closed under here-intersection*, respectively, in [Eiter et al. 2005; Cabalar and Ferraris 2007].

— in the class of generalised logic programs, if $(X, Y) \in S$ implies $(Y, Y) \in S$;
— in the class of disjunctive logic programs, if S is elementary in the class of generalised programs and if $(X, Y) \in S$ and $(Z, Z) \in S$ where $Y \subseteq Z$ then $(X, Z) \in S$; and
— in the class of normal logic programs, if S is elementary in the class of disjunctive programs and if $(X, Z), (Y, Z) \in S$ then $(X \cap Y, Z) \in S$.

With this we are done: We can apply Theorems 4.8 and 4.9, obtaining a representation result for AGM-style revision in these three classes of answer set programs.

6.4. Literal Revision

Our last instance of the general approach is of independent interest, in that it illustrates that AGM revision is definable even in extremely weak (albeit non-trivial) logics. To motivate this instance, we can ask "what is the weakest system that might reasonably be called a logic?" and then examine the associated AGM-style revision function in that logic. Arguably, for a system to be considered a non-trivial logic, it needs some notion of inconsistency expressible in the language. This could be given by a designated atom, such as $\perp$ in Horn logic, or it could be given in terms of a notion of negation. In this latter case, a set of formulas A is inconsistent if some formula and its negation are derivable from A. To this end, assume that an agent's knowledge is comprised of *facts* only, where a fact is an atom or a negated atom, and consequently in which an agent's knowledge is given by a set of literals. We refer to the resulting approach to revision as *literal revision*.

We need to first specify the three components of the general framework. As before, our language will be based on a finite set of atoms $\mathcal{P}$. The sentences of our language $\mathcal{L}_L$ will be sets of literals definable from $\mathcal{P}$. Hence, for $\mathcal{P} = \{p, q, r\}$ sentences include $\{p, \neg q\}$ and $\{p, \neg p, q\}$ which, as before, we can abbreviate as $p\bar{q}$ and $p\bar{p}q$. The set of possible worlds $\mathcal{M}$ will be the set of propositional interpretations over $\mathcal{P}$. The function f is defined as one would expect: for sentence ϕ, $f(\phi)$ is just the set of interpretations at which ϕ is true.

A sentence ϕ is inconsistent just if ϕ contains complementary literals, and a set of sentences A is inconsistent just if the union of members of A contains complementary literals. If a set of sentences A is inconsistent then $Cn(A) = \mathcal{L}_L$; and if A is consistent then $Cn(A) = \mathbf{P}(\cup A)$ where $\mathbf{P}(X)$ is the power set of X. For two sets of sentences A and B, we can define a notion of logical consequence by: $A \models B$ iff A is inconsistent or $(\cup A) \supseteq (\cup B)$.

In general, an arbitrary set of worlds $S \subseteq \mathcal{M}$ will not be elementary. For example, for $\mathcal{P} = \{p, q, r\}$, there is no set of sentences whose models is precisely $\{p\bar{q}r, \bar{p}qr\}$. It is straightforward to show that a set of worlds $S \subseteq \mathcal{M}$ is elementary just if $[\bigcap_{w \in S} w] = S$. Thus for example, if $S = \{\bar{p}\bar{q}r, p\bar{q}r, \bar{p}qr, pqr\}$ then $[\bigcap_{w \in S} w] = [\{r\}] = \{\bar{p}\bar{q}r, p\bar{q}r, \bar{p}qr, pqr\} = S$. In contrast, $S' = \{\bar{p}\bar{q}r, p\bar{q}r, \bar{p}qr\}$ is not elementary; we have $[\bigcap_{w \in S'} w] = [\{r\}] = \{\bar{p}\bar{q}r, p\bar{q}r, \bar{p}qr, pqr\} \neq S'$. Given this, our representation results apply and so we obtain a class of AGM-style revision functions in this approach. It can be noted that while the formal system is trivial, the resulting set of revision functions is not; for example, for $\mathcal{P} = \{p, q, r\}$ and $K = Cn(\{p, q\})$, the following is a faithful regular preorder defining a revision function:

$$\begin{matrix} pqr \\ pq\bar{r} \end{matrix} \preceq \begin{matrix} \bar{p}qr \\ p\bar{q}r \\ \bar{p}\bar{q}r \end{matrix} \preceq \begin{matrix} p\bar{q}\bar{r} \\ \bar{p}q\bar{r} \end{matrix} \preceq \bar{p}\bar{q}\bar{r}.$$

Consequently, we obtain, for example, that $K * \{\neg p\} = \{\neg p, r\}$. As well, the example illustrates a subtlety about the approach mentioned earlier: neither the set of worlds $\{\bar{p}qr, p\bar{q}r, \bar{p}\bar{q}r\}$ nor $\{p\bar{q}\bar{r}, \bar{p}q\bar{r}\}$ as they appear in the total preorder are elementary.[21] However, we don't run into trouble in defining revision in this preorder, since the preorder is nonetheless regular; for example, there is no set of sentences A such that $\min([A], \preceq) = \{p\bar{q}\bar{r}, \bar{p}q\bar{r}\}$.

Literal revision, while very basic, is of interest in at least two respects. First, it highlights aspects of the general approach while, second, it may also be of independent interest. With regards to the first point, literal revision demonstrates that AGM-style belief revision obtains in a very weak framework. The revision postulates are satisfied in this approach, and the semantic approach of regular faithful rankings captures literal belief revision. In a certain sense also, these results show that the AGM approach per se can be decoupled from the underlying logic, in that the AGM approach can be obtained even assuming essentially no meaningful underlying logic.

As well, literal revision may be of independent interest, since there has been some interest in *proper* knowledge bases [Levesque 1998], where a proper knowledge base is equivalent to a set of literals. Arguably a proper knowledge base is the simplest kind of knowledge base that allows open world reasoning. So, to the extent that proper knowledge bases are interesting, it is an interesting question to ask how change can be managed in such knowledge bases. Literal revision then addresses revision with respect to proper knowledge bases and demonstrates that meaningful revision operators that adhere to the AGM approach are definable.

7. DISCUSSION

In this section we briefly examine some of the underlying assumptions, results, and implications of the approach.

At the outset we suggested that the approach extends the AGM framework to any system that might reasonably be called a logic, with the caveat that for us a logic was defined in terms of a set of models and a satisfaction relation between models and formulas. Consequently, the approach is inapplicable for a system for which a model theory has not yet been developed, or for one with a non-standard notion of satisfaction of formulas.

However, it is interesting to note that the central constructions of the AGM approach are expressed essentially independently of a model theory. The approaches to constructing contraction functions are expressed in terms of certain maximal subsets of a belief set (in the case of remainders) or in terms of an ordering over the formulas in a belief set (for epistemic entrenchment). In Grove's system-of-spheres construction of revision functions, "possible worlds" are in fact maximal consistent sets of formulas. To be sure, approaches have been developed with respect to models, best known being the Katsuno-Mendelzon formulation of revision. However, in the Katsuno-Mendelzon approach, a crucial assumption is made that the language, and the set of models, is finite; consequently every set of models is representable by a formula.

We follow the Katsuno-Mendelzon approach in assuming that the set of models, or possible worlds, is finite. Of course it proves to be the case that in arbitrary logics it no longer holds that every set of models is representable by a formula; this leads to the notions in our approach of elementary sets of worlds and regular preorders. So for us the assumption of finiteness helps accentuate the difference between classical AGM belief revision and our approach.

On the postulational side, the schema (Acyc) provides a counterpart to regular preorders. Just as one does not need to distinguish regular preorders in the AGM approach, so too is the (Acyc) schema not required there. It can be noted also that

[21] For instance, the set $\{p\bar{q}\bar{r}, \bar{p}q\bar{r}, \bar{p}\bar{q}\bar{r}, pq\bar{r}\}$ is the least set of worlds containing $p\bar{q}\bar{r}$ and $\bar{p}q\bar{r}$ that is elementary.

the acyclicity postulate (Acyc) differs from the other AGM postulates, in that it is schematically infinite, and it specifies a schema for every $n \geq 3$. This raises the question of whether (Acyc) might be replaced by a finite schema. However, Yaggie and Turán [2015] show that the class of Horn belief revision operators require (Acyc), and cannot be characterized by finitely many postulate schemas. Since Horn belief revision is an instance of our approach, this means that a schematically infinite postulate is unavoidable. This is perhaps not surprising, since the (Loop) schema in distance-based revision similarly cannot be finitely characterised [Schlechta 2004; Ben-Naim 2006].

8. CONCLUSION

In this paper we have investigated belief revision in arbitrary logics, and we have shown that AGM-style revision can be obtained even when extremely little is assumed of the underlying language and its semantics. This is done by adding a postulate schema (Acyc) to the usual set of AGM postulates, on the one hand, and adding a constraint of *regularity* to preorders over possible worlds on the other hand. Both of these additions are redundant in the original AGM approach. Subsequently, a representation result established a correspondence between operators satisfying the postulates, and operators defined via minimal worlds in regular faithful rankings. The approach is also shown to be compatible with the general Darwiche-Pearl postulates for iterated revision, and several instances of the framework are given to illustrate the approach.

Consequently, AGM revision is extended to arbitrary logics, or at least to those based on a notion of model and satisfaction of formulas. Conceptually then, revision in such logics can in theory be addressed no differently than in classical logic. Expressed slightly differently, the AGM approach provides *constraints* on a rational belief operator; what our results show is that (rational) belief revision is definable essentially within any logic. This may not help immediately in the development of a specific revision operator; however, the approach may provide a guide to the formulation of specific revision operators in fragments of classical logic (such as Horn logic and description logics), and non-classical logics (such as modal logics and extended logic programs). In part, this is due to the fact that our representation result is applicable to any "reasonable" logic; thus once one has specified a language, set of models, and satisfaction relation, and supplied an appropriate notion of regularity, the representation results (Theorems 4.8 and 4.9) apply. To be sure, an appropriate logic-specific characterisation of elementary sets of worlds may be non-obvious; however our formal results offer the possibility of a very significant short cut in developing a representation result for logics (such as, for example, in description logics or modal logics) for which revision functions have not been developed.

These results may also help to better understand the overall landscape of belief change, particularly the interrelation of different belief change operators. In the classical AGM approach, belief revision and contraction are essentially two sides of the same coin, in that revision and contraction are interdefinable via the Levi and Harper identities. However, when the underlying logic is weaker than classical propositional logic, these identities generally fail. Thus, when the underlying logic is weaker than classical propositional logic, revision and contraction become distinct, independent change operations. Of interest then is to determine what relations exist between revision and contraction in the context of arbitrary logics.

A. PROOF OF THEOREMS 4.9 AND 5.1

A.1. Proof of Theorem 4.9

Let $K \subseteq \mathcal{L}$ be an arbitrary theory. We shall progressively construct the preorder $\preceq$ alluded to in the statement of the theorem. First we define, using K and $*$, a binary relation $\sqsubseteq$ over $\mathcal{M}$ for which we show that $[K*A] = \min([A], \sqsubseteq)$ for all $A \subseteq \mathcal{L}$. In general, $\sqsubseteq$ is neither transitive nor total (although it is reflexive). The transitive closure of $\sqsubseteq$, denoted $\preceq_0$, is clearly a preorder, but in general it is not total. We therefore construct a series of extensions of $\preceq_0$, denoted $\preceq_1, \preceq_2, \cdots$, that preserve the minimal elements of $[A]$ for all $A \subseteq \mathcal{L}$. The union of this series is denoted $\preceq$ and it will be shown to be a total preorder having all the desired properties.

In progressing from $\sqsubseteq$ to $\preceq$ we shall prove a number of supplementary results that will help us establish the main line of the argument.

First some notation. For any two worlds $w_1, w_2 \in \mathcal{M}$, we define

$$B(w_1, w_2) = t(\{w_1\}) \cap t(\{w_2\}).$$

Clearly, $w_1, w_2 \in [B(w_1, w_2)]$. Moreover, according to the following result, $B(w_1, w_2)$ is the strongest set of sentences consistent with both w_1 and w_2:

LEMMA A.1. *Let $A \subseteq \mathcal{L}$ be any set of sentences and $w_1, w_2 \in \mathcal{M}$ any two worlds. If $w_1, w_2 \in [A]$, then $[B(w_1, w_2)] \subseteq [A]$.*

PROOF. Assume that $w_1, w_2 \in [A]$. Let w_3 be an arbitrary world in $[B(w_1, w_2)]$ and assume towards a contradiction that $w_3 \notin [A]$. Then for some $\phi \in A$, $w_3 \not\models \phi$. On the other hand, since $w_1, w_2 \in [A]$, it follows that $w_1 \models \phi$ and $w_2 \models \phi$; hence $\phi \in t(\{w_1\}) \cap t(\{w_2\})$. Since $[B(w_1, w_2)] = [t(\{w_1\}) \cap t(\{w_2\})]$, we derive that $w_3 \in [t(\{w_1\}) \cap t(\{w_2\})]$, and consequently $w_3 \models \phi$. This of course contradicts our earlier conclusion. □

We now define the binary relation $\sqsubseteq$ over $\mathcal{M}$ as follows:

$$w_1 \sqsubseteq w_2 \text{ iff } w_1 \in [K * B(w_1, w_2)].$$

As usual, $\sqsubset$ denotes the strict part of $\sqsubseteq$; that is, $w_1 \sqsubset w_2$ iff $w_1 \sqsubseteq w_2$ and $w_2 \not\sqsubseteq w_1$.

LEMMA A.2. *Let w_1, w_2 be any two worlds such that $w_1 \sqsubseteq w_2$ and let $A \subseteq \mathcal{L}$ be a set of sentences such that $w_1 \in [A]$ and $w_2 \in [K * A]$. Then we have that $w_1 \in [K * A]$.*

PROOF. Let A be any set of sentences such that $w_1 \in [A]$ and $w_2 \in [K*A]$. Then clearly $B(w_1, w_2)$ is consistent with $K * A$. Hence by (K*7) and (K*8) we derive that $K*(A \cup B(w_1, w_2)) = (K*A) + B(w_1, w_2)$. Moreover, from $w_2 \in [K*A]$ and (K*2), it follows that $w_2 \in [A]$. From $w_1, w_2 \in [A]$ and Lemma A.1, it follows that $[B(w_1, w_2))] \subseteq [A]$. Hence, $[A \cup B(w_1, w_2)] = [A] \cap [B(w_1, w_2)] = [B(w_1, w_2))]$. Therefore by (K*6), $K * (A \cup B(w_1, w_2)) = K * B(w_1, w_2)$, and thus $K*B(w_1, w_2) = (K*A) + B(w_1, w_2)$. This, together with $w_1 \sqsubseteq w_2$, entails $w_1 \in [K * A]$. □

LEMMA A.3. *For all $A \subseteq \mathcal{L}$,* $\min([A], \sqsubseteq) = [K * A]$.

PROOF.
$\underline{LHS \subseteq RHS}$
Let $A \in \mathcal{L}$ be any set of sentences and assume towards a contradiction that there is a $w_1 \in \min([A], \sqsubseteq)$ such that $w_1 \notin [K * A]$. From $w_1 \in \min([A], \sqsubseteq)$ it follows that A is consistent, and therefore, by (K*5), $[K * A] \neq \emptyset$. Let w_2 be any world in $[K * A]$. By Lemma A.2 we derive that $w_1 \not\sqsubseteq w_2$. This again entails that $w_2 \not\sqsubseteq w_1$ (for otherwise w_1 wouldn't be minimal in $[A]$). Hence by the definition of $\sqsubseteq$, $w_1, w_2 \notin [K * B(w_1, w_2)]$. Since $B(w_1, w_2)$ is consistent, from (K*5) it follows that there is a world $w_3 \in [K * B(w_1, w_2)]$. Clearly then, $B(w_1, w_3)$ is consistent with $K * B(w_1, w_2)$, and therefore by (K*7) and (K*8), $K * (B(w_1, w_2) \cup B(w_1, w_3)) = (K * B(w_1, w_2)) + B(w_1, w_3)$.

Next we show that $[B(w_1, w_3)] \subseteq [B(w_1, w_2)]$. Assume towards a contradiction that for some $r \in [B(w_1, w_3)]$, $r \notin [B(w_1, w_2)]$. Then for some $\phi \in B(w_1, w_2)$, $r \not\models \phi$. This again entails that $\phi \notin t(\{w_3\})$, and therefore $w_3 \not\models \phi$. Notice however that from **(K*2)**, $\phi \in K * B(w_1, w_2)$, which of course contradicts $w_3 \in [K * B(w_1, w_2)]$. Hence we have shown that $[B(w_1, w_3)] \subseteq [B(w_1, w_2)]$.
From $[B(w_1, w_3)] \subseteq [B(w_1, w_2)]$, it follows that $[B(w_1, w_2) \cup B(w_1, w_3)] = [B(w_1, w_3)]$. Together with **(K*6)** we then derive that $K * B(w_1, w_3) = (K * B(w_1, w_2)) + B(w_1, w_3)$. Hence it follows that $w_3 \in [K * B(w_1, w_3)]$ and consequently, $w_3 \sqsubseteq w_1$. On the other hand from $w_3 \in [K * B(w_1, w_2)]$ and $w_1 \notin [K * B(w_1, w_2)]$, we derive from Lemma A.2 that $w_1 \not\sqsubseteq w_3$; hence, $w_3 \sqsubset w_1$.
Finally notice that from $w_1, w_2 \in [A]$, it follows that $[B(w_1, w_2)] \subseteq [A]$. Then, since we have shown that $[B(w_1, w_3)] \subseteq [B(w_1, w_2)]$, we derive that $w_3 \in [A]$. This however contradicts our assumption that w_1 is minimal in $[A]$ with respect to $\sqsubseteq$.

$\underline{RHS \subseteq LHS}$
Let $A \subseteq \mathcal{L}$ be any set of sentences and let w_1 be any world in $[K*A]$. We show that w_1 is $\sqsubseteq$-minimal in $[A]$. Let w_2 be any world in $[A]$. Clearly, since $w_1 \in [K * A]$, $B(w_1, w_2)$ is consistent with $K * A$, and consequently, by **(K*7)** and **(K*8)**, $K * (A \cup B(w_1, w_2)) = (K * A) + B(w_1, w_2)$. Moreover, since $w_1, w_2 \in [A]$, it follows that $[B(w_1, w_2)] \subseteq [A]$, and therefore, $[A \cup B(w_1, w_2)] = [B(w_1, w_2)]$. Hence by **(K*6)**, $K * B(w_1, w_2) = K * (A \cup B(w_1, w_2)) = (K * A) + B(w_1, w_2)$. Consequently, from $w_1 \in [K * A]$ we derive that $w_1 \in [K * B(w_1, w_2)]$, and therefore, $w_1 \sqsubseteq w_2$. Since w_2 was chosen arbitrarily, it follows that $w_1 \in \min([A], \sqsubseteq)$. $\square$

LEMMA A.4. *If* $w_1 \sqsubseteq w_2 \sqsubseteq \ldots \sqsubseteq w_n \sqsubseteq w_1$ *then* $w_1 \sqsubseteq w_n$.

PROOF. If $n = 1$, the lemma is trivially true.
Let $w_1, w_2, \ldots, w_n$ be any sequence of worlds, with $n > 1$, such that $w_1 \sqsubseteq w_2 \sqsubseteq \ldots \sqsubseteq w_n \sqsubseteq w_1$.
Then $w_1 \in [K*B(w_1, w_2)]$, $w_2 \in [K*B(w_2, w_3)], \ldots, w_{n-1} \in [K*B(w_{n-1}, w_n)]$, and $w_n \in [K * B(w_1, w_n)]$. Hence,

$$\begin{array}{lcl}
K * B(w_2, w_3) & \text{is consistent with} & B(w_1, w_2) \\
 & \vdots & \\
K * B(w_{n-1}, w_n) & \text{is consistent with} & B(w_{n-2}, w_{n-1}) \\
K * B(w_1, w_n) & \text{is consistent with} & B(w_{n-1}, w_n) \\
 & \text{and} & \\
K * B(w_1, w_2) & \text{is consistent with} & B(w_1, w_n)
\end{array}$$

Then by (Acyc) we derive that $K * B(w_1, w_n)$ is consistent with $B(w_1, w_2)$. Consequently, by **(K*7)** and **(K*8)**, $K * (B(w_1, w_n) \cup B(w_1, w_2)) = K * B(w_1, w_n)) + B(w_1, w_2)$.
On the other hand, since $K * B(w_1, w_2)$ is consistent with $B(w_1, w_n)$, **(K*7)** and **(K*8)** entail that $K*(B(w_1, w_n) \cup B(w_1, w_2)) = (K*B(w_1, w_2)) + B(w_1, w_n)$. Hence, from $w_1 \sqsubseteq w_2$, it follows that $w_1 \in [K * (B(w_1, w_n) \cup B(w_1, w_2))]$. Consequently, since $K * (B(w_1, w_n) \cup B(w_1, w_2)) = K * B(w_1, w_n)) + B(w_1, w_2)$, we conclude that $w_1 \in [K * (B(w_1, w_n)]$, and therefore $w_1 \sqsubseteq w_n$. $\square$

LEMMA A.5. *For any* $A \subseteq \mathcal{L}$*, if* $w \in \min([A], \sqsubseteq)$ *and* $w' \in [A]$*, then* $w \sqsubseteq w'$.

PROOF. Assume on the contrary that for some $A \subseteq \mathcal{L}$, there are $w, w' \in \mathcal{M}$ such that $w \in \min([A], \sqsubseteq)$, $w' \in [A]$, and $w \not\sqsubseteq w'$. Clearly then $w' \not\sqsubseteq w$. Consequently, $w, w' \notin [K * B(w, w')]$.
Since $B(w, w')$ is consistent, from (K*5) it follows that $[K * B(w, w')] \neq \emptyset$. Let r be any world in $[K * B(w, w')]$. Clearly $r \neq w$ and $r \neq w'$. From (K*2) it follows that $r \in [B(w, w')]$ and therefore by Lemma A.1, $r \in [A]$.
Next observe that $[B(w, r)] \subseteq [B(w, w')]$. To see this consider any world $r' \in [B(w, r)]$ and let ϕ be any sentence in $B(w, w')$. Since $w, r \in [B(w, w')]$ we derive that $w \models \phi$ and $r \models \phi$. Hence $\phi \in B(w, r)$. Then from $r' \in [B(w, r)]$ we derive that $r' \models \phi$. This again entails that $r' \in [B(w, w')]$. Therefore $[B(w, r)] \subseteq [B(w, w')]$.
From $[B(w, r)] \subseteq [B(w, w')]$ and (K*6) we then derive that $K * (B(w, w') \cup B(w, r)) = K * B(w, r)$. On the other hand from $r \in [K * B(w, w')]$ and $(K * 7)$ - $(K * 8)$ we derive that $K * (B(w, w') \cup B(w, r)) = (K * B(w, w')) + B(w, r))$. Combining the above it follows that $[K * B(w, r)] = [K * B(w, w')] \cap [B(w, r)]$. Hence, given that $r \in [K * B(w, w')]$ and $w \notin [K * B(w, w')]$, we derive that $r \in [K * B(w, r)]$ and $w \notin [K * B(w, r)]$. That is, $r \sqsubset w$. Since, as we have shown earlier, $r \in [A]$, this contradicts our initial assumption that w is $\sqsubseteq$-minimal in $[A]$. □

Let us now define $\preceq_0$ to be the transitive closure of $\sqsubseteq$; that is, $w \preceq_0 w'$ iff there exist worlds $u_1, \ldots, u_n$, such that $w \sqsubseteq u_1 \sqsubseteq \cdots \sqsubseteq u_n \sqsubseteq w'$. By construction, $\preceq_0$ is reflexive and transitive; that is, $\preceq_0$ is a partial preorder. Moreover,

LEMMA A.6. *For any $A \subseteq \mathcal{L}$, $\min([A], \preceq_0) = [K * A]$.*

PROOF. Let A be any set of sentences in $\mathcal{L}$. Given Lemma A.3 it suffices to show that $\min([A], \preceq_0) = \min([A], \sqsubseteq)$.
From Lemma A.5 it follows immediately that $\min([A], \sqsubseteq) \subseteq \min([A], \preceq_0)$. For the converse, let w be any element of $\min([A], \preceq_0)$. Consider any $w' \in [A]$ such that $w' \sqsubseteq w$. Since $w \in \min([A], \preceq_0)$ it follows that $w \preceq_0 w'$. Hence there exist $u_1, \ldots, u_n \in \mathcal{M}$ such that $w \sqsubseteq u_1 \sqsubseteq \cdots \sqsubseteq u_n \sqsubseteq w'$. Consequently, $w \sqsubseteq u_1 \sqsubseteq \cdots \sqsubseteq u_n \sqsubseteq w' \sqsubseteq w$. Therefore by Lemma A.4, $w \sqsubseteq w'$. This shows that $w \in \min([A], \sqsubseteq)$. □

An immediate corollary of Lemmas A.3, A.5, A.6 is the following:

COROLLARY A.7. *For all $A \subseteq \mathcal{L}$, if $w \in \min([A], \preceq_0)$ and $w' \in [A]$, then $w \preceq_0 w'$.*

If $\preceq_0$ happens to be total, then in view of the above results it is easy to verify that it satisfies all the properties required by the theorem. Assume therefore that $\preceq_0$ is not total. Then there are pairs of worlds that are incomparable with respect to $\preceq_0$. Given that there are only finitely many worlds in $\mathcal{M}$, there are also only finitely many incomparable pairs of worlds with respect to $\preceq_0$. Let $S_1, \ldots S_m$ be an *enumeration* of these incomparable pairs of world. We shall denote the elements of S_i as w_1^i and w_2^i; that is, $S_i = \{w_1^i, w_2^i\}$.[22] Moreover, we pick arbitrarily a world $w \in \mathcal{M}$ and we define $w_1^0 = w_2^0 = w$.

Next we shall construct a series of preorders $\preceq_1, \cdots \preceq_m$, each an extension of its predecessor, that preserves the properties reported in Lemma A.6 and Corollary A.7. The union of this series, denoted $\preceq$, will be shown to have all the desired properties.

[22]It makes no difference which of the two worlds in S_i is assigned the smaller subscript; the choice is arbitrary.

First one more definition. We define g to be a functions that maps any preorder $\preceq_i$ into a natural number $g(\preceq_i)$ as follows:

$$g(\preceq_i) \;=\; \begin{cases} 0 & \text{if } \preceq_i \text{ is total} \\ \text{the smallest number } k \text{ such that} & \\ w_1^k, w_2^k \text{ are incomparable wrt } \preceq_i & \text{otherwise} \end{cases}$$

With the aid of the above definition, we recursively define the series of preorders $\preceq_1, \cdots, \preceq_m$ as follows:

$$\preceq_{i+1} = \text{the transitive closure of } \preceq_i \cup \{(w_1^{g(\preceq_i)}, w_2^{g(\preceq_i)})\}.$$

Clearly all $\preceq_i$ are preorders. Moreover,

LEMMA A.8. *For all $i \geq 0$ and any $A \subseteq \mathcal{L}$,*
(i) $\min([A], \preceq_i) = [K * A]$.
(ii) *if $w \in \min([A], \preceq_i)$ and $w' \in [A]$, then $w \preceq_i w'$.*

PROOF. We prove the lemma by induction on i. For $i = 0$, the lemma follows from Lemma A.6 and Corollary A.7. Assume that the lemma is true for all $0 \leq i \leq k$ (*induction hypothesis*). Next we show that it holds for $i = k + 1$ (*induction step*).
If $\preceq_i$ is total then by construction $\preceq_{i+1} = \preceq_i$. Hence, since by the induction hypothesis the conditions (i)–(ii) are satisfied for $\preceq_i$, they are also satisfied for $\preceq_{i+1}$. Assume therefore that $\preceq_i$ is not total.
Let $A \subseteq \mathcal{L}$ be an arbitrary set of sentences. To prove Condition (i) it suffices to show, due to the induction hypothesis, that $\min([A], \preceq_{i+1}) = \min([A], \preceq_i)$. If $[A] = \emptyset$, then this is clearly true. Assume therefore that $[A] \neq \emptyset$. Then by (K*5), $[K * A] \neq \emptyset$, and therefore by the induction hypothesis, $\min([A], \preceq_i) \neq \emptyset$.
First we show that $\min([A], \preceq_i) \subseteq \min([A], \preceq_{i+1})$. Let w be any world in $\min([A], \preceq_i)$. Then by Condition (ii) of the induction hypothesis it follows that $w \preceq_i r$ for all $r \in [A]$. Since $\preceq_{i+1}$ is an extension of $\preceq_i$ we derive that $w \preceq_{i+1} r$ for all $r \in [A]$. Hence $w \in \min([A], \preceq_{i+1})$, which again shows that $\min([A], \preceq_i) \subseteq \min([A], \preceq_{i+1})$.
For the converse we shall prove the contrapositive. Let r be any world such that $r \notin \min([A], \preceq_i)$. We will show that $r \notin \min([A], \preceq_{i+1})$. If $r \notin [A]$ this is trivially true. Assume therefore that $r \in [A]$. Let z be any world in $\min([A], \preceq_i)$. Clearly $r \notin \min([A], \preceq_i)$ entails $r \not\preceq_i z$. Next we show that $r \not\preceq_{i+1} z$. Assume on the contrary that $r \preceq_{i+1} z$. Then, since $r \not\preceq_i z$, if follows by the construction of $\preceq_{i+1}$ and the transitivity of $\preceq_i$, that $r \preceq_i w_1^{g(\preceq_i)}$ and $w_2^{g(\preceq_i)} \preceq_i z$. Moreover by the induction hypothesis, Condition (ii), $z \preceq_i r$. Hence, $w_2^{g(\preceq_i)} \preceq_i z \preceq_i r \preceq_i w_1^{g(\preceq_i)}$, and consequently by transitivity, $w_2^{g(\preceq_i)} \preceq_i w_1^{g(\preceq_i)}$, which of course contradicts the definition of $w_1^{g(\preceq_i)}, w_2^{g(\preceq_i)}$ as the pair of worlds with the smallest index among those that are *incomparable* wrt $\preceq_i$. Thus we have shown that $r \not\preceq_{i+1} z$. On the other hand from $z \preceq_i r$ it follows that $z \preceq_{i+1} r$. Hence, since $z \in [A]$, we derive that $r \notin \min([A], \preceq_{i+1})$. Therefore $\min([A], \preceq_{i+1}) \subseteq \min([A], \preceq_i)$.
We have thus shown that $\preceq_{i+1}$ satisfies Condition (i). For Condition (ii), consider any $w \in \min([A], \preceq_{i+1})$ and let w' be any world in $[A]$. Since, as already shown, $\min([A], \preceq_{i+1}) = min([A], \preceq_i)$, we derive that $w \in \min([A], \preceq_i$

). Moreover, by Condition (ii) of the induction hypothesis, $w \preceq_i w'$. Hence, since $\preceq_{i+1}$ is an extension of $\preceq_i$, $w \preceq_{i+1} w'$. $\square$

We now define $\preceq$ to be the union of $\preceq_i$ for all $0 \leq i \leq m$:

$$\preceq \;=\; \bigcup_{i=0}^{m}(\preceq_i)$$

First we show that $\preceq$ is a preorder; that is, reflexive and transitive. Reflexivity is straightforward: since $\preceq_0$ is reflexive and $\preceq_0 \subseteq \preceq$, then $\preceq$ is also reflexive. For transitivity, let w_1, w_2, w_3 be any three worlds such that $w_1 \preceq w_2 \preceq w_3$. Then for some $i, j \geq 0$, $w_1 \preceq_i w_2$ and $w_2 \preceq_j w_3$. Let k be the greatest of the two numbers i, j. Then by the construction of the series $\preceq_0, \ldots, \preceq_m$, both preorders $\preceq_i$ and $\preceq_j$ are subsets of $\preceq_k$. Hence $w_1 \preceq_k w_2 \preceq_k w_3$, and therefore, $w_1 \preceq_k w_3$. Since $\preceq_k \subseteq \preceq$ we derive $w_1 \preceq w_3$.

Next we show that $\preceq$ is total. Assume on the contrary that there are two worlds r, r' that are incomparable wrt to $\preceq$. Since $\preceq_0 \subseteq \preceq$, it follows that r, r' are also incomparable wrt $\preceq_0$. Hence for some $i \geq 0$, $S_i = \{r, r'\}$. Observe that by the definition of g, we have $g(\preceq_{i+1}) > i$. Hence worlds in S_1, are comparable wrt $\preceq_{i+1}$; and so are the worlds in S_2, in S_3, ..., in S_i. That is $r \preceq_{i+1} r'$ or $r' \preceq_{i+1} r$. Since $\preceq$ extends $\preceq_{i+1}$ we derive that $r \preceq r'$ or $r' \preceq r$. Thus $\preceq$ is total, and hence it fulfills the first requirement, namely (F1), for being faithful to K.

To complete the proof we need to show that $\preceq$, also satisfies (F2) – (F3), as well as ($\preceq *$). We start with the latter. In fact we shall prove something slightly stronger than ($\preceq *$); namely that for all $A \subseteq \mathcal{L}$, $[K * A] = \min([A], \preceq)$.

Let A be any set of sentence in $\mathcal{L}$. If $[A] = \emptyset$ then from (K*2) we immediately derive $[K * A] = \min([A], \preceq) = \emptyset$. Assume therefore that $[A] \neq \emptyset$.

Consider any $w \in [K * A]$ and let r be any world in $[A]$. Then by Lemma A.8, $w \preceq_0 r$, and since $\preceq_0 \subseteq \preceq$, we derive that $w \preceq r$. This entails that $w \in \min([A], \preceq)$. Hence $[K * A] \subseteq \min([A], \preceq)$.

For the converse, let r be any world in $\min([A], \preceq)$. Since $[A] \neq \emptyset$, from (K*5) we get that $[K * A] \neq \emptyset$. Let w be any world in $[K * A]$. Clearly, by (K*2), $w \in [A]$, and since as already shown $\preceq$ is total, from $r \in \min([A], \preceq)$ we derive that $r \preceq w$. Hence, for some $i \geq 0$, $r \preceq_i w$. Moreover, from Lemma A.8 and $w \in [K * A]$, it follows that $w \in \min([A], \preceq_i)$. Therefore from $r \preceq_i w$ we derive that $r \in \min([A], \preceq_i)$. Using Lemma A.8 again we derive $r \in [K * A]$ as desired.

We have thus shown that for all $A \subseteq \mathcal{L}$, $\min([A], \preceq) = [K * A]$. This clearly proves ($\preceq *$). Moreover, combined with (K*5), it also proves (F3). Finally, by setting $A = \emptyset$, from $\min([A], \preceq) = [K * A]$ and (K*3) – (K*4), we derive (F2) as well. $\square$

A.2. Proof of Theorem 5.1

As already noted, Theorem 4.8 also holds for the extended framework developed in this section. Hence we can define the revision function $*$ alluded to in the present theorem in terms of total preorders over possible worlds.

More precisely, let us denote by $\preceq$ any regular faithful ranking with respect to S. For every set of sentences $A \subseteq \mathcal{L}$ we shall construct a total preorder $\preceq_A$ over $\mathcal{M}$, which will be shown to be faithful and regular with respect to $t(min([A], \preceq))$. The initial preorder $\preceq$ along with the family of preorders $\{\preceq_A\}_{A \subseteq \mathcal{L}}$ is then used to define a revision function $*$ which is subsequently shown to satisfy (DP1) - (DP4).

Consider an arbitrary set of sentences A. If A is inconsistent, then we define $\preceq_A$ to be identical to $\preceq$. If on the other hand A is consistent, then $\preceq_A$ is defined as follows:

$$w \preceq_A w' \text{ iff } w \in \min([A], \preceq), \text{ or } w \preceq w' \text{ and } w' \notin \min([A], \preceq).$$

According to the above definition, to construct $\preceq_A$, one starts with $\preceq$ and simply moves the minimal A-worlds (with respect to $\preceq$) to the beginning of the ranking; everything else is unchanged. We note that this construction is not new. It was proposed by Boutilier [Boutilier 1993; 1996] in his treatment of iterated revision, which he called *natural revision*.

Since $\preceq$ is a total preorder, it is not hard to verify that $\preceq_A$ is also a total preorder over $\mathcal{M}$, regardless of whether A is consistent or not. Moreover by construction it follows immediately that $\preceq_A$ is faithful with respect to $t(min([A], \preceq))$. Next we show that $\preceq_A$ is also regular.

If A is inconsistent then $\preceq_A = \preceq$ and hence **(F3)** follows immediately. Assume therefore that A is consistent. To prove **(F3)** we need to show that for any nonempty set of sentences C, $\min([C], \preceq_A)$ is elementary. From the construction of $\preceq_A$, there are only two cases to consider: either $\min([C], \preceq_A) = \min([C], \preceq)$ or $\min([C], \preceq_A) \subseteq \min([C], \preceq)$.[23] In the first case, since by assumption we have that $\min([C], \preceq)$ is elementary, so is $\min([C], \preceq_A)$. For the second case, from the definition of $\preceq_A$, we have $\min([C], \preceq_A) = \min([C], \preceq) \cap \min([A], \preceq)$. By assumption both $\min([C], \preceq)$ and $\min([A], \preceq)$ are elementary, so there are sets of sentences B_1 and B_2 such that $[B_1] = \min([C], \preceq)$ and $[B_2] = \min([A], \preceq)$. However, by Proposition 4.2 we have that $[B_1] \cap [B_2] = [B_1 \cup B_2]$. Hence $\min([C], \preceq) \cap \min([A], \preceq)$ is elementary, and so $\min([C], \preceq_A)$ is elementary.

Now define $*$ to be any revision function that assigns the preorder $\preceq$ to S, and $\preceq_A$ to $S * A$ for every set of sentences A. According to Theorem 4.8, $*$ satisfies the postulates **(K*1) - (K*8)**, **(Acyc)**. Hence what's left to be shown is that $*$ also satisfies **(DP1) - (DP4)**.

For **(DP1)**, assume that $B, A \subseteq \mathcal{L}$ are such that $A \subseteq Cn(B)$. If B is inconsistent then **(DP1)** trivially holds. Assume therefore that B is consistent. We distinguish between two cases. First suppose that $B \cup bel(S*A)$ is consistent. Then by **(K*7) - (K*8)** we derive that $bel(S*A) + B = bel(S*(B \cup A))$, and by **(K*3) - (K*4)** we get that $bel(S*A) + B = bel((S*A)*B)$. Moreover from **(K*6)** we derive that $bel(S*(B \cup A)) = bel(S*B)$. Combining the above it follows that $bel((S*A)*B) = bel(S*A) + B = bel(S*(B \cup A)) = bel(S*B)$ as desired. Now for the second case, assume that $B \cup bel(S*A)$ is inconsistent. Then no B-world belongs to $min([A], \preceq)$. Hence, by the construction of $\preceq_A$ it follows that the restriction of $\preceq$ to $[B]$ is identical with the restriction of $\preceq_A$ to $[B]$. Therefore $min([B], \preceq_A) = min([B], \preceq)$ and consequently, once again, $bel((S*A)*B) = bel(S*B)$ as desired.

For **(DP2)**, assume that $B, A \subseteq \mathcal{L}$ are such that $B \cup A$ is inconsistent. If A is inconsistent, then clearly $\preceq_A = \preceq$, and therefore $min([B], \preceq_A) = min([B], \preceq)$. This again entails $bel((S*A)*B) = bel(S*B)$ as desired. Hence assume that A is consistent. From $B \cup A$ being inconsistent it follows that no A-world satisfies B. Then by the construction of $\preceq_A$, it is not hard to see that the restriction of $\preceq_A$ to $[B]$ is identical to the restriction of $\preceq$ to $[B]$. This again entails that $min([B], \preceq_A) = min([B], \preceq)$ and therefore, once again, $bel((S*A)*B) = bel(S*B)$.

For **(DP3)**, assume that $B, A \subseteq \mathcal{L}$ are such that $A \subseteq bel(S*B)$. If B is inconsistent, then **(DP3)** follows immediately from **(K*1) - (K*2)**. Assume therefore that B is consistent. Then by **(K*5)**, so is $bel(S*B)$. Hence from $A \subseteq bel(S*B)$ we derive that A is also consistent. Moreover, from $A \subseteq bel(S*B)$ we derive that $min([B], \preceq) \subseteq [A]$. This again entails that any B-world that is $\preceq$-minimal in $[A]$, is also $\preceq$-minimal in $[B]$. Hence, by the construction of $\preceq_A$, it follows that the restriction of $\preceq_A$ to $[B]$ is identical to the restriction of $\preceq$ to $[B]$. This again entails that $min([B], \preceq_A) = min([B], \preceq)$ and therefore $bel((S*A)*B) = bel(S*B)$ as desired.

[23]These cases arise from $\min([C], \preceq) \cap \min([A], \preceq) = \emptyset$ and $\min([C], \preceq) \cap \min([A], \preceq) \neq \emptyset$ respectively.

Finally, for (DP4), assume that $B, A \subseteq \mathcal{L}$ are such that $A \cup bel(S * B)$ is consistent. Clearly then, by (K*2) both B and A are consistent. Moreover, from $A \cup bel(S * B)$ being consistent we derive that there exists an A-world in $min([B], \preceq)$. This entails that any B-world that is $\preceq$-minimal in $[A]$, is also $\preceq$-minimal in $[B]$. Hence, by the construction of $\preceq_A$, it follows that the restriction of $\preceq_A$ to $[B]$ is identical to the restriction of $\preceq$ to $[B]$. This again entails that $min([B], \preceq_A) = min([B], \preceq)$ and therefore $bel((S * A) * B) = bel(S * B)$. Hence, since $A \cup bel(S * B)$ is consistent, so is $A \cup bel((S * A) * B)$ $\square$

ACKNOWLEDGMENT

We thank the reviewers for their helpful and pertinent comments.

REFERENCES

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

A.R. Anderson and N.D. Belnap Jr. 1975. *Entailment: The Logic of Relevance and Necessity, Vol. I*. Princeton University Press, Princeton.

F. Baader, D. Calvanese, D. McGuiness, D. Nardi, and P Patel-Schneider (Eds.). 2007. *The Description Logic Handbook* (second ed.). Cambridge University Press, Cambridge.

R. Baumann and G. Brewka. 2015. AGM Meets Abstract Argumentation: Expansion and Revision for Dung Frameworks. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*. AAAI Press, 2734–2740.

Jonathan Ben-Naim. 2006. Lack of Finite Characterizations for the Distance-Based Revision. In *Proceedings of the Tenth International Conference on the Principles of Knowledge Representation and Reasoning*, P. Doherty, J. Mylopoulos, and C. Welty (Eds.). Lake District, UK, 239–248.

R. Booth, T. Meyer, and I.J. Varzinczak. 2009. Next Steps in Propositional Horn Contraction. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Pasadena, CA, 702–707.

R. Booth, T. Meyer, I. Varzinczak, and R. Wassermann. 2011. On the Link between Partial Meet, Kernel, and Infra Contraction and its Application to Horn Logic. *Journal of Artificial Intelligence Research* 42 (2011), 31–53.

C. Boutilier. 1993. Revision Sequences and Nested Conditionals. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Chambery, France, 519–531.

C. Boutilier. 1996. Iterated revision and minimal change of conditional beliefs. *Journal of Logic and Computation* 25 (1996), 262–305.

Gerhard Brewka, Thomas Eiter, and Miroslaw Truszczyński. 2011. Answer set programming at a glance. *Commun. ACM* 54, 12 (2011), 92–103. http://dblp.uni-trier.de/db/journals/cacm/cacm54.html#BrewkaET11

Gerhard Brewka, Jean-Guy Mailly, and Stefan Woltran. 2016. Translation-Based Revision and Merging for Minimal Horn Reasoning. In *Proceedings of the European Conference on Artificial Intelligence*. IOS Press, 734–742.

P. Cabalar and P. Ferraris. 2007. Propositional Theories are Strongly Equivalent to Logic Programs. *Theory and Practice of Logic Programming* 7, 6 (2007), 745–759.

S. Coste-Marquis, S. Konieczny, J.-G. Mailly, and P. Marquis. 2014. On the revision of argumentation systems: minimal change of arguments statuses. In *Proceedings of the Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*, Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter (Eds.). AAAI Press, 72–81.

N. Creignou, O. Papini, R. Pichler, and S. Woltran. 2014. Belief revision within fragments of propositional logic. *J. Comput. System Sci.* 80, 2 (2014), 427–449. DOI:http://dx.doi.org/10.1016/j.jcss.2013.08.002

A. Darwiche and J. Pearl. 1997. On the Logic of Iterated Belief Revision. *Artificial Intelligence* 89 (1997), 1–29.

J.P. Delgrande. 2008. Horn Clause Belief Change: Contraction Functions. In *Proceedings of the Eleventh International Conference on the Principles of Knowledge Representation and Reasoning*, G. Brewka and J. Lang (Eds.). AAAI Press, Sydney, Australia, 156–165.

J.P. Delgrande and P. Peppas. 2011. Revising Horn Theories. In *Proceedings of the International Joint Conference on Artificial Intelligence*. Barcelona, Spain, 839–844.

J.P. Delgrande and P. Peppas. 2015. Belief Revision in Horn Theories. *Artificial Intelligence* 218 (2015), 1 – 22. http://www.sciencedirect.com/science/article/pii/S0004370214001155

J.P. Delgrande, P. Peppas, and S. Woltran. 2013. AGM-Style Belief Revision of Logic Programs under Answer Set Semantics. In *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR13) (Lecture Notes in Artificial Intelligence)*, P. Cabalar and T.C. Son (Eds.), Vol. 8148. Springer Verlag, Corunna, Spain, 264–276.

J.P. Delgrande and R. Wassermann. 2013. Horn Clause Contraction Functions. *Journal of Artificial Intelligence Research* 48 (November 2013), 475–511.

M. Diller, A. Haret, T. Linsbichler, S. Rümmele, and S. Woltran. 2018. An Extension-Based Approach to Belief Revision in Abstract Argumentation. *International Journal of Approximate Reasoning* 93 (2018), 395–423.

P.M. Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77, 2 (1995), 321–358.

T. Eiter, H. Tompits, and S. Woltran. 2005. On Solution Correspondences in Answer Set Programming. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI 2005)*. 97–102.

D. Gabbay, O. Rodrigues, and A. Russo. 2008. Belief Revision in Non-Classical Logics. *The Review of Symbolic Logic* 1 (2008), 267–304. Issue 03. DOI:http://dx.doi.org/10.1017/S1755020308080246

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

M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. 2012. *Answer Set Solving in Practice*. Morgan & Claypool Publishers.

M. Gelfond and V. Lifschitz. 1988. The Stable Model Semantics for Logic Programming. In *Proceedings of the Fifth International Conference and Symposium of Logic Programming (ICLP'88)*, R. Kowalski and K. Bowen (Eds.). The MIT Press, 1070–1080.

A. Grove. 1988. Two Modellings for Theory Change. *Journal of Philosophical Logic* 17 (1988), 157–170.

S. O. Hansson. 1999. *A Textbook of Belief Dynamics*. Kluwer Academic Publishers.

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

H. Katsuno and A. Mendelzon. 1992. On the Difference between Updating a Knowledge Base and Revising it. In *Belief Revision*, P. Gärdenfors (Ed.). Cambridge University Press, Cambridge, 183–203.

S. Kraus, D. Lehmann, and M. Magidor. 1990. Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence* 44, 1-2 (1990), 167–207.

D. Lehmann, M. Magidor, and K. Schlechta. 2001. Distance Semantics for Belief Revision. *Journal of Symbolic Logic* 66, 1 (2001), 295–317.

H.J. Levesque. 1998. A Completeness Result for Reasoning with Incomplete First-Order Knowledge Bases. In *Proceedings of the Sixth International Conference on the Principles of Knowledge Representation and Reasoning*, A. Cohn, L. Schubert, and S. Shapiro (Eds.). Morgan Kaufmann Publishers, 14–23.

D. Lewis. 1973. *Counterfactuals*. Harvard University Press, Cambridge, Mass.

P. Peppas. 1996. Well Behaved and Multiple Belief Revision. In *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI-96)*. Wiley, 90–94.

P. Peppas. 2004. The Limit Assumption and Multiple Revision. *Journal of Logic and Computation* 14, 3 (2004), 355–371.

P. Peppas. 2008. Belief Revision. In *Handbook of Knowledge Representation*, F. van Harmelen, V. Lifschitz, and B. Porter (Eds.). Elsevier Science, San Diego, USA, 317–359.

G. Restall and J.K. Slaney. 1995. Realistic Belief Revision. In *Proceedings of the First World Congress in the Fundamentals of Artificial Intelligence*. Paris, 367–378.

M. Ribeiro and R. Wassermann. 2014. Minimal Change in AGM Revision for Non-Classical Logics. In *Proceedings of the Fourteenth International Conference on the Principles of Knowledge Representation and Reasoning*. AAAI Press, Vienna, Austria, 657–660.

Karl Schlechta. 2004. *Coherent Systems*. Elsevier, Amsterdam.

N. Schwind and K. Inoue. 2013. Characterization Theorems for Revision of Logic Programs. In *Proceedings of the 12th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2013) (Lecture Notes in Artificial Intelligence)*, P. Cabalar and T.C. Son (Eds.), Vol. 8148. Springer, 485–498.

B. Selman and H. Kautz. 1996. Knowledge compilation and theory approximation. *J. ACM* 43, 2 (1996), 193–224.

A. Tarski. 1936. On Some Fundamental Concepts of Metamathematics [1930]. In *Logic, Semantics, Methamathematics. Papers from 1923 to 1938*. Clarendon Press, Oxford, 30–36.

H. Turner. 2003. Strong equivalence made easy: Nested expressions and weight constraints. *Theory and Practice of Logic Programming* 3, 4 (2003), 609–622.

R. Wassermann. 2011. On AGM for Non-Classical Logics. *Journal of Philosophical Logic* 40, 2 (2011), 271–294.

Jon Yaggie and György Turán. 2015. Characterizability in Horn Belief Revision: An Application of Finite Model Theory. In *Proceedings of the European Workshop on Logics in Artificial Intelligence (JELIA 2015) (Lecture Notes in Artificial Intelligence)*, Loizos Michael and Antonis C. Kakas (Eds.). Springer Verlag, 497–511.

D. Zhang and N. Foo. 2001. Infinitary Belief Revision. *Journal of Philosophical Logic* 30, 6 (2001), 525–570.

Z. Zhuang and Maurice Pagnucco. 2010. Horn Contraction via Epistemic Entrenchment. In *Logics in Artificial Intelligence - 12th European Conference (JELIA 2010) (Lecture Notes in Artificial Intelligence)*, T. Janhunen and I. Niemelä (Eds.), Vol. 6341. Springer Verlag, 339–351.

Z. Zhuang and M. Pagnucco. 2011. Transitively Relational Partial Meet Horn Contraction. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*. Barcelona, Spain, 1132–1138.

Z. Zhuang and M. Pagnucco. 2012. Model Based Horn Contraction. In *Proceedings of the Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*. Rome, Italy, 169–178.

Z. Zhuang, M. Pagnucco, and Y. Zhang. 2013. Definability of Horn Revision from Horn Contraction. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*. Beijing, China, 1205–1211.

Received Month Year; revised Month Year; accepted Month Year