

Using Conditional Independence for Belief Revision

Matthew James Lynn,¹ James P. Delgrande,¹ Pavlos Peppas²

¹Simon Fraser University, Canada

²University of Patras, Greece

mlynn@cs.sfu.ca, jim@cs.sfu.ca, pavlos@upatras.gr

Abstract

We present an approach to incorporating qualitative assertions of conditional irrelevance into belief revision, in order to address the limitations of existing work which considers only unconditional irrelevance. These assertions serve to enforce the requirement of minimal change to existing beliefs, while also suggesting a route to reducing the computational cost of belief revision by excluding irrelevant beliefs from consideration. In our approach, a knowledge engineer specifies a collection of multivalued dependencies that encode domain-dependent assertions of conditional irrelevance in the knowledge base. We consider these as capturing properties of the underlying domain which should be taken into account during belief revision. We introduce two related notions of what it means for a multivalued dependency to be taken into account by a belief revision operator: partial and full compliance. We provide characterisations of partially and fully compliant belief revision operators in terms of semantic conditions on their associated faithful rankings. Using these characterisations, we show that the constraints for partially and fully compliant belief revision operators are compatible with the AGM postulates. Finally, we compare our approach to existing work on unconditional irrelevance in belief revision.

1 Introduction

Belief revision is concerned with the situation in which an agent is confronted with a new fact to incorporate into its belief set. If the new fact is inconsistent with the current belief set, the challenge is to revise these beliefs so that as many of the current beliefs as possible are retained while incorporating the new fact and maintaining consistency. This process is formalised as a *belief revision operator* $*$ which takes a current knowledge base K and a formula for revision ϕ and produces a revised knowledge base $K * \phi$.

In order to formalise the requirement that revision should result in a minimal change to existing beliefs, a number of authors have turned to *irrelevance*, suggesting that those beliefs irrelevant to the formula for revision should remain unchanged (Gardenfors 1990). This also has the potential advantage of opening a pathway to more efficient belief revision operators, by being able to exclude irrelevant beliefs from the revision process. However, so far, these notions of

irrelevance have been extremely strict, considering beliefs as irrelevant only when there is no connection, however indirect, between them.

To see the issue, consider the following situation: an agent is informed that refrigerators require power, power is generated in the local area by wind turbines, and wind turbines kill birds. It would seem that information about birds would be independent of information concerning refrigerators; however, this is not the case, given the link between refrigerators and birds mediated by wind turbines. Consequently, existing approaches would consider refrigerators relevant to birds. However, when revising our beliefs about birds there would seem to be no reason for our beliefs about refrigerators to change. Hence, it seems we need a more nuanced and less restrictive notion of irrelevance.

This situation has a parallel in probability theory. In practice, random variables are rarely independent. However, they are frequently conditionally independent. As a result, Bayesian networks have been developed to exploit conditional independence properties, thereby overcoming the otherwise seemingly-intractable complexity of probabilistic inference (Pearl 2014).

In this paper we take a suitable analogue of conditional independence for determining which beliefs may be considered irrelevant to others in a given context. We then apply this notion to belief revision, and we study those revision operators which comply with this formulation of conditional independence. Our approach is given in terms of the Katsuno-Mendelzon approach for belief revision. In our approach, we assume that conditional independence is a property of the underlying domain, and we consequently assume that a knowledge engineer has provided a collection of such conditional independence assertions. These assertions can then be taken into account in the belief revision process. To this end, we study two related notions of what it means for a belief revision operator to take into account conditional independencies. We provide postulates that characterise conditional independence in revision, and which generalise previous approaches to (non-conditional, absolute) independence. Furthermore, we provide representation results, giving conditions on faithful rankings which correspond to the sets of postulates characterising conditional independence in revision.

The next section covers background material: we first

present useful definitions and notation, after which we give background material on belief revision, including existing approaches to independence in belief revision, along with conceptions of conditional independence in logic. Section 3 shows that multivalued dependencies can be used to represent background knowledge of conditional irrelevance, and introduces the classes of *partially compliant* and *fully compliant* belief revision operators which take these into account. In Section 4 we provide representation results for partially and fully compliant belief revision operators in terms of semantic conditions on faithful rankings. Finally, Section 5 discusses related work, and future work, after which we have a brief conclusion.

2 Background Material

2.1 Preliminaries and Notation

Let $V = \{p, q, r, \dots\}$ be a finite set of propositional variables, arbitrary subsets of which are denoted by X, Y , and Z . We sometimes juxtapose these subsets to represent unions, e.g. $XY = X \cup Y$. The relative complement $V - X$ will be denoted by \bar{X} . Every subset X of V induces a propositional language $L(X)$ consisting of formulae constructed from the elements of X by applying the propositional connectives \neg, \wedge, \vee , and \rightarrow . We write L for the entire propositional language $L(V)$.

Lower case Greek letters $\phi, \psi, \gamma, \dots$ will be used to range over formulae in a propositional language, with K playing a special role of a formula thought of as representing the knowledge base of an agent.

Also associated to every subset X of V is the set Ω_X of functions $v : X \rightarrow \{T, F\}$ referred to as *models* or *possible worlds* over X . We will freely think of these possible worlds as either these functions, or as conjunctions of the literals satisfied by them. Hence, for us, $\{x \mapsto T, y \mapsto F\}$ is the same thing as $x \wedge \neg y$. Given a possible word u over V alongside a subset X of V , we write u_X for the reduct of u to a possible world over X , that is the unique function $u_X : X \rightarrow \{T, F\}$ agreeing with u on X .

When ϕ is a formula we write $[\phi]$ for the set of models over V satisfying ϕ , so that $[\phi] \subseteq \Omega_V$. We write $\phi \vdash \psi$ to indicate $[\phi] \subseteq [\psi]$, and $\phi \equiv \psi$ to indicate $[\phi] = [\psi]$.

We write $V(\phi)$ for the minimal set of propositional variables for which there exists a formula ψ logically equivalent to ϕ containing only occurrences of variables in $V(\phi)$, for instance $V(q \wedge (p \vee \neg p)) = \{q\}$.

2.2 Projections of a Propositional Formula

In order to speak about components of a knowledge base K expressed in various subvocabularies we will introduce the following analogue of the *projection operator* from the relational algebra (Abiteboul, Hull, and Vianu 1995).

Definition 2.1. *If ϕ is a propositional formula, and $X \subseteq V$, then the **projection** ϕ_X of ϕ onto X is defined up to logical equivalence as the formula ϕ_X such that*

$$[\phi_X] = \{u \in \Omega_V \mid \exists v \in [\phi], v_X = u_X\}.$$

Example 2.1. *The projection of $(p \rightarrow q) \wedge (q \rightarrow r)$ onto $\{p, q\}$ is $(p \rightarrow q)$, whereas the projection of $(p \rightarrow q) \wedge (q \rightarrow r)$ onto $\{q, r\}$ is $(q \rightarrow r)$.*

Regarding a set of possible worlds as tuples in a relation, it follows that ϕ_X defines the set of worlds resulting from projecting this “relation” onto the “attributes” in X , then taking the Cartesian product of this with all possible interpretations of the remaining variables. This operator also appears as the notion of a uniform interpolant, a model-theoretic reduct (Hodges 1993), or as the dual of a forgetting operator¹ (Delgrande 2017). For our purposes, we will rely on the following property of projections:

Theorem 2.1. *If $\phi \vdash \psi$ and $V(\psi) \subseteq X$ then $\phi \vdash \phi_X$ and $\phi_X \vdash \psi$.*

2.3 Revision Operators and Faithful Rankings

A belief revision operator, as formalised by Alchourron, Gärdenfors, and Makinson (1985), is a binary function $*$ which maps a belief set K and a formula ϕ and produces a revised belief set $K * \phi$ in a manner satisfying the *AGM postulates*. These postulates attempt to capture the requirement that $K * \phi$ must include ϕ alongside as many beliefs from K as possible, while maintaining consistency. In other words, $K * \phi$ results from a minimal change to the existing belief set K which results in ϕ being believed. Note that belief revision captures an agent revising its beliefs about the present state of affairs, whereas updating its beliefs when the state of the world changes is the subject of *belief update operators*, cf. (Katsuno and Mendelzon 1991a) or (Peppas 2008).

In our setting of a finite vocabulary, we can simplify matters by working instead with the Katsuno-Mendelzon approach wherein the belief sets K and $K * \phi$ are represented as single formulas, and the AGM postulates are rephrased in the following manner (Katsuno and Mendelzon 1991b).

Definition 2.2. *A binary function $* : L \times L \rightarrow L$ is a **belief revision operator** if it satisfies the following **Katsuno-Mendelzon postulates**:*

- R1.** $K * \psi \vdash \psi$;
- R2.** *If $K \wedge \phi$ is satisfiable then $K * \phi \equiv K \wedge \phi$;*
- R3.** *If ϕ is satisfiable then $K * \phi$ is satisfiable;*
- R4.** *If $K_1 \equiv K_2$ and $\phi_1 \equiv \phi_2$ then $K_1 * \phi_1 \equiv K_2 * \phi_2$;*
- R5.** $(K * \phi) \wedge \psi \vdash K * (\phi \wedge \psi)$;
- R6.** *If $(K * \phi) \wedge \psi$ is satisfiable then $K * (\phi \wedge \psi) \vdash (K * \phi) \wedge \psi$.*

It is worthwhile noting that some authors only require a belief revision operator to satisfy the *basic postulates* (R1) through to (R4), and refer to (R5) and (R6) as the *supplementary postulates*.

Katsuno and Mendelzon (1991b) show that belief revision operators satisfying (R1) through to (R6) can be characterised as determining $K * \phi$ by selecting those worlds in $[\phi]$ which are minimally implausible with respect to a ranking on worlds. To this end, they introduce binary relations \leq_K on worlds referred to as *faithful rankings* wherein $u \leq_K v$ means that v is at least as implausible as u from the perspective of an agent knowing only K .

Definition 2.3. *A **faithful ranking** for K is a binary relation \leq_K on possible worlds which satisfies the following properties:*

¹In the sense that $\phi_Y \equiv \text{forget}(\phi, V - Y)$.

1. $w \leq_K w'$ and $w' \leq_K w''$ implies $w \leq_K w''$.
2. Either $w \leq_K w'$ or $w' \leq_K w$.
3. $w \leq_K w'$ for all w' if and only if $w \models K$.

A family of faithful rankings $\{\leq_K\}_K$ such that $K_1 \equiv K_2$ implies $\leq_{K_1} = \leq_{K_2}$ is called a **faithful assignment**.

If W is a set of possible worlds and \leq is a faithful ranking, we write $\min(W, \leq)$ for the set of worlds in W which are minimal under \leq . That is to say, $x \in \min(W, \leq)$ if and only if $x \in W$ and $x \leq y$ for all $y \in W$.

Theorem 2.2 (Katsuno and Mendelzon, 1991b). *A binary function $* : L \times L \rightarrow L$ is a belief revision operator if and only if there exists a faithful assignment $\{\leq_K\}_K$ where for every K it follows that $[K * \phi] = \min([\phi], \leq_K)$.*

2.4 Relevance in Belief Revision

Although the general consensus is that a belief revision operator must satisfy the KM postulates, these postulates place few constraints on the behaviour of belief revision operators. For instance, they fail to rule out the belief revision operator defined by setting $K * \phi = K \wedge \phi$ if $K \wedge \phi$ is consistent and $K * \phi = \phi$ otherwise². This is in tension with the objective of belief revision to preserve as many of the original beliefs as possible.

In (Parikh 1999) the issue of minimal change is addressed via considering an additional postulate asserting that whenever the knowledge base is divisible into two unrelated components, then revision by a formula pertaining to only one of those components should leave the other component unchanged. For a KM belief revision operator $*$, Parikh's postulate can be expressed as follows:

P If $K \equiv K_1 \wedge K_2$ where $V(K_1) \subseteq X_1$, $V(K_2) \subseteq X_2$, $X_1 \cap X_2 = \emptyset$, and ϕ is such that $V(\phi) \subseteq X_1$ then

$$K * \phi \equiv (K_1 \circledast \phi) \wedge K_2$$

where \circledast is a belief revision operator for the language X_1 .

The statement of Parikh's postulate admits a weak reading wherein \circledast varies as a function of K , as well as a strong reading wherein \circledast is fixed. In order to clarify this situation, Peppas et al. (2004) introduced the postulate (P1) corresponding to the weak reading of (P), and the postulate (P2) which in combination with (P1) corresponds to the strong reading of (P). We state these as follows in the KM setting:

P1. If $V(K_1) \cap V(K_2) = \emptyset$ and $V(\phi) \subseteq V(K_1)$ then

$$((K_1 \wedge K_2) * \phi)_{V(K_2)} \equiv K_2.$$

P2. If $V(K_1) \cap V(K_2) = \emptyset$ and $V(\phi) \subseteq V(K_1)$ then

$$((K_1 \wedge K_2) * \phi)_{V(K_1)} \equiv (K_1 * \phi)_{V(K_1)}.$$

If we interpret a syntax-splitting $K \equiv K_1 \wedge K_2$ with $V(K_1) \cap V(K_2) = \emptyset$ as meaning that the beliefs represented by K_1 and K_2 are irrelevant, then we obtain the following intuitive readings for (P1) and (P2). The postulate (P1) requires that when revising K by a formula ϕ whose vocabulary is disjoint from K_2 , then only the part of K relevant to

²Consider the rankings \leq_K where $u \leq_K v$ for all $u, v \notin [K]$.

ϕ may be modified during revision, and hence K_2 is left unchanged. The postulate (P2) further requires that K_2 cannot influence the changes made to K_1 in any way.

Parikh (1999) only shows Parikh's postulate is consistent with the basic postulates for belief revision. Using these clarified postulates, Peppas et al. (2015) develop a characterisation of those belief operators satisfying (P1) and (P2), and show that Dalal's belief revision operator satisfies both the basic and supplementary KM postulates as well as (P1) and (P2). Subsequent work has extended these results to epistemic states (Kern-Isberner and Brewka 2017), to belief contraction operators (Haldimann, Kern-Isberner, and Beierle 2020), to epistemic entrenchments and selection functions (Aravanis, Peppas, and Williams 2019), and to preferential entailment relations (Kern-Isberner, Beierle, and Brewka 2020).

Rather than considering belief revision operators that satisfy (P1), Delgrande and Peppas (2018) consider belief revision operators which satisfy an analogue of Parikh's postulate for only certain theories and a subset of possible syntax splittings. The idea is that the knowledge engineer will specify a number of *irrelevance assertions*, which are expressions of the form $\sigma \rightarrow Y^3$ which intuitively state that whenever a knowledge base entails the formula σ then beliefs over Y must be treated as irrelevant to beliefs over \bar{Y} , and belief revision operators will be required to comply with these assertions in the following sense:

Definition 2.4. *A belief revision operator $*$ complies with $\sigma \rightarrow Y$ at K when either $K \not\models \sigma$ or for every consistent ϕ with $V(\phi) \subseteq Y$ the following postulate is satisfied:*

R If $K \vdash \neg\phi$ then $K * \phi \equiv (K * \phi)_Y \wedge K_{\bar{Y}}$.

For a belief revision operator $*$ induced from a faithful assignment $\{\leq_K\}_K$, Delgrande and Peppas (2018) show that complying with $\sigma \rightarrow Y$ is equivalent to stating that, for every K entailing σ , the following postulates are satisfied:

- S1.** If $u_Y = v_Y$, $K \vdash \neg u_Y$, and $K_{\bar{Y}} \not\models \neg u$ then $u \leq_K v$;
- S2.** If $u_Y = v_Y$, $K \vdash \neg u_Y$, $K_{\bar{Y}} \not\models \neg u$, and $K_{\bar{Y}} \vdash \neg v$ then $u <_K v$;

2.5 Conditional Independence

Parikh's postulate, and the majority of approaches descending from it, suffers from the limitation that the knowledge base must be able to be split into disjoint components in order for the postulate to apply. This limitation is already noted in (Chopra and Parikh 2000) which attempts to overcome this limitation by introducing the notion of a *belief structure*, which splits a knowledge base into a number of compartments which may overlap in vocabulary. However this compartmentalisation is fixed which can lead to information being lost.

This situation has an analogue in probability theory, where unconditional independence is a powerful but rarely applicable assumption. Rather, it is conditional independence which arises most frequently, and in fact has become

³For the reader familiar with multivalued dependencies, the similarity of this notation was a deliberate choice in (Delgrande and Peppas 2018).

a central component of modern probabilistic modelling and inference.

Inspired by probability theory, Darwiche (1997) introduces a notion of conditional logical independence together with a number of equivalent characterisations tailored for different reasoning problems. We will adopt the following definition, adapted from (Lang and Marquis 1998) and (Lang, Liberatore, and Marquis 2002).

Definition 2.5. *If X , Y_1 , and Y_2 are pairwise disjoint subsets of V and K is a propositional formula over V then Y_1 and Y_2 are **conditionally independent given X modulo K** when for any world u and formulae ϕ_1 and ϕ_2 with $V(\phi_1) \subseteq Y_1$ and $V(\phi_2) \subseteq Y_2$ such that $K \wedge u_X \vdash \phi_1 \vee \phi_2$ either $K \wedge u_X \vdash \phi_1$ or $K \wedge u_X \vdash \phi_2$.*

Example 2.2. *The sets $\{p\}$ and $\{r\}$ are conditionally independent given $\{q\}$ modulo $K := (p \rightarrow q) \wedge (q \rightarrow r)$. This follows from Lemma 3.1 below. To verify this for a specific case, let u be an arbitrary possible world and consider that $K \wedge u_{\{q\}} \vdash \neg p \vee r$. Either $u(q) = F$ in which case $K \wedge u_{\{q\}} \vdash \neg p$, or $u(q) = T$ in which case $K \wedge u_{\{q\}} \vdash r$, as required.*

Taking inspiration instead from database theory, we can regard the worlds satisfying a propositional formula K as constituting a database table wherein the attributes are the propositional variables in V . Then, we may consider the notion of a multivalued dependency:

Definition 2.6. *A propositional formula K satisfies the **multivalued dependency** $X \twoheadrightarrow Y$ if and only if*

$$K \equiv K_{XY} \wedge K_{\bar{Y}}.$$

Example 2.3. *The formula $K = (p \rightarrow q) \wedge (q \rightarrow r) \wedge (q \wedge r \rightarrow s)$ satisfies the multivalued dependencies $\{q\} \twoheadrightarrow \{p\}$ and $\{q\} \twoheadrightarrow \{r, s\}$.*

In the next section we show that multivalued dependencies, Darwiche’s conditional logical independence, and Parikh’s syntax-splittings are deeply interconnected.

3 Compliance with Multivalued Dependencies

Delgrande and Peppas (2018) point out that Parikh’s (1999) approach has a number of drawbacks: it assumes that irrelevance is completely determined by the current beliefs of an agent, it assumes that every syntax-splitting must be taken into account during belief revision, and it assumes that beliefs must be expressed over disjoint vocabularies in order to qualify as irrelevant. They argue that irrelevance is a domain-dependent phenomena, and represent this knowledge of the domain as a collection of irrelevance assertions which a belief revision operator is then required to comply with. This addresses the first and second drawback, but it leaves the issue of Parikh’s original postulate considering only unconditional independence, which is an unrealistically strong condition to expect to hold often.

Consider even a seemingly clear situation, such as a knowledge base containing knowledge about birds and knowledge about refrigerators. These topics would seem to

be independent. However, suppose we have that refrigerators require power, power is generated in the local area by wind turbines, and wind turbines often kill birds. Now, the ability to split the knowledge base is gone. However, we can observe that if the only link between birds and refrigerators passes through the language of wind turbines, then when revising knowledge about birds, our knowledge concerning refrigerators is not impacted, provided that our knowledge of wind turbines is unaffected.

In our approach, the knowledge engineer will represent their understanding of conditional irrelevance between components of the knowledge base as a collection of multivalued dependencies. The intuitive interpretation being that a multivalued dependency $X \twoheadrightarrow Y$ is satisfied when the only connections between knowledge over Y and knowledge over \bar{Y} pass through X . In other words, were an agent given complete information about X , its beliefs about Y and \bar{Y} would be independent. In our example scenario, knowledge about turbines comprises the only connection between birds and refrigerators, so the knowledge engineer would represent this via the multivalued dependencies $TurbineVocabulary \twoheadrightarrow BirdVocabulary$ and $TurbineVocabulary \twoheadrightarrow RefrigeratorVocabulary$.

When a knowledge base K satisfies $X \twoheadrightarrow Y$, it follows that $K \equiv K_{XY} \wedge K_{\bar{Y}}$. In the case $X = \emptyset$, the equivalence $K \equiv K_{XY} \wedge K_{\bar{Y}}$ amounts to a syntax-splitting as used by (Parikh 1999). Using Craig’s (1957) Interpolation Theorem, we can show the following equivalence between multivalued dependencies, Darwiche’s logical conditional independence, and a generalisation of Parikh’s syntax-splittings.

Lemma 3.1 (Splitting Criterion). *If Y_1 , Y_2 , and X are pairwise disjoint sets of propositional variables then for any propositional formulae K_1 and K_2 such that $V(K_1) \subseteq Y_1X$ and $V(K_2) \subseteq Y_2X$ it follows that Y_1 and Y_2 are independent given X modulo $K_1 \wedge K_2$.*

Proof Sketch. Using $K_1 \wedge K_2 \wedge u_X \vdash \phi_1 \vee \phi_2$ derive $K_1 \wedge u_X \wedge \neg\phi_1 \vdash \phi_2 \vee \neg K_2 \vee \neg u_X$. Apply Craig’s Interpolation Theorem to get an interpolant over X , and observe u_X must satisfy the interpolant or its negation. \square

The Splitting Criterion can be regarded as a special case of Darwiche’s results on *structured databases*, which are graphs similar to Bayesian networks whose vertices are labelled by components of a knowledge base in such a way that conditional independencies may be read directly off the graph itself (Darwiche 1997; Darwiche and Pearl 1994).

Lemma 3.2 (Projection Criterion). *Given a propositional formula K and disjoint sets Y_1 , Y_2 , and X of propositional variables, it follows that Y_1 and Y_2 are independent given X modulo K if and only if $K_{Y_1X} \wedge K_{Y_2X} \equiv K_{Y_1Y_2X}$ holds.*

Combining the Splitting and Projection Criteria, we arrive at the following result:

Theorem 3.1. *If X and Y are disjoint subsets of V and K is a propositional formula, then the following are equivalent:*

1. *There exists K_1 and K_2 with $V(K_1) \subseteq XY$ and $V(K_2) \subseteq \bar{Y}$ such that $K \equiv K_1 \wedge K_2$.*
2. *K satisfies $X \twoheadrightarrow Y$.*
3. *Y and $V - (XY)$ are independent given X modulo K .*

3.1 Partially Compliant Operators

Once a knowledge engineer has gathered a collection of multivalued dependencies to capture the conditional irrelevance properties of the domain, these are incorporated into the belief revision process by requiring that a belief revision operator *comply* with each of the multivalued dependencies. We introduce two notions of compliance, the first of which is *partial compliance*:

Definition 3.1. *If X and Y are disjoint subsets of V then a belief revision operator $*$ **partially complies with** $X \twoheadrightarrow Y$ if the following postulate holds:*

PCR. *If K is consistent and satisfies $X \twoheadrightarrow Y$, $V(\phi) \subseteq Y$, and ϕ is consistent then*

$$K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}.$$

In other words, any belief revision operator partially complying with $X \twoheadrightarrow Y$ must, when revising a knowledge base satisfying $X \twoheadrightarrow Y$ by a consistent formula over Y , leave the \overline{Y} component of the knowledge base unchanged. Returning to our example, supposing our knowledge base K satisfies $TurbineVocabulary \twoheadrightarrow BirdVocabulary$ and we revise by some formula ϕ in the bird vocabulary, we would have that knowledge over $BirdVocabulary$ is preserved. In particular, as $RefrigeratorVocabulary \subseteq BirdVocabulary$, our beliefs concerning the relationship between turbines and refrigerators could not be changed during revision by any formula ϕ only referring to birds.

As an immediate corollary of the Splitting Criterion, it follows that a belief revision operator partially complying with $X \twoheadrightarrow Y$ must preserve the satisfaction of $X \twoheadrightarrow Y$ when revising by a consistent formula over Y .

Theorem 3.2. *If $*$ is a belief revision operator which partially complies with $X \twoheadrightarrow Y$, K satisfies $X \twoheadrightarrow Y$, and $V(\phi) \subseteq Y$ then $K * \phi$ satisfies $X \twoheadrightarrow Y$.*

3.2 Fully Compliant Operators

Consider again an agent aware of wind turbines killing birds, and powering refrigerators, but with no knowledge directly linking birds and refrigerators. Suppose that this agent is given the new fact that birds have evolved a fear response to wind turbines, and are consequently no longer being killed by them. Revising by this new fact using a belief revision operator which partially complies with $TurbineVocabulary \twoheadrightarrow BirdVocabulary$ may result in changes to the agent's beliefs about refrigerators, as the postulate (PCR) does not constrain belief revision by formulae involving $TurbineVocabulary$. However, since the new knowledge is consistent with the fact that turbines power refrigerators, it seems that there is no reason why knowledge about refrigerators should be changed. This motivates us to consider a stronger notion of compliance which constrains revision by knowledge involving the shared vocabulary about wind turbines.

Definition 3.2. *If X and Y are disjoint subsets of V then a belief revision operator $*$ **fully complies with** $X \twoheadrightarrow Y$ if the following postulate holds:*

CR. *If K is consistent and satisfies $X \twoheadrightarrow Y$, $V(\phi) \subseteq XY$, and $\phi \wedge K_{\overline{Y}}$ is consistent then*

$$K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}.$$

In other words, any belief revision operator fully complying with $X \twoheadrightarrow Y$ must, when revising a knowledge base satisfying $X \twoheadrightarrow Y$ by a formula over XY which is consistent with the \overline{Y} component of the knowledge, leave the \overline{Y} component of the knowledge base unchanged. With our running example, the belief that birds evade wind turbine is expressed over $BirdVocabulary \cup TurbineVocabulary$, and consistent with the agent's beliefs over $BirdVocabulary$, and thus revising by this new belief using a belief revision operator which fully complies with $TurbineVocabulary \twoheadrightarrow BirdVocabulary$ will result in the beliefs about refrigerators being left unchanged.

It follows from the Splitting Criterion that a belief revision operator fully complying with $X \twoheadrightarrow Y$ must preserve the satisfaction of $X \twoheadrightarrow Y$ in appropriate circumstances:

Theorem 3.3. *If $*$ is a belief revision operator which fully complies with $X \twoheadrightarrow Y$, K satisfies $X \twoheadrightarrow Y$, $V(\phi) \subseteq XY$, and $\phi \wedge K_{\overline{Y}}$ is consistent then $K * \phi$ satisfies $X \twoheadrightarrow Y$.*

Requiring that a belief revision operator fully comply with $X \twoheadrightarrow Y$ is stronger than requiring that it partially comply with $X \twoheadrightarrow Y$, for the reason that (CR) applies to a broader class of formulae. Consequently, we obtain the following relationship between full and partial compliance:

Theorem 3.4. *If X and Y are disjoint subsets of V and $*$ is a belief revision operator which fully complies with $X \twoheadrightarrow Y$, then $*$ partially complies with $X \twoheadrightarrow Y$.*

3.3 Sources of Multivalued Dependencies

Our approach assumes that a knowledge engineer has specified a collection of multivalued dependencies as part of the domain knowledge, rather than these multivalued dependencies being automatically extracted from the initial knowledge base. This avoids requiring that a belief revision operator comply with spurious multivalued dependencies which just happen to hold. This also avoids the cost of having to determine all potential satisfied multivalued dependencies prior to revision, which is particularly important as checking whether a single multivalued dependency holds is known to be in Π_2^p (Lang, Liberatore, and Marquis 2002).

This raises the question of how a knowledge engineer should select an appropriate collection of multivalued dependencies. This question, in the analogous setting of irrelevance assertions, is discussed in (Delgrande and Peppas 2018) which suggests a number of possible sources: knowledge about the domain (e.g. birds and refrigerators are unrelated), a causal theory, a Bayesian network, or some structural features of a knowledge base which the knowledge engineer deems essential.

In our setting, we can make this a bit more precise. Using the notion of a symbolic causal network introduced by Darwiche and Pearl (1994), it follows from (Darwiche 1997) that conditional independence properties can be read off directly from these networks just as they are for Bayesian networks in probability theory (Pearl 2014). Any multivalued

dependency obtained by this method will be non-spurious since it would arise from the causal structure of the domain, as given in the causal network. We believe further investigation of revision operators which comply with the entire structure of a symbolic causal network is worthwhile.

4 Representation via Faithful Rankings

Belief revision operators which partially, or fully, comply with a multivalued dependency can be characterised semantically in terms of conditions on their corresponding faithful rankings. Using these characterisations, we can construct compliant belief revision operators, and gain insight into the epistemic aspect of compliance.

4.1 Partially Compliant Revision Operators

Belief revision operators which partially comply which a multivalued dependency $X \twoheadrightarrow Y$ can be represented via faithful assignments which *partially respect* $X \twoheadrightarrow Y$ in the following sense:

Definition 4.1. A faithful assignment $\{\leq_K\}_K$ *partially respects* $X \twoheadrightarrow Y$ if for every K either K does not satisfy $X \twoheadrightarrow Y$ or \leq_K satisfies the following conditions:

PCS1. If $u_{XY} = v_{XY}$, $K \vdash \neg u_Y$, $u \in [K_{\bar{Y}}]$, and $v <_K u$ then there exists w such that $w_Y = u_Y$ and $w <_K v$.

PCS2. If $K_{\bar{Y}} \vdash \neg v$ then there exists a world $u \in [K_{\bar{Y}}]$ such that $u_Y = v_Y$ and $u <_K v$.

Condition (PCS1) states that when worlds u and v with $u_{XY} = v_{XY}$ are ruled out by K on the basis of u_Y , yet u is consistent with $K_{\bar{Y}}$, then either u is at least as plausible as v or there is some world w with $w_Y = u_Y$ strictly more plausible than both u and v .

Condition (PCS2) further states that a possible world v inconsistent with $K_{\bar{Y}}$ is always less plausible than some possible world u satisfying $K_{\bar{Y}}$, and furthermore such a u may be obtained from v by modifying only the variables in \bar{Y} . This has the important consequence that whenever v is minimal in $[v_Y]$ then $v \in [K_{\bar{Y}}]$.

Theorem 4.1. If $*$ is a belief revision operator which partially complies with $X \twoheadrightarrow Y$, then there exists a faithful assignment $\{\leq_K\}_K$ which partially respects $X \twoheadrightarrow Y$, such that $[K * \phi] = \min([\phi], \leq_K)$ for all K and ϕ .

Proof Sketch. Choose any faithful ranking via Theorem 2.2. Verify (PCS1) by considering $K * u_Y$, and (PCS2) by considering $K * v_Y$. \square

Theorem 4.2. If $\{\leq_K\}_K$ is a faithful assignment which partially respects $X \twoheadrightarrow Y$, then the binary function defined by $[K * \phi] = \min([\phi], \leq_K)$ is a belief revision operator which partially complies with $X \twoheadrightarrow Y$.

Proof Sketch. Show $(K * \phi)_{XY} \wedge K_{\bar{Y}} \vdash K * \phi$ using (PCS1), and $K * \phi \vdash (K * \phi)_{XY} \wedge K_{\bar{Y}}$ using (PCS2). \square

4.2 Fully Compliant Revision Operators

As with (PCR), the postulate (CR) can be characterised in terms of conditions (CS1), (CS2), and (CS3) on faithful rankings. The stronger nature of (CR) will result in (CS1) and (CS2) appearing much closer to the original conditions (S1) and (S2) introduced in (Delgrande and Peppas 2018).

Definition 4.2. A faithful assignment $\{\leq_K\}_K$ *fully respects* $X \twoheadrightarrow Y$ if for every K either K does not satisfy $X \twoheadrightarrow Y$ or \leq_K satisfies the following conditions:

CS1. If $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, and $K_{\bar{Y}} \not\vdash \neg u$ then $u \leq_K v$.

CS2. If $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, $K_{\bar{Y}} \not\vdash \neg u$, and $K_{\bar{Y}} \vdash \neg v$ then $u <_K v$.

CS3. If $K \vdash \neg u_{XY}$, $K \vdash \neg v_{XY}$, and $K_{\bar{Y}} \not\vdash \neg u_{XY}$ and $K_{\bar{Y}} \vdash \neg v_{XY}$ then there exists w with $w_{XY} = u_{XY}$ and $w <_K v$.

Conditions (CS1) and (CS2) together state that for possible worlds u such that $K \vdash \neg u_{XY}$ it follows that u is minimally implausible among those worlds in $[u_{XY}]$ if and only if $u \in [K_{\bar{Y}}]$. Condition (CS3) is rather involved, but under the assumption of (CS1) it can be shown to be equivalent to the following (CS3') which states minimally implausible worlds consistent with $K_{\bar{Y}}$ are always strictly preferred to worlds inconsistent with K_X :

Theorem 4.3. Assuming (CS1) holds, condition (CS3) is equivalent to the following (CS3') condition:

CS3' If $u \in [K_{\bar{Y}}]$ and $v \notin [K_X]$ then $u <_K v$.

Demonstrating that a belief revision operator fully complying with $X \twoheadrightarrow Y$ results in the conditions (CS1), (CS2), and (CS3) being satisfied for the corresponding faithful rankings proceeds along lines strongly reminiscent to Theorem 2 of (Delgrande and Peppas 2018).

Theorem 4.4. If $*$ is a belief revision operator which fully complies with $X \twoheadrightarrow Y$, then there exists a faithful assignment $\{\leq_K\}_K$ which fully respects $X \twoheadrightarrow Y$ such that $[K * \phi] = \min([\phi], \leq_K)$ for all K and ϕ .

Proof Sketch. Choose any faithful ranking via Theorem 2.2. Verify (CS1) and (CS2) by considering $K * u_{XY}$, and (CS3) by considering $K * (u_{XY} \vee v_{XY})$. \square

Theorem 4.5. If $\{\leq_K\}_K$ is a faithful assignment which fully respects $X \twoheadrightarrow Y$, then the binary function defined by $[K * \phi] = \min([\phi], \leq_K)$ is a belief revision operator which fully complies with $X \twoheadrightarrow Y$.

Proof Sketch. Show $(K * \phi)_{XY} \wedge K_{\bar{Y}} \vdash K * \phi$ using (CS1), and show $K * \phi \vdash (K * \phi)_{XY} \wedge K_{\bar{Y}}$ using (CS2) and (CS3). \square

4.3 Existence of Fully Compliant Operators

Using these representation results, we can demonstrate that partial and full compliance with $X \twoheadrightarrow Y$ is compatible with the postulates for belief revision. As full compliance entails partial compliance, it suffices to show that for any multivalued dependency $X \twoheadrightarrow Y$ there exists a belief revision operator which fully complies with $X \twoheadrightarrow Y$.

Theorem 4.6. *If X and Y are disjoint then there exists a belief revision operator $*$ which fully complies with $X \twoheadrightarrow Y$.*

Proof Sketch. Intuitively, given a multivalued dependency $X \twoheadrightarrow Y$ one may construct a faithful assignment $\{\leq_K\}_K$ such that each ranking \leq_K arranges the worlds into three levels: the lowest level consisting of worlds satisfying K , the second level consisting of worlds satisfying $K_{\overline{Y}}$ but not K , and the third level consisting of worlds not satisfying $K_{\overline{Y}}$. It follows that (CS1), (CS2), and (CS3) are satisfied by any such faithful ranking, and therefore the corresponding belief revision operator satisfies the AGM postulates and is fully compliant with $X \twoheadrightarrow Y$. \square

5 Discussion

5.1 Related Work

The approach of (Delgrande and Peppas 2018) is closest to our work, which raises the question of whether the independence assertions studied there are related to the conditional independence assertions considered here. Clearly our multivalued dependencies have no mechanism for encoding the selective behaviour of the condition σ in an assertion $\sigma \twoheadrightarrow Z$ unless σ is tautologous, in which case it becomes equivalent to the multivalued dependency $\emptyset \twoheadrightarrow Z$.

In the reverse direction, suppose a multivalued dependency $X \twoheadrightarrow Y$ were encoded via an independence assertion $\sigma \twoheadrightarrow Z$. There are two natural-appearing approaches to consider:

1. If $Z = Y$ then when revising K with $K \vdash \sigma$ by ϕ with $V(\phi) \subseteq Z = Y$ it would follow that $K * \phi \equiv (K * \phi)_Y \wedge K_{\overline{Y}}$. Hence, we would have $K * \phi$ satisfies $\emptyset \twoheadrightarrow Y$. This is far too strong, for this means that all beliefs relating X and Y have been lost in the revision process, whereas we know that (PCR) and (CR) would result in them having been preserved.
2. If $Z = XY$ then when revising K with $K \vdash \sigma$ by ϕ with $V(\phi) \subseteq Z = XY$ it would follow that $K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{XY}}$. Hence, we would have $K * \phi$ satisfies $\emptyset \twoheadrightarrow XY$. This is again far too strong, for this means that all beliefs relating X and Y have been lost in the revision process, whereas we know that (PCR) and (CR) would result in them having been preserved.

Neither of these are tenable, which suggests that conditional independence assertions cannot in general simulate the multivalued dependencies we consider in this work.

5.2 Future Work

There are a number of opportunities for future work deriving from the above. One immediate observation is that although we demonstrate the classes of operators partially complying, or fully complying, with an arbitrary multivalued dependency are non-empty, we have not demonstrated that any reasonable-looking, “natural” belief revision operator reside within these classes. Hence, the question remains of finding interesting belief revision operators which satisfy our postulates.

Another line of inquiry would be to ask how we can take advantage of partial or full compliance to reduce the

computational cost of belief revision. One possibility is to develop efficient representations for rankings analogous to Bayesian networks for probability distributions (Pearl 2014) or ranking functions (Spohn 2012), which use the ranking conditions (CS1), (CS2), and (CS3) to factor a ranking into smaller components.

When a knowledge engineer specifies that a belief revision operator must comply with a large number of multivalued dependencies, the question of determining which multivalued dependencies constrain the operator in computing $K * \phi$ becomes challenging: determining whether K satisfies $X \twoheadrightarrow Y$ is shown to be coNP-complete in (Lang and Marquis 1998). This suggests that it would be valuable to adopt a special representation such as the symbolic causal networks introduced in (Darwiche and Pearl 1994), the structured databases of (Darwiche 1997), or perhaps an analogue of the B-structures in (Chopra and Parikh 2000). Investigating alternative notions of independence, such as the path-relevance from (Makinson 2009), or the conditional independence for ranking functions in (Spohn 2012), also might reduce this cost.

Finally, it would be interesting to investigate whether these postulates can be extended to nonmonotonic logics in a manner analogous to the extension of Parikh’s syntax splitting paradigm in (Kern-Isberner, Beierle, and Brewka 2020).

6 Conclusion

The central challenge of belief revision is to efficiently and plausibly restore consistency to a knowledge base after incorporating a contradictory proposition, and in a manner which causes only minimal changes to existing beliefs. With the standard postulates for belief revision failing to rule out rather pathologically-destructive or bizarre operators, the problem of formalising this requirement of minimality remains an ongoing challenge. We believe that enforcing the requirement that irrelevant beliefs are unchanged is an important aspect of minimal change.

In this work we have extended the previous work on unconditional independence in belief revision to accommodate conditional independence in the form of multivalued dependencies. We have introduced two notions by which a belief revision operator may comply with a multivalued dependency, and characterised these postulates in terms of conditions on faithful rankings. Further, we have endorsed the perspective of (Delgrande and Peppas 2018) that conditional independencies should be provided by the knowledge engineer, rather than read off of the knowledge base. This both avoids enforcing spurious conditional independencies, and means that our operators are not required to carry out the expensive task of checking for conditional independence themselves.

Our hope is that these postulates will assist in identifying those belief revision operators which can be truly said to result in minimal changes to existing beliefs, and that these operators will admit computationally efficient implementations by merit of being able to limit the amount of work required to perform revisions.

Acknowledgements

We would like to thank our reviewers for their insightful comments, as well as the Natural Sciences and Engineering Research Council of Canada for financial support.

References

- Abiteboul, S.; Hull, R.; and Vianu, V. 1995. *Foundations of Databases*, volume 8. Addison-Wesley Reading.
- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 510–530.
- Aravanis, T.; Peppas, P.; and Williams, M.-A. 2019. Full Characterization of Parikh’s Relevance-Sensitive Axiom for Belief Revision. *Journal of Artificial Intelligence Research*, 66: 765–792.
- Chopra, S.; and Parikh, R. 2000. Relevance sensitive belief structures. *Annals of Mathematics and Artificial Intelligence*, 28(1): 259–285.
- Craig, W. 1957. Three uses of the Herbrand-Gentzen Theorem in Relating Model Theory and Proof Theory. *The Journal of Symbolic Logic*, 22(3): 269–285.
- Darwiche, A. 1997. A Logical Notion of Conditional Independence: Properties and Applications. *Artificial Intelligence*, 97(1-2): 45–82.
- Darwiche, A.; and Pearl, J. 1994. Symbolic Causal Networks. In *AAAI*, 238–244.
- Delgrande, J.; and Peppas, P. 2018. Incorporating Relevance in Epistemic States in Belief Revision. In *International Conference on Principles of Knowledge Representation and Reasoning*.
- Delgrande, J. P. 2017. A Knowledge Level Account of Forgetting. *Journal of Artificial Intelligence Research*, 60: 1165–1213.
- Gärdenfors, P. 1990. Belief Revision and Relevance. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 349–365. Philosophy of Science Association.
- Haldimann, J. P.; Kern-Isberner, G.; and Beierle, C. 2020. Syntax Splitting for Iterated Contractions. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 17, 465–475.
- Hodges, W. 1993. *Model Theory*. Cambridge University Press.
- Katsuno, H.; and Mendelzon, A. O. 1991a. On the Difference between Updating a Knowledge Base and Revising It. In Allen, J. F.; Fikes, R.; and Sandewall, E., eds., *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR’91)*. Cambridge, MA, USA, April 22-25, 1991, 387–394. Morgan Kaufmann.
- Katsuno, H.; and Mendelzon, A. O. 1991b. Propositional Knowledge Base Revision and Minimal Change. *Artificial Intelligence*, 52: 263–294.
- Kern-Isberner, G.; Beierle, C.; and Brewka, G. 2020. Syntax Splitting= Relevance+ Independence: New Postulates for Nonmonotonic Reasoning from Conditional Belief Bases. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, volume 17, 560–571.
- Kern-Isberner, G.; and Brewka, G. 2017. Strong Syntax Splitting for Iterated Belief Revision. In *IJCAI*, 1131–1137.
- Lang, J.; Liberatore, P.; and Marquis, P. 2002. Conditional Independence in Propositional Logic. *Artificial Intelligence*, 141(1-2): 79–121.
- Lang, J.; and Marquis, P. 1998. Complexity Results for Independence and Definability. In *Proc. the 6th International Conference on Knowledge Representation and Reasoning*, 356–367.
- Makinson, D. 2009. Propositional relevance through letter-sharing. *J. Appl. Log.*, 7(4): 377–387.
- Parikh, R. 1999. Beliefs, Belief Revision, and Splitting Languages. *Logic, Language and Computation*, 2(96): 266–268.
- Pearl, J. 2014. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Elsevier.
- Peppas, P. 2008. Belief revision. *Foundations of Artificial Intelligence*, 3: 317–359.
- Peppas, P.; Chopra, S.; and Foo, N. Y. 2004. Distance Semantics for Relevance-Sensitive Belief Revision. In Dubois, D.; Welty, C. A.; and Williams, M., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004)*, Whistler, Canada, June 2-5, 2004, 319–328. AAAI Press.
- Peppas, P.; Williams, M.-A.; Chopra, S.; and Foo, N. 2015. Relevance in Belief Revision. *Artificial Intelligence*, 229: 126–138.
- Spohn, W. 2012. *The Laws of Belief: Ranking Theory and its Philosophical Applications*. Oxford University Press.