Proceedings of the Sixteenth International Conference on
Principles of Knowledge Representation and Reasoning (KR 2018)

Incorporating Relevance in Epistemic States in Belief Revision

James Delgrande,[1] **Pavlos Peppas**[2,3]
[1] School of Computing Science, Simon Fraser University, Burnaby, BC V5A1S6, Canada
[2] Centre for AI, School of Software, University of Technology Sydney, NSW 2007, Australia
[3] Dept of Business Administration, University of Patras, Patras 26500, Greece
jim@cs.sfu.ca, pavlos.peppas@uts.edu.au

Abstract

We present an account of *relevance* in belief revision where, intuitively, one wants to only consider the relevant part of an agent's epistemic state in a revision. We assume that *relevance* is a domain-specific notion, and that (ir)relevance assertions are given as part of the agent's epistemic state. Such assertions apply in a given context, and are of the form "in the case that formula σ holds, the Y part of the agent's epistemic state is independent of the rest of the epistemic state", where Y is part of the signature of the language. Two approaches are given, one in which (in semantic terms) conditions are placed on a faithful ranking on possible worlds to enforce the (ir)relevance assertions, and a second in which the possible worlds characterising the agent's beliefs may be modified in a revision. These approaches are shown to yield the same resulting belief set. Corresponding postulates and a representation result are given. The overall approach is compared to that of Parikh's for language splitting as well as with multivalued dependencies in relational databases.

1 Introduction

The area of belief revision addresses the problem in which an agent is to incorporate a new belief into its belief corpus. The dominant approach in belief revision, and in belief change as a whole, is the *AGM approach* (Alchourrón, Gärdenfors, and Makinson 1985), named after its original developers. A number of factors make revision a difficult and subtle problem: Since it is desirable to maintain consistency in a belief set if possible, and since a new belief may conflict with the agent's prior beliefs, some of these prior beliefs may need to be removed before the new belief can be consistently incorporated. The problem of *which* beliefs to discard is not a purely logical problem, and there will be domain constraints and epistemic factors to be taken into account in a practical revision system. To this end, a guiding principle is that change should be *minimal*, in that as few beliefs as possible are given up.

However, the notion of what constitutes minimal change has proven to be a challenging problem; for example, see (Rott 2000). In fact, not only should change be minimal, but also in a revision one would want to deal with only that (minimal) part of the knowledge base that is *relevant* to the formula for revision. Indeed, as information repositories grow progressively larger, it is crucial that a belief change be restricted to a "local" portion of a knowledge base. Consequently, an account of *relevance* in belief change is crucial, both conceptually as well as computationally.

In this paper we present an account of *relevance* in belief revision. We take relevance to be a *domain dependent* notion, in which assertions of (ir)relevance are part of the background theory of a domain. For example, we may have the assertion that, in the context of a party, whether Alice is there is independent of whether Bob is there (because they don't know each other). Hence if we revise by the fact that Alice was there, our beliefs about Bob's presence should remain unaffected. This is in contrast with most previous work on relevance in belief change, which addresses the case where an agent's contingent beliefs can be split between two disjoint languages. Our goal here is, in revising a belief set by a formula ϕ, to specify what part of the agent's beliefs are *irrelevant* with respect to ϕ; then these irrelevant beliefs will be among the beliefs that an agent will retain following revision. While our goals are foremost conceptual and representational, such an approach would have the potential for significant computational benefits in a practical system. Informally, given that a task (like revision) in the worst case grows exponentially more difficult with an increase in the size of a language, then one might expect an exponential improvement (in the best case) in restriciting revision to a portion of the knowledge base.

The above example with Bob and Alice illustrates another aspect of our approach, that assertions of relevance are generally not universal, but rather hold within a given context. The concept of conditional independence from probability theory provides a good analogy, given that independence assumptions in probability theory are generally not unconditional, and they assert something about the overall structure of a domain. However, as will become clear later, our approach does not really have much in common with (a qualitative counterpart to) probability theory, but rather will be closer to notions of independence in databases, specifically with multivalued dependencies.

In our approach, we begin with a standard AGM-style framework, enhanced with a set of context-dependent irrelevance assertions forming part of the background theory governing an agent's epistemic state. On the postulational side we give a postulate that states, for an assertion of irrelevance, that the agent's belief state following revision can be split into two parts, one reflecting those beliefs that have changed, and the other part consisting of beliefs unchanged from the initial belief set. We provide a semantic constraint on faithful rankings that exactly captures this postulate, and prove a representation result tying these two notions together. In a secondary approach, the worlds characterising the agent's belief set are *modified*, with the "irrelevant" part of a world left unchanged. To our knowledge this is the first time that revision has been based on the modification of possible worlds. We show that the two approaches coincide to the extent that in each case the agent's belief set is the same.

The next section collects notation used in the paper. This is followed by a summary of background material, including a brief introduction to belief revision and a survey of previous work in the area. Section 4 discusses issues and intuitions concerning relevance and our approach, while Section 5 gives the formal details. Section 6 discusses relevant related issues, after which there is a brief concluding section.

2 Notation

Define $P = \{A, B, C, \dots\}$ to be a finite nonempty set of propositional variables. In the approach, we need to distinguish the *name* of a propositional variable from its *value* (or truth assignment) in a model. Upper case English letters from the start of the alphabet are used for the former, and lower case for the latter. For a propositional variable such as A, one can think of a as $A \leftarrow true$ and $\neg a$ as $A \leftarrow false$.[1]

The symbol $\perp$ is not included in P and is used to denote contradiction; thus the truth value of $\perp$ is always set to false. The propositional language built from $P \cup \{\perp\}$ with the standard Boolean connectives is denoted by L. For a subset Y of P, we denote by L_Y the propositional language built from $Y \cup \{\perp\}$ with the standard Boolean connectives. Observe that even if $Y = \emptyset$, the sub-language L_Y is not empty since it contains sentences like $\perp$, $\neg\perp$, $\perp \vee \neg\perp$, etc.

For any sentence $\phi \in L$ we denote by P_ϕ the smallest subset of P with which we can write a sentence that is logically equivalent to ϕ (this set is unique). For example, for $\phi = a \wedge (b \vee \neg b)$, $P_\phi = \{A\}$. If ϕ is inconsistent or a tautology we take P_ϕ to be the empty set. L_ϕ is the language built from $P_\phi \cup \{\perp\}$ with the standard Boolean connectives.[2]

The deductive closure of a set of formulas Γ is denoted $\mathcal{C}n(\Gamma)$. Sometimes, for a specific set of formulas we will simply list the formulas, for example, writing $\mathcal{C}n(\{\phi, \psi\})$ as $\mathcal{C}n(\phi, \psi)$. An agent's set of beliefs will be modelled by a theory, also referred to as a *belief set*. Conventionally K (K', etc.) is used to denote a belief set. Thus a belief set is a set of formulas K such that $K = \mathcal{C}n(K)$.

For a theory K of L, and a set $Y \subseteq P$, K_Y denotes the restriction of K to L_Y; i.e. $K_Y = K \cap L_Y$. Clearly, since K is a theory, K_Y is also a theory. $\overline{Y}$ is the complement of Y with respect to P, that is, $\overline{Y} = P \setminus Y$. For any two sets of atoms $X, Y \subseteq P$ we will often use L_{XY} as an abbreviation of $L_{X \cup Y}$. Finally, for a sentence $\phi \in L$, by K_ϕ we shall denote the restriction of K to L_ϕ; i.e. $K_\phi = K \cap L_\phi$.

$\mathcal{M}$ is the set of models, or possible worlds, over P. Individual possible worlds are denoted by w, r, possibly super- or subscripted. (Mnemonically, r is used to denote worlds inconsistent with a belief set K whereas w is used for worlds consistent with K; however w and r are simply possible worlds.) For a formula (set of formulas) ϕ, $[\phi]$ is the set of worlds at which ϕ is true. We will identify a model with the set of literals satisfied at that model. Then, for possible world w, we identify w_X with the *reduct*[3] of w with respect to $X \subseteq P$, i.e. that part of w over X. For a set of possible worlds W, we define $W_X = \{w_X \mid w \in \mathcal{M}\}$. For a possible world w, we may refer to w_X as *the X part of* w, and similarly for belief sets, sets of worlds, etc.

3 Background

3.1 Belief Revision

In the AGM approach to belief change (Alchourrón, Gärdenfors, and Makinson 1985; Peppas 2008), an agent's beliefs are modelled by a belief set in which the underlying logic is assumed to contain classical propositional logic. This allows syntactic details, such as how a formula is expressed, to be ignored, and a focus placed on the logical contents of an agent's beliefs. In belief revision, a formula ϕ is incorporated into a belief set K such that the result is consistent (unless ϕ is itself inconsistent). Thus, if $K \vdash \neg\phi$, some formulas must be dropped from K before ϕ can be consistently added. In developing an account of revision, a simpler operator of *expansion* is first defined, where the expansion of K by ϕ is given by:

$$K + \phi \doteq \mathcal{C}n(K \cup \{\phi\}).$$

Formally, a revision operator $*$ maps a belief set K and formula ϕ to a revised belief set $K * \phi$. The AGM postulates for revision are as follows:

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

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

(K*3) $K * \phi \subseteq K + \phi$

(K*4) If $\neg\phi \notin K$ then $K + \phi \subseteq K * \phi$

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

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

(K*7) $K * (\phi \wedge \psi) \subseteq K * \phi + \psi$

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

Motivation for these postulates can be found in (Gärdenfors 1988; Peppas 2008).

[1] This notation is borrowed from relational databases, where the name of an attribute (such as *Age*) in a schema is distinguished from its value (like 35) in a tuple in a relation instance.

[2] Much of the notation and terminology used in this article is quite common in the belief revision literature. Part of our notation and terminology related to sub-languages has been adopted from (Parikh 1999) and/or (Peppas et al. 2015).

[3] The reduct is a standard notion from model theory (Hodges 1997; Chang and Keisler 2012), and should not be confused with the notion of *reduct* in answer set programmming.

Katsuno and Mendelzon (1991) have shown that a necessary and sufficient condition for constructing an AGM revision operator is that there is a function that associates a total preorder on the set of possible worlds with any belief set K, as follows:[4]

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

1. *If* $w_1 \in [K]$ *then* $w_1 \preceq_K w_2$
2. *If* $w_1 \in [K]$ *and* $w_2 \notin [K]$*, then* $w_1 \prec_K w_2$.

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

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

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

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

Since we will deal with revision with respect to a single belief set, for simplicity we drop the K subscript on a faithful ranking, and just use $\preceq$.

3.2 Belief Change and Relevance

The problem of relevance is an important, if not crucial, topic for many areas of AI; for example, a special issue of the Artificial Intelligence Journal on the topic (Subramanian, Greiner, and J. Pearl, editors 1997) contained papers dealing with relevance in causality, conditional independence, and inference, as would be expected, as well as in areas such as machine learning, game playing, and computer vision. As noted in the introduction, an account of relevance in belief change is required not just as a component in a conceptual framework, but is also essential for a general computational account: In a large knowledge base, it is totally infeasible to take the entirety of an agent's beliefs into consideration and, instead, belief change must be restricted to a local, relevant, part of a knowledge base.

In belief revision, Gärdenfors (1990) was the first to consider the problem of relevance; the key criterion he considered was the following:

If a belief state is revised by a sentence ϕ*, then all sentences in* K *that are independent of the validity of* ϕ *should be retained in the revised state of belief.*

Subsequently, Fariñas del Cerro and Herzig (1996) showed how AGM contraction and a given dependence relation are interdefinable, in much the same way as contraction and an epistemic entrenchment relation are interdefinable; this direction is further explored in (Oveisi et al. 2017).

[4] Katsuno and Mendelzon use formulas instead of belief sets. Since we deal with finite languages, the difference is immaterial.

Parikh's (1999) work on *language splitting* is perhaps the best-known approach to incorporating relevance in belief revision. Here, if a belief set K can be expressed in disjoint languages L_1 and L_2, then for epistemic input $\phi \in L_1$, revision can be restricted to the L_1 part of K. Parikh expresses this in the following postulate:

(P) If $K = Cn(\chi, \psi)$ where χ, ψ are sentences of disjoint sublanguages L_1, L_2 respectively, and $\phi \in L_1$, then $K * \phi = (Cn_{L_1}(\chi) \circ \phi) + \psi$, where $\circ$ is a revision operator of the sublanguage L_1.

Parikh's postulate encodes the intuition that an agent typically organises her beliefs into *independent compartments* that refer to different subject areas; for example, beliefs about life in Mars, would be held separately from beliefs about one's family history. Thus changes in one compartment should not effect other compartments.

Parikh's idea was extended in (Chopra and Parikh 2000) to allow limited interaction between compartments. Given a propositional language L, Chopra and Parikh define a *belief structure* $\mathcal{B}$ on L to be a set of ordered pairs $\mathcal{B} = \{(L_1, T_1), \ldots, (L_n, T_n)\}$, such that $\bigcup L_i = L$ and each T_i is a consistent theory over the sublanguage L_i. At first glance a belief structure $\mathcal{B}$ looks equivalent to a belief set K split into n "local" compartments $T_1, \ldots, T_n$; i.e. $K = Cn(T_1, \ldots, T_n)$. There are however two crucial differences between the two: (i) the sublanguages L_i need not be disjoint, and (ii) even with mutually inconsistent local theories T_i, the belief structure $\mathcal{B}$ does not necessarily collapse to absurdity. This is accomplished with the use of *four* truth values in evaluating the epistemic attitude of $\mathcal{B}$ towards a sentence ϕ; see (Chopra and Parikh 2000) for details.

For the revision of the belief structure $\mathcal{B} = \{(L_1, T_1), \ldots, (L_n, T_n)\}$ by a sentence ϕ, Chopra and Parikh offer two proposals, both of which are relevance-sensitive. The first approach revises each local theory T_i individually by the fraction of ϕ that can be expressed in L_i. More precisely, define the *i-shadow* of ϕ, denoted ϕ_i, to be a sentence in L_i such that $Cn(\phi_i) = Cn(\phi) \cap L_i$. Then $\mathcal{B} * \phi$ is defined as $\{(L_1, T_1 *_1 \phi_1), \ldots, (L_n, T_n *_n \phi_n)\}$, where each $*_i$ is a local revision function defined over the sublanguage L_i. A weakness of this approach is that the compartments never change; they remain fixed even in the face of new information that makes connections between two or more compartments. Moreover, breaking up the new information ϕ into its shadows $\phi_1, \ldots, \phi_n$ may lead to information loss; i.e. $\phi_1, \ldots, \phi_n$ is not always logically equivalent to ϕ.

The second approach is more flexible with compartments and does not suffer any information loss. Here, the revision of $\mathcal{B}$ by ϕ is performed in three steps. Firstly, local theories T_i that are relevant to ϕ are identified (i.e. all the T_is that share at least one variable with the minimal language of ϕ). Secondly, the relevant local theories are merged into a new theory H, and likewise their corresponding sublanguages L_i are merged in L_H. Thirdly, $(L_H, H * \phi)$ replaces all relevant pairs (L_i, T_i) in $\mathcal{B}$, while the remaining elements of $\mathcal{B}$ are left unchanged.

$\mathcal{B}$-structures and their revision methods are important extensions of Parikh's language splitting approach, but they lie outside the AGM framework. Firstly, they are based on a multivalued logic. Secondly, as noted in (Chopra and Parikh 2000), only the first six AGM postulates are taken into consideration; the remaining two AGM postulates, known as the *supplementary* postulates are omitted.

Extensions of Parikh's approach within the AGM framework have also been proposed. Kourousias and Makinson (2007) extend Parikh's result to the infinite case. Makinson (2009) shows that, while relevance is syntax-independent, it is not language-independent; as well the central definition is relaxed, allowing for extra-logical sources of relevance. In another direction, Peppas, Fotinopoulos, and Seremataki (2008) show that the Parikh postulate (P) is in conflict with the Darwiche and Pearl postulates for iterated belief revision. As well, (Peppas et al. 2015) provides a representation theorem for axiom (P) in terms of new constraints on faithful rankings. Moreover they show that axiom (P) is consistent with the full set of AGM postulates for revision. (Aravanis, Peppas, and Williams 2017) provides a characterisation of Parikh's axiom (P) in terms of constraints on an epistemic entrenchment, while Kern-Isberner and Brewka (2017) generalise the notion of syntax-spliting to epistemic states; they also study ordinal conditional functions (Spohn 1988) in this context.

3.3 Multivalued Dependencies

It proves to be the case that our approach is related to the notion of multivalued dependencies in relational database theory; see (Fagin 1977; Beeri, Fagin, and Howard 1977) for the original papers, or (Abiteboul, Hull, and Vianu 1995) or any of the many excellent texts on databases. In brief, a multivalued dependency is an assertion about a database relation that says that if the values of a set of attributes (given by X below) is fixed then the values of some attributes (Y below) are independent of the other attributes. The definition of a multivalued dependency is given as follows, expressed using our notation:

Definition 2 ((Abiteboul, Hull, and Vianu 1995)) *If P is a set of attributes, then a* multivalued dependency *over P is an expression of the form $X \twoheadrightarrow Y$, where $X, Y \subseteq P$. A relation W over P* satisfies *$X \twoheadrightarrow Y$ if $W \models XY \bowtie X(P \setminus Y)$, where $\bowtie$ is the natural join operator.*

The requirement $W \models XY \bowtie X(P \setminus Y)$ can be equivalently expressed by stipulating that for any pair of tuples that agree on the X values, say xy_1z_1 and xy_2z_2, if these tuples are in a relation instance W, then xy_1z_2 is also a tuple in W. Consequently, given a fixed value for X (viz. x), the Y part of the relation is independent of the Z part.

In practice, if a multivalued dependency $X \twoheadrightarrow Y$ holds over a relation R, this means that R can be split into two relations R_1 and R_2 over XY and $X(R \setminus Y)$ respectively such that no information is lost. As well, there is an axiomatisation of the class of multivalued dependencies.

4 The Approach: Intuitions

4.1 Relevance and Epistemic States

Parikh's (1999) approach and subsequent work assumes that (ir)relevance is a function of the agent's contingent beliefs: if it happens that the agent's beliefs can be expressed in disjoint languages, then this may be exploited in a revision. However, this criterion allows some instances where relevance via language splitting should not be employed. Consider the earlier example of a party for which the agent believes both Alice (a) and Bob (b) were there; thus $K = Cn(a, b)$. Now consider the following two cases:

1. Alice and Bob are a couple; they are *always* together. So the KB should reflect this, naively $K' = K + (a \equiv b)$.
2. Alice and Bob are total strangers; they have never met.

In the first case, a and b are clearly dependent, whereas in the second they are independent. However, obviously $K = K'$ in the first case, because we're dealing with belief sets. Thus, by Parikh's criterion, in both cases a and b are independent, and so if we learn that Alice was not there we would still conclude that Bob was present. In the first case this is clearly counterintuitive.

As well, there are cases where the agent's beliefs cannot be expressed in disjoint sublanguages and yet employing relevance in revision is justifiable. Consider the following example, where we assume that beliefs concerning a person's car are independent of beliefs regarding their phone. We believe that Daphne has a Honda or an iPhone; hence there is no splitting based on these contingent beliefs. If we later learn that Daphne has a different car, our beliefs concerning her phone(s) would presumably remain unchanged. Thus, Postulate (P) is of no help in dealing with relevance in this case. From these examples, we see that information in an agent's belief set is insufficient to determine relevance. We go beyond this and take relevance as a feature not of a (contingent) belief set, but as a feature of the agent's underlying belief state or epistemic state. Hence, relevance influences an agent's revision *policy*.

Last, it can be argued that for a sufficiently general language, there may be no instances in which (P) can be applied. To see this, consider a first-order setting and consider any two objects that appear to be unrelated, for example a bird somewhere in Canada, denoted $bird_{132}$, and the Greek philosopher Plato, denoted by the constant $plato$. These two individuals would seem to have nothing to do with each other, and so learning something about one should have nothing to do with the other. However, any theory with $bird_{132}$ and $plato$ will presumably have atomic formulas like $Bird(bird_{132})$ and $Human(plato)$. It would seem that these two facts should be independent with respect to any revision involving either. Yet, no language splitting into L_1 and L_2 with $Bird(bird_{132}) \in L_1$, $Human(plato) \notin L_1$ and with $Bird(bird_{132}) \notin L_2$, $Human(plato) \in L_2$ will allow the derivation $\neg Bird(plato)$. Consequently, language splitting appears to be too weak in this case to handle such presumably-unrelated assertions. That is, in *reasoning* that a bird and a historical individual belong to different categories, there must be language elements that allow one to reason from one notion to the other, and thus a chain of language elements relating the two notions. However, any such chain of course defeats any language splitting.

4.2 Independence Assertions

An *independence assertion*[5] or *irrelevance assertion* is an ordered pair (σ, Y) such that Y is a set of propositional variables ($Y \subseteq P$), and $\sigma \in L$ is a consistent sentence. We shall often denote the independence assertion (σ, Y) as $\sigma \twoheadrightarrow Y$. If $\sigma = \top$ we shall call $\sigma \twoheadrightarrow Y$ an *unconditional independence assertion*.

Intuitively, for the independence assertion $\sigma \twoheadrightarrow Y$, the formula σ specifies some precondition or context, and Y, being part of the signature, specifies some "subject matter". Then, the assertion $\sigma \twoheadrightarrow Y$ states that if $\sigma \in K$, then revision by any proposition $\phi \in L_Y$, where $\neg\phi \in K$, is independent of $L_{\overline{Y}}$. So this is more general than multivalued dependencies in databases (from where we have borrowed the notation) in which an assertion would be given as $X \twoheadrightarrow Y$; i.e. for *any* complete, consistent set of literals from X, Y is independent of $P \setminus Y$. It is similarly more general than Pearl's (1988) approach to independence.

Here is an example, adapted from (Pearl 1988). Let $P = \{A, J, M\}$. If there is a fire alarm (A), then Mary and John (who live within earshot) are expected to phone (M and J respectively). If there is no alarm then John or Mary calling are independent (since they don't know each other). Thus: $\neg a \twoheadrightarrow J$ (or equivalently $\neg a \twoheadrightarrow M$ – the two assertions have the same effect).

An important assumption in our approach is that relevance assertions are employed only when the formula for revision conflicts with the agent's belief set. This assumption is not required by the formalism, but rather is arguably more desirable. Consider the earlier example where if there is no alarm then John phoning is independent of Mary phoning, and where $K = \mathcal{C}n(\neg a, \neg j \vee \neg m)$. Assume that we learn that John called. There are two possibilities: First, the fact that John called is consistent with the agent's knowledge; and in this case the AGM approach stipulates that revision is the same as expansion. We would have that $K * j = \mathcal{C}n(\neg a, j, \neg m)$. Second, if we decide that the M part of the belief set is independent of the J part then, via our notion of irrelevance, we would have that $K * j = (K_{A,J} * j) + K_{A,M} = \mathcal{C}n(\neg a, j)$. The first alternative is preferable for several reasons: the result is more intuitive; it complies with the AGM approach; and it entails that relevance assertions are used only to guide a *change* in the agent's beliefs, and not an *expansion* of the beliefs.

Consider then where we have an assertion $\sigma \twoheadrightarrow Y$. The general idea is that in revising by $\phi \in L_Y$ where $K \vdash \sigma \wedge \neg\phi$, the $\overline{Y}$ part of the agent's belief set remains unchanged while revision only affects the Y part of the agent's beliefs. In fact, there are two ways in which this can be expressed semantically via a faithful ranking.

First, in a faithful ranking the $\overline{Y}$ part of the K worlds can be "duplicated" at the minimum ϕ worlds. This in turn will require that $\sigma \twoheadrightarrow Y$ have the properties of a multivalued dependency at this set of worlds (hence our adoption of this notation). This requires a constraint that can be seen as extending the conditions of a faithful ranking to the set of $\overline{Y}$ "subworlds". A representation result shows that the class of functions captured by the AGM postulates along with our relevance postulate is exactly that given by faithful rankings with the additional constraints.

The second alternative is simple but intuitive: In a revision, the Y part of a belief set is modified, and so the possible worlds characterising K are *modified* so that ϕ is believed. (This to our knowledge is the first time in which belief revision is carried out by "editing" possible worlds, rather than modifying a faithful ranking.) We show that the agent's resulting belief set is the same as that obtained via the first approach.

5 The Approach

We shall say that a revision function $*$ *complies* with the independence assertion $\sigma \twoheadrightarrow Y$ at a theory K, iff either $\sigma \notin K$ or for every consistent sentence ϕ, the following condition holds:

(R) If $\phi \in L_Y$ and $\neg\phi \in K$, then $K * \phi = (K * \phi)_Y + K_{\overline{Y}}$.

Thus for an independence assertion $\sigma \twoheadrightarrow Y$, if K contains σ and entails $\neg\phi$ where $\phi \in L_Y$, then the revision of K by ϕ is made up of $K * \phi$ restricted to L_Y, together with that part of K over the non-Y part of the language, that is, over $L_{\overline{Y}}$. For a set Θ of independence assertions we say that a revision function $*$ *complies* with Θ at a theory K iff $*$ complies with each one of the independence assertions in Θ at K. If $*$ complies with Θ at *all* theories of L, we shall simply say that $*$ complies with Θ. In Section 5.3 we show that the effects of axiom (P) can be captured with appropriately defined independence assertions; hence our work generalises Parikh's relevance-sensitive revision.

We say that a set of independence assertions Θ is *consistent* with the AGM postulates iff there is at least one AGM revision function that complies with it. In what follows we provide representation results for (R), and we study necessary and sufficient conditions under which a set of independence assertions Θ is consistent with the AGM postulates.

5.1 Characterisation via Faithful Preorders

Condition (R) can be characterised semantically in terms of the following constraints over faithful rankings. Let $\sigma \twoheadrightarrow Y$ be an independence assertion and assume that K contains σ; then we have:

(S1) If $r_Y = r'_Y$, $\neg r_Y \in K$, and $\neg r \notin K_{\overline{Y}}$, then $r \preceq r'$.

(S2) If $r_Y = r'_Y$, $\neg r_Y \in K$, $\neg r \notin K_{\overline{Y}}$, and $\neg r' \in K_{\overline{Y}}$, then $r \prec r'$.

Conditions (S1) - (S2) essentially say the following. If two worlds r, r' agree over the variables in Y and with none of the K-worlds, then their comparative plausibility depends on their agreement with the K-world over the remaining variables. In particular, (S1) says, if there exists a K-world that agrees with r over all $\overline{Y}$ variables (which follows from $\neg r \notin K_{\overline{Y}}$), then r is at least as plausible as r'. If in addition,

[5] The focus on *independence*, rather than *dependence*, assertions is discussed in Section 6.

(S2) says, there is no K-world that agrees with r' over the $\overline{Y}$ variables (which follows from $\neg r' \in K_{\overline{Y}}$), then r is in fact strictly more plausible than r'. Another way to think of conditions (S1) - (S2) is the following. Suppose that we partition the set of possible worlds into classes based on their values over Y; i.e. two worlds r, r' belong to the same class iff they agree over the variables in Y. Then (S1) and (S2) tell us that the minimal (or most plausible) worlds in each class are the ones that agree with some K-world over the remaining variables.

These conditions can be regarded as generalising those of a faithful ranking. In (S1) and (S2), if one takes $Y = \emptyset$ and assumes that $\neg r_{\emptyset} \in K$ holds vacuously, then these conditions reduce to the two conditions in Definition 1.

The next result shows that AGM revision functions augmented with Postulate (R) are captured exactly by faithful rankings that conform to (S1) and (S2).

Theorem 2 *Let $\sigma \twoheadrightarrow Y$ be an independence assertion, K a theory that contains σ, $\preceq$ a preorder faithful to K, and $*$ the revision function induced from $\preceq$ at K. Then $*$ satisfies (R) iff $\preceq$ satisfies (S1)-(S2).*

Proof
$(\Rightarrow)$
Assume that $*$ satisfies (R). For (S1) let r, r' be any two worlds such that $r_Y = r'_Y$, $\neg r_Y \in K$ and $\neg r \notin K_{\overline{Y}}$ (or equivalently $r \in [K_{\overline{Y}}]$). From (R) we derive that $[K * r_Y] = [(K * r_Y)_Y] \cap [K_{\overline{Y}}]$. Clearly, $[(K * r_Y)_Y] = [r_Y]$ and therefore $[K * r_Y] = [r_Y] \cap [K_{\overline{Y}}]$. Hence from $r \in [K_{\overline{Y}}]$ we derive that $r \in [K * r_Y]$. This again entails that $r \in min([r_Y], \preceq)$. Moreover from $r_Y = r'_Y$ it follows that $r' \in [r_Y]$. Hence from $r \in min([r_Y], \preceq)$ we derive that $r \preceq r'$ as desired.

The proof of (S2) is almost identical to that for (S1). In particular, consider any two worlds r, r' such that $r_Y = r'_Y$, $\neg r_Y \in K$, $\neg r \notin K_{\overline{Y}}$ (or equivalently $r \in [K_{\overline{Y}}]$), and $\neg r' \in K_{\overline{Y}}$ (or equivalently $r' \notin [K_{\overline{Y}}]$). Following exactly the same line of argument as before we conclude that $r \in min([r_Y], \preceq)$. Moreover, by (R), $[K * r_Y] = [(K * r_Y)_Y] \cap [K_{\overline{Y}}]$. Then from $r' \notin [K_{\overline{Y}}]$ we derive that $r' \notin min([r_Y], \preceq)$. Given that $r' \in [r_Y]$ (since $r_Y = r'_Y$) we then conclude that $r \prec r'$ as desired.
$(\Leftarrow)$
Assume that $\preceq$ satisfies (S1) - (S2). Consider any sentence $\phi \in L$. If ϕ is inconsistent, then (R) is trivially true. Assume therefore that $\phi \not\models \perp$, and moreover that $\phi \in L_Y$, and $K \models \neg\phi$. We show that $min([\phi], \preceq) = [\bigvee min([\phi], \preceq)_Y] \cap [K_{\overline{Y}}]$.

First we show that $LHS \subseteq RHS$. Consider any world $r \in min([\phi], \preceq)$. Clearly $r \in [\bigvee min([\phi], \preceq)_Y]$. Hence what's left to show is that $r \in [K_{\overline{Y}}]$. Assume on the contrary that $r \notin [K_{\overline{Y}}]$. Let w be any world in $[K]$. Define r' to be the world that agrees with r over Y and with w over $\overline{Y}$. From $\phi \in L_Y$ and $r \in [\phi]$ we derive that $r_Y \models \phi$. Since $r_Y = r'_Y$ we then derive that $r'_Y \models \phi$, or equivalently, that $\neg\phi \models \neg r'_Y$. Consequently, from $\neg\phi \in K$, we derive that $\neg r'_Y \in K$. Then, from (S2) we get that $r' \prec r$, which however contradicts $r \in min([\phi], \preceq)$. This shows that $min([\phi], \preceq) \subseteq [\bigvee min([\phi], \preceq)_Y] \cap [K_{\overline{Y}}]$.

For the converse, let r be any world in $[\bigvee min([\phi], \preceq)_Y] \cap [K_{\overline{Y}}]$. Since $r \in [\bigvee min([\phi], \preceq)_Y]$, there exists a world $r' \in min([\phi], \preceq)$ such that $r'_Y = r_Y$. Moreover, since $\phi \in L_Y$ and $r' \in [\phi]$, we derive that $r'_Y \models \phi$, or equivalently, that $\neg\phi \models \neg r'_Y$. Hence from $r'_Y = r_Y$ it follows that $\neg\phi \models \neg r_Y$. Then from $\neg\phi \in K$ we get $\neg r_Y \in K$. Hence from (S1) it follows that $r \preceq r'$. This again entails that $r \in min([\phi], \preceq)$ as desired. $\Box$

5.2 Characterisation via Local Revisions

We next consider a second way in which independence assertions may be used in a revision function. Informally, for a formula $\phi \in L_Y$ satisfying certain properties, the idea is that the Y part (i.e. over the subsignature $Y \subseteq P$) of the possible worlds in K are modified, while the $\overline{Y}$ part is left unchanged, with the result that ϕ is believed in the set of modified possible worlds.

Let $Q \subseteq P$ be a nonempty set of propositional variables and L_Q the sublanguage of L built over the variables in Q. We shall call any AGM revision function $\circ$ over L_Q a *local* revision function. Moreover, as a contrast to local revisions, we shall often call AGM revision functions defined over the entire language L, *global* revision functions.

Relevance and local revision are tightly connected. In particular, let $\sigma \twoheadrightarrow Y$ be an independence assertion and K a consistent theory satisfying σ. Moreover let $\circ$ be a local revision function defined over L_Y. Using $\circ$ we can (partially) construct a global revision function as follows:

(LR) If $\phi \in L_Y$ and $\neg\phi \in K$, then $K * \phi = (K_Y \circ \phi) + K_{\overline{Y}}$.

The above definition has a very intuitive appeal. It says that to revise K by ϕ one simply modifies the relevant part of K to accommodate ϕ, and then adds the non-relevant part of K. Viewed semantically, (LR) essentially says that the revision of K by ϕ is performed by taking each K-world w and modifying only its Y-part (to bring about ϕ), while leaving the rest of w unchanged.

Clearly (LR) and (R) look very similar. There is however a feature of relevance-sensitive revision that goes unnoticed with (R) and is made explicit with (LR): the local modifications to K (or to its worlds) to bring about ϕ, are performed in accordance to the same principles that govern global revision:

Theorem 3 *Let $\sigma \twoheadrightarrow Y$ be an independence assertion, $*$ a global revision function, and K a theory that contains σ. Then $*$ satisfies (R) at K iff there is a local revision function $\circ$ over L_Y such that (LR) holds at K for all consistent $\phi \in L_Y$.*

Proof
$(\Rightarrow)$
Assume that $*$ satisfies (R) at K. Define the function $\circ$ over the language L_Y as follows: $K_Y \circ \phi = (K * \phi)_Y$, for all $\phi \in L_Y$. Clearly (LR) follows immediately for all consistent $\phi \in L_Y$. Hence to complete this part of the proof we need to show that $\circ$ satisfies the AGM postulates. We

do so by constructing a faithful preorder $\preceq_Y$ to K_Y over the worlds in $\mathcal{M}_Y$, such that $[K_Y \circ \phi]_Y = min([\phi]_Y, \preceq_Y)$.[6]

Since $*$ is an AGM revision function, there exists a preorder over $\mathcal{M}$ that is faithful to K and such that $[K * \phi] = min([\phi], \preceq)$ for all $\phi \in L$. Based on $\preceq$ we construct $\preceq_Y$ as follows: for any $r, r' \in \mathcal{M}_Y$, $r \preceq_Y r'$, iff there is a $z \in [r]$ such that $z \preceq z'$, for all $z' \in [r']$.

It suffices to show that $\preceq_Y$ is a total preorder, it is faithful to K_Y, and moreover, for all $\phi \in L_Y$, $[K_Y \circ \phi]_Y = min([\phi]_Y, \preceq_Y)$.

Clearly $\preceq_Y$ is reflexive. For transitivity, let r, r', r'' be any three worlds in $\mathcal{M}_Y$ such that $r \preceq_Y r' \preceq_Y r''$. Let z, z', z'' be $\preceq$-minimal worlds in $[r]$, $[r']$, $[r'']$ respectively. Then, by the construction of $\preceq_Y$ we have that $z \preceq z' \preceq z''$; hence $z \preceq z''$. Since z'' is $\preceq$-minimal in $[r'']$, it follows that $r \preceq_Y r''$ as desired.

For totality, let r, r' be any two worlds in $\mathcal{M}_Y$ such that $r' \not\preceq_Y r$. Then for every $z' \in [r']$ there is a $z \in [r]$ such that $z' \not\preceq z$, or equivalently, $z \prec z'$. Let u be a $\preceq$-minimal element in $[r]$. We then derive that $u \prec z'$ for all $z' \in [r']$. Hence $r \preceq_Y r'$ as desired.

For faithfulness, if K_Y is inconsistent then faithfulness trivially holds. Assume therefore that K_Y is consistent. Then so is K (by the definition of K_Y). Consider now any world $r \in [K_Y]_Y$, and let r' be any world in $\mathcal{M}_Y$. It is not hard to verify that there exists a world $z \in [K]$ such that $z \in [r]$. Since $z \in [K]$, and given that $\preceq$ is faithful to K, it follows that $z \preceq z'$ for all $z' \in [r']$. Hence $r \preceq_Y r'$. Consider next a $r'' \in \mathcal{M}_Y$ such that $r'' \notin [K_Y]$. Then for all $z'' \in [r'']$, $z'' \notin [K]$. Hence, since $\preceq$ is faithful to K, it follows that $z \prec z''$ for all $z'' \in [r'']$. This again entails that $r \prec_Y r''$. Thus $\preceq_Y$ is faithful to K_Y as desired.

Finally we prove that $[K_Y \circ \phi]_Y = min([\phi]_Y, \preceq_Y)$, or equivalently that $[(K * \phi) \cap L_Y]_Y = min([\phi]_Y, \preceq_Y)$. It is not hard to verify that $[(K * \phi) \cap L_Y]_Y = \{z \cap L_Y : z \in [K * \phi]\}$. Hence, since $[K * \phi] = min([\phi], \preceq)$, it suffices to show that $\{z \cap L_Y : z \in min([\phi], \preceq)\} = min([\phi]_Y, \preceq_Y)$.

Assume that $r \in \{z \cap L_Y : z \in min([\phi], \preceq)\}$; i.e., there is a $z \in min([\phi], \preceq)$ such that $r = z \cap L_Y$. Clearly, $r \in \mathcal{M}_Y$. Moreover, since $\phi \in L_Y$ and $z \models \phi$, it follows that $r \in [\phi]_Y$. Consider now any $r' \in [\phi]_Y$ and let z' be any world in $[r']$. Since $r' \models \phi$ it follows that $z' \models \phi$. Hence since z is $\preceq$-minimal in $[\phi]$ we derive that $z \preceq z'$. Hence $r \preceq_Y r'$, and therefore, since r' was chosen as an arbitrary member of $[\phi]_Y$, it follows that $r \in min([\phi]_Y, \preceq_Y)$. Thus we have shown that $\{z \cap L_Y : z \in min([\phi], \preceq)\} \subseteq min([\phi]_Y, \preceq_Y)$.

For the converse, let r be any element of $min([\phi]_Y, \preceq_Y)$. Let u be a $\preceq$-minimal element of $[r]$. Clearly, since $r \models \phi$, so does u; i.e. $u \in [\phi]$. Moreover by construction $r = u \cap L_Y$. Hence to show that $r \in \{z \cap L_Y : z \in min([\phi], \preceq)\}$ it suffices to show that u is $\preceq$-minimal in $[\phi]$. Let u' be an arbitrary world in $[\phi]$. Define $r' = u' \cap L_Y$. Since $u' \models \phi$ and $\phi \in L_Y$, it follows that $r' \in [\phi]_Y$. Hence $r \preceq_Y r'$. This again entails that $u \preceq u'$. Since u' was chosen as an arbitrary ϕ-world, it follows that u is $\preceq$-minimal in $[\phi]$ as desired.

($\Leftarrow$)

Assume that there is a local revision function $\circ$ over the language L_Y such that (LR) is satisfied at K. To prove that (R) holds at K it suffices to show that for all $\phi \in L_Y$, $(K * \phi)_Y = K_Y \circ \phi$.[7] To this aim, observe that for any $A \subseteq L_Y$, any $B \subseteq L_{\overline{Y}}$, and any sentence $\psi \in L_Y$, if $\psi \notin Cn(A)$, then $\psi \notin A + B$.

Now, for $(K * \phi)_Y \subseteq K_Y \circ \phi$, let ψ be any sentence in $(K * \phi)_Y$. Then clearly, $\psi \in L_Y$ and $\psi \in K * \phi$. Therefore by (LR), $\psi \in (K_Y \circ \phi) + K_{\overline{Y}}$. Since $\psi \in L_Y$, by the observation above, we derive that $\psi \in K_Y \circ \phi$. Hence we have shown that $(K * \phi)_Y \subseteq K_Y \circ \phi$.

For the converse, let ψ be any sentence in $K_Y \circ \phi$. Clearly $\psi \in L_Y$. Moreover from (LR), it follows that $\psi \in K * \phi$. Therefore, since $\psi \in L_Y$, it follows immediately that $\psi \in (K * \phi)_Y$ as desired. $\square$

5.3 Relationship with Parikh's Postulate

In this section we examine the relationship between independence assertions and Parikh's approach to relevance-sensitive revision. We recall that there are two different readings to the axiom proposed by Parikh (see (Peppas et al. 2015)). Herein we shall focus on the weaker version of Parikh's axiom, which is more general and less controversial:[8]

(wP) If $K = Cn(\chi, \psi)$, $L_\chi \cap L_\psi = \emptyset$, and $\phi \in L_\chi$, then $(K * \phi) \cap L_\psi = K \cap L_\psi$.

To relate (wP) to our condition (R), we need to borrow some more definitions from (Parikh 1999): For a theory K, we say that a partition $S = \{S_1, \ldots, S_n\}$ of P is a K-*splitting* iff there exist sentences $\alpha_1, \ldots, \alpha_n$, such that $\alpha_i \in L_{S_i}$ for all $1 \leq i \leq n$, and moreover $K = Cn(\alpha_1, \ldots, \alpha_n)$. We call S the *finest* K-splitting iff S is a K-splitting and it refines every other K-splitting.[9] Parikh, (Parikh 1999), has shown that for every theory K there is only one finest K-splitting.

To every partition $S = \{S_1, \ldots, S_n\}$ of P we shall assign the following set $\gamma(S)$ of independence assertions:

$$\gamma(S) = \{\top \twoheadrightarrow Y : \text{ for some } Q \subseteq S,\ Y = \cup Q\}.$$

In other words, $\gamma(S)$ consists of all independence assertions $\top \twoheadrightarrow Y$ such that Y is the union of some elements of S.

Theorem 4 *Let K be a theory, $S = \{S_1, \ldots, S_n\}$ the finest K-splitting, and $*$ an AGM revision function. Then $*$ satisfies (wP) at K iff $*$ complies with $\gamma(S)$ at K.*

Proof.

($\Rightarrow$)

Assume that $*$ satisfies (wP) at K, and let $\top \twoheadrightarrow Y$ be any independence assertion in $\gamma(S)$. By construction, there

[6] Recall that $\mathcal{M}$ denotes the set of possible worlds over the language L, while $\mathcal{M}_Y$ denotes the set of possible worlds over the language L_Y. Likewise, for a set of possible worlds U, by U_Y we denote the restriction (or: reduct) of the U-worlds to the sublanguage L_Y.

[7] Recall that by definition, $K_Y \circ \phi$ is a theory of L_Y; hence $K_Y \circ \phi = (K_Y \circ \phi)_Y$.

[8] In (Peppas et al. 2015), axiom (wP) is called (R1).

[9] A partition S refines another partition S', iff for every $S_i \in S$, there is $S'_j \in S'$, such that $S_i \subseteq S'_j$ (see (Peppas et al. 2015)).

is a $Q \subseteq S$ such that $Y = \cup Q$. Since S is a K-splitting, it follows that there exist sentence χ, ψ such that $\chi \in L_Y$, $\psi \in L_{\overline{Y}}$, and $K = Cn(\chi, \psi)$. Consider now any sentence $\phi \in L_Y$ such that $\neg\phi \in K$. From (wP) it follows that $(K * \phi)_{\overline{Y}} = K_{\overline{Y}}$. Moreover, (wP) entails that $\{Y, \overline{Y}\}$ is also a $K * \phi$-splitting,[10] and therefore, $K * \phi = (K * \phi)_Y + (K * \phi)_{\overline{Y}}$. Consequently we derive that $K * \phi = (K * \phi)_Y + K_{\overline{Y}}$ as desired.

($\Leftarrow$)

Assume that $*$ complies with $\gamma(S)$ at K. We show that $*$ satisfies (wP) at K. Let χ, ψ be any two sentences such that $K = Cn(\chi, \psi)$ and $L_\chi \cap L_\psi = \emptyset$. Then clearly $\{P_\chi, \overline{P_\chi}\}$ is a K-splitting and consequently there is a subset Q of the finest K-splitting S, such that $P_\chi = \cup Q$. Define $Y = P_\chi$. Clearly then the independence assertion $\top \twoheadrightarrow Y$ belongs to $\gamma(S)$, and therefore $*$ complies with $\top \twoheadrightarrow Y$ at K. Consider now any sentence $\phi \in L_\chi$. To prove (wP) we need to show that $(K * \phi)_{\overline{Y}} = K_{\overline{Y}}$.

If $\neg\phi \notin K$ then $K*\phi = K+\phi$, and consequently $K*\phi = Cn(\chi \wedge \phi, \psi)$. Since $\chi \wedge \phi \in Y$, we derive immediately that $(K * \phi)_{\overline{Y}} = Cn(\psi)_{\overline{Y}} = K_{\overline{Y}}$ as desired. Assume therefore that $\neg\phi \in K$. Then since $*$ complies with $\top \twoheadrightarrow Y$ at K, we derive that $K * \phi = (K * \phi)_Y + K_{\overline{Y}}$. Hence, once again, $(K * \phi)_{\overline{Y}} = K_{\overline{Y}}$ as desired. $\square$

5.4 Relationship with (Database) Multivalued Dependencies

It is not too surprising that there are connections between our approach and that of multivalued dependencies in relational databases. Specifically, a faithful ranking corresponds to a total preorder over the set of possible worlds $\mathcal{M}$, where the worlds in each set are equally plausible. Each such set of worlds can be thought of as an instance of a database relation. the relation's set of attributes is given by the (names of the) propositional variables in P, and each world can be thought of as a relation tuple. Then it proves to be the case that each such set of worlds, with the exception of $[K]$, which was excluded by fiat, conforms to each of the independence assertions, viewed as multivalued dependencies. This is expressed in the following result:

Theorem 5 *Let K be a theory that contains σ, and $*$ a revision function that complies with $\sigma \twoheadrightarrow Y$. Let X be the minimal set of atoms required to express σ, and let $x \in [\sigma]_X$. Then for any $\phi \in L_Y$ where $\neg\phi \in K$ we have that $xy_1z_1, xy_2z_2 \in [K * \phi]$ implies that $xy_1z_2 \in [K * \phi]$.*

Proof

Let K be a theory that contains σ, $*$ a revision function that complies with $\sigma \twoheadrightarrow Y$, and $\neg\phi \in K$ where $\phi \in L_Y$. Let $x \in [\sigma]_X$ where X is the minimal set of atoms required to express σ. Assume that $xy_1z_1, xy_2z_2 \in [K * \phi]$.

We have from (R) that $K * \phi = (K * \phi)_Y + K_{\overline{Y}}$, and so $[K * \phi] = [(K * \phi)_Y + K_{\overline{Y}}]$, or

$$[K * \phi] = [(K * \phi)_Y] \cap [K_{\overline{Y}}]. \qquad (1)$$

[10]This follows immediately from the results reported in (Peppas et al. 2015)

Since $xy_1z_1 \in [K*\phi]$ we obtain that $xy_1z_1 \in [(K*\phi)_Y]$, and so $y_1 \in [(K * \phi)_Y]_Y$.[11] However, $y_1 \in [(K * \phi)_Y]_Y$ implies that $xy_1z_2 \in [(K * \phi)_Y]$.

Similarly $xy_2z_2 \in [K * \phi]$ implies via (1) that $xy_2z_2 \in [K_{\overline{Y}}]$. Consequently $xz_2 \in [K_{\overline{Y}}]_{\overline{Y}}$, which in turn implies that $xy_1z_2 \in [K_{\overline{Y}}]$.

We thus have that $xy_1z_2 \in [(K*\phi)_Y]$ and $xy_1z_2 \in [K_{\overline{Y}}]$, which via (1) implies that $xy_1z_2 \in [K * \phi]$, which was to be shown. $\square$

Our notion of independence assertions is more general than that of database multivalued dependencies, since for $\sigma \twoheadrightarrow Y$, σ may be an arbitrary formula. However, the above result indicates that in the case where, for every independence assertion $\sigma \twoheadrightarrow Y$, σ is equivalent to a conjunction of literals, we inherit the formal results concerning multivalued dependencies. (How these results may extend to an arbitrary formula σ is a topic for future research.) In particular, there is an axiomatisation of such dependencies, and so one can talk about a dependency being derived from others. Moreover, for a set of dependencies Θ, and set of atoms X, the set $dep(X) = \{Y \subseteq P \mid \Theta \vdash X \twoheadrightarrow Y\}$ is a Boolean algebra of sets for P. Hence for each such X, one just needs to keep track of the $\subseteq$-minimal elements of $dep(X)$. For details, see any of the references of Section 3.3.

6 Discussion

In this section we consider two pragmatic issues having to do with the approach. First we discuss the emphasis on encoding *irrelevance*, as opposed to *relevance* assertions. Second, since irrelevance assertions must be specified as part of a background theory, the question arises as to where such assertions may come from.

Why irrelevance (and not relevance) assertions? There are two broad ways in which (ir)relevance properties can be specified: one can state what is relevant, and then everything not mentioned is irrelevant; or one can specify what is irrelevant, and then everything else is potentially relevant. An example of the first category is Bayesian nets, where the structure of a network determines what is relevant to what (i.e. conditionally dependent), and then from this (via the Bayesian network independence assumption or using the Markov blanket, e.g. (Russell and Norvig 2010)), one can determine certain conditional independence relations. An example of the second category is multivalued dependencies in database systems.

Our approach is in the second category. The main reason for this is that it is more conservative. The AGM approach is intended to be as general as possible with respect to belief change, and our approach is in keeping with this stance. That is, our approach allows one to incorporate further domain-specific information via independence assertions, to further constrain the class of rational revision functions. However, if, in encoding a domain, an assertion of irrelevance is missed then one simply has less domain-specific

[11]Recall for a set of possible worlds W that W_Y is the set of worlds over subsignature Y.

information to work with. This is in contrast with a Bayesian network, for example, where if a causal relation is missed, then the set of inferable conditional independence relations will increase.

Where do assertions of (ir)relevance come from? Notions of independence and (ir)relevance are an important aspect of information regarding a domain. Hence in our view, in constructing a rational, knowledge-based agent that can change its beliefs, specifying such assertions will be an important part of the axiomatisation of the underlying domain. An important question then is, where do these assertions come from and how may they be determined? Fortunately, there is no lack of possibilities:

1. As a base case, independence assertions may be specified directly, based on knowledge of the domain, as is done in relational databases – for example, that information about a person's phone is independent of their cars.
2. Independence assertions may be determined from a given physical or causal theory. For example, Amir and McIlraith (2005) give an account of an espresso maker, where variables $water$ and $steam$ in their theory allow a theory to be partitioned; for our approach, fixed values of these variables will in turn yield independence assertions.
3. Such assertions can be extracted from a structure that contains implicit independence information, for example, from Bayesian or Markov networks. Thus, in a Bayesian net, independence is given via the Bayes net independence assumption and the Markov blanket, notions that can readily be translated into independence assertions.

   Similarly, in a *belief base*, which need not be deductively closed, the way in which information is represented may indicate what items are related to what. For example, the belief bases $\{a, b\}$ and $\{a \wedge b\}$ express the same information, but the first may provide evidence that a and b are unrelated, whereas in the second $a \wedge b$ constitutes a single, perhaps indivisible, item of information.
4. Much more speculatively, it may be that in a nonclassical logic or fragment of a classical logic, independence may be assumed (perhaps by default) in the structure of a knowledge base.

   Specifically, in a description logic, independence may be implicit in some fashion with respect to the subsumption hierarchy or role structure. Consider for example a structure such as

   $$(\exists R_1.C_1) \sqcap (\exists R_2.C_2)$$

   which expresses the concept that is R_1 related to concept C_1, and similarly for R_2, C_2. It may be that one can assert or assume that a change to the first term is independent of the second.

These considerations give rise to a related issue, that of determining the *scope* of an independence assertion; i.e. the set of atoms that ought to appear in its right hand side. Consider the example above, where we take information about a person's phone to be independent of that about their cars. The notion of "information about a person's phone" includes items such as the phone's colour, make, and battery level. Hence a correct encoding of the independence of a person's phone from his/her car, would require an independence assertion whose right hand side includes all of these items. Specifying such domain-specific related information may be nontrivial, but in principle it can be specified, as indeed is done in the admittedly-simpler context of relational database systems.

Moreover, here again it may be possible to automate the extraction of such related information. Thus, in Item 2 above, fixing the values of $water$ and $steam$ results in a partitioned theory in which the propositional atoms in a partition would make up a suitable "subject" for the right-hand side of an independence assertion, yielding an assertion, for example, of the form $steam \twoheadrightarrow \langle\text{set of atoms}\rangle$. Similarly, in Item 3 above, the Bayesian network independence assumption — that for a variable X, given values for its parents, that variable is independent of its nondescendants — would seem to lead to independence assertions of the form $parents(X) \twoheadrightarrow descendants(X)$.

So to conclude, What we have done in this paper is to argue that independence assertions are necessary in a general account of belief revision, and to show how a set of independence assertions may be incorporated in the AGM approach. However, a full account of relevance in AI and in reasoning in general remains the subject of ongoing research.

7 Conclusion

In belief change, as in many areas of AI, *relevance* or *independence* is an important notion, both conceptually and computationally. For belief revision, we take relevance to be a guiding principle when a formula for revision is inconsistent with an agent's beliefs. (If a formula is consistent with an agent's beliefs, then the success postulate seems to say all that needs to be said.)

We have argued that *relevance* is a domain-specific notion, and that (ir)relevance assertions are given as part of the agent's epistemic state, expressing knowledge regarding the structure of the domain. Such assertions need not hold universally, but rather apply in a given context, specifying for $\sigma \twoheadrightarrow Y$ that "in the case where σ holds, the Y part of the agent's epistemic state is independent of the rest of the epistemic state". In the main approach, conditions were placed on a faithful ranking on possible worlds to enforce the (ir)relevance assertions. As well, corresponding postulates and a representation result were given. In a second approach, the possible worlds characterising the agent's beliefs are modified to effect revision. Interestingly, these approaches, based on somewhat differing intuitions, nonetheless yield the same belief set following revision.

The approach to relevance can also be seen as contributing to the desiderata of *minimal change* underlying the AGM approach. That is, the AGM approach provides constraints on how a rational agent may revise its beliefs; independence assertions give a means for specifying additional domain-specific constraints on a revision function. Hence the AGM approach is augmented and, by allowing the specification of additional domain-specific information, the approach may provide a step toward the development of specific, practical revision systems.

Acknowledgements

The authors are grateful to Theofanis Aravanis for valuable comments on this work. The first author gratefully acknowledges financial support from the Natural Sciences and Engineering Research Council of Canada.

References

Abiteboul, S.; Hull, R.; and Vianu, V. 1995. *Foundations of Databases: The Logical Level*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.

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

Amir, E., and McIlraith, S. A. 2005. Partition-based logical reasoning for first-order and propositional theories. *Artificial Intelligence* 162(1-2):49–88.

Aravanis, T.; Peppas, P.; and Williams, M.-A. 2017. Epistemic entrenchment characterization of Parikh's axiom. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 772–778.

Beeri, C.; Fagin, R.; and Howard, J. H. 1977. A complete axiomatization for functional and multivalued dependencies in database relations. In *Proceedings of the 1977 ACM SIGMOD International Conference on Management of Data*, SIGMOD '77, 47–61. New York, NY, USA: ACM.

Chang, C. C., and Keisler, H. J. 2012. *Model Theory*. Dover Publications, third edition.

Chopra, S., and Parikh, R. 2000. Relevance sensitive belief structures. *Annals of Mathematics and Artificial Intelligence* 28(1-4):259–285.

Fagin, R. 1977. Multivalued dependencies and a new normal form for relational databases. *ACM Transactions on Database Systems* 2(3):262–278.

Fariñas del Cerro, L., and Herzig, A. 1996. Belief change and dependence. In *Proceedings of the Sixth conference on Theoretical Aspects of Rationality and Knowledge*, 147–161. Morgan Kaufmann.

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

Gärdenfors, P. 1990. Belief revision and relevance. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1990, 349–365.

Hodges, W. 1997. *A Shorter Model Theory*. Cambridge, UK: Cambridge University Press.

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

Kern-Isberner, G., and Brewka, G. 2017. Strong syntax splitting for iterated belief revision. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1131–1137.

Kourousias, G., and Makinson, D. 2007. Parallel interpolation, splitting, and relevance in belief change. *The Journal of Symbolic Logic* 72(3):994–1002.

Makinson, D. 2009. Propositional relevance through letter-sharing. *Journal of Applied Logic* 7(4):377–387. Special Issue: Formal Models of Belief Change in Rational Agents.

Oveisi, M.; Delgrande, J. P.; Pelletier, J.; and Popowich, F. 2017. Kernel contraction and base dependence. *Journal of Artificial Intelligence Research* 60:97–148.

Parikh, R. 1999. Beliefs, belief revision, and splitting languages. In Moss, L.; Ginzburg, J.; and de Rijke, M., eds., *Logic, Language and Computation, Vol 2*. CSLI Publications. 266–278.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Mateo, CA: Morgan Kaufman.

Peppas, P.; Williams, M.-A.; Chopra, S.; and Foo, N. 2015. Relevance in belief revision. *Artificial Intelligence* 229:126–138.

Peppas, P.; Fotinopoulos, A.; and Seremetaki, S. 2008. Conflicts between relevance-sensitive and iterated belief revision. In Ghallab, M.; Spyropoulos, C. D.; Fakotakis, N.; and Avouris, N. M., eds., *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI-08)*, 85–88. IOS Press.

Peppas, P. 2008. Belief revision. In van Harmelen, F.; Lifschitz, V.; and Porter, B., eds., *Handbook of Knowledge Representation*. San Diego, USA: Elsevier Science. 317–359.

Rott, H. 2000. Two dogmas of belief revision. *Journal of Philosophy* 97:503–522.

Russell, S. J., and Norvig, P. 2010. *Artificial Intelligence: A Modern Approach*. Pearson Education, third edition.

Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume II. Kluwer Academic Publishers. 105–134.

Subramanian, D.; Greiner, R.; and J. Pearl, editors. 1997. Special issue on relevance. *Artificial Intelligence* 97(1-2).