Using Conditional Independence for Belief Revision

Matthew James Lynn[1], **James P. Delgrande**[1], **Pavlos Peppas**[2]

[1] Simon Fraser University, Canada
[2] University of Patras, Greece

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

Abstract

We present an approach for incorporating qualitative conditional independence into belief revision. Our stance is that, as with probability, conditional independence arises far more frequently than the unconditional independence studied in previous work. Our approach uses multivalued dependencies to represent domain-dependent conditional independence assertions. In particular, the multivalued dependency $X \twoheadrightarrow Y$ expresses that assertions over the subvocabularies Y and $\overline{Y}$ are independent whenever complete information is known about the subvocabulary X. We introduce the class of partially compliant revision operators, wherein revising a KB satisfying $X \twoheadrightarrow Y$ by a formula expressed over Y results in the part of the KB expressed over $\overline{Y}$ remaining unchanged. This helps ensure that partially compliant revision operators result in minimal changes to existing beliefs, as irrelevant existing beliefs are left unchanged. Furthermore, we identify a subclass of partially compliant operators, called fully compliant operators, for which the same is true when revising by a formula expressed over XY rather than just Y. For both classes, we provide representation results which characterise compliance semantically in terms of faithful rankings. Finally, we compare our use of multivalued dependencies to existing work on independence 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 general 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 introduces the class of belief revision operators which *partially comply* with a multivalued dependency, and characterises partial compliance in terms of faithful rankings. Section 4 studies the stronger property of *full compliance* with a multivalued dependency, again with a characterisation in terms of faithful rankings. In Section 5 we examine and clarify the relationship between logical conditional independence, multivalued dependencies, and syntax splitting. Finally, Section 6 discusses our approach, 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 $\overline{X}$. 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 function $u_X : X \rightarrow \{T, F\}$ agreeing with u.

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 operator[1] (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. (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 1991).

Definition 2.2. *A binary function $* : L \times L \rightarrow L$ is a **belief revision operator** if it satisfies the following **basic 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 $K_1 * \phi_1 \equiv K_2 * \phi_2$.*

*We will say that a belief revision operator $*$ satisfies the **supplementary postulates** when it satisfies the following:*

R5. $(K * \phi) \wedge \psi \vdash K * (\phi \wedge \psi)$;

R6. *If $(K * \phi) \wedge \psi$ is satisfiable then $K * (\phi \wedge \psi) \vdash (K * \phi) \wedge \psi$.*

Unless we explicitly specify that a belief revision operator satisfies the supplementary postulates, we will assume only that the basic postulates are satisfied. Note that this partitioning of the Katzuno-Mendelzon postulates into basic and supplementary postulates exactly mirrors the organisation of the original AGM postulates into basic and supplementary postulates.

When working with belief revision operators satisfying the basic and supplementary KM postulates, Katsuno and Mendelzon (1991) show that we may semantically characterise the belief revision operator 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

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

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:*

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 $w \models K$.*

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 1991))**.** *A binary function $* : L \times L \to L$ is a belief revision operator satisfying the supplementary postulates if and only if for every K there exists a faithful ranking $\leq_K$ for K such 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 $K * \phi = \phi$ otherwise[2]. This is in tension with the objective of belief revision to preserve as many of the original beliefs as possible.

In (Parikh 1999) the notion 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.(2015) introduced the following variations (P1) and (P2) of (P) which we state here 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 $((K_1 \wedge K_2) * \phi)_{V(K_1)} \equiv (K_1 * \phi)_{V(K_1)}$.

Intuitively, (P1) states that when revising K by ϕ, only the part of K relevant to ϕ is revised. The role of (P2) is to ensure that whenever K_1 and K_2 agree on the beliefs relevant to ϕ, then the revisions $K_i * \phi$ change this part in the same way.

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 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 Pappas (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* $\sigma \twoheadrightarrow Y$[3], 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 $\sigma \twoheadrightarrow Y$ at K when either $K \nvdash \sigma$ or for every consistent ϕ with $V(\phi) \subseteq Y$ the following postulate is satisfied:*

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

For a belief revision operator $*$ induced from a family of faithful rankings $\{\leq_K\}_K$, Delgrande and Pappas (2018) show that complying with $\sigma \twoheadrightarrow Y$ 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 $K_{\overline{Y}} \nvdash \neg u$ then $u \leq_K v$;

S2. If $u_Y = v_Y$, $K \vdash \neg u_Y$, $K_{\overline{Y}} \nvdash \neg u$, and $K_{\overline{Y}} \vdash \neg v$ 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 a central component of modern probabilistic modelling and inference.

Inspired by probability theory, Darwiche (1997) introduces a notion of conditional logical independence together

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

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

with a number of equivalent characterisations tailored for different reasoning problems. We will adopt the following notion, 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 $K := (p \to q) \wedge (q \to r)$. This follows from Theorem 5.2 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*** *$X \twoheadrightarrow Y$ when for any models v and u of K such that $v_X = u_X$ there exists a model w of K such that $w_Y = v_Y$ and $w_{\overline{Y}} = u_{\overline{Y}}$.*

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

In Section 5 we show that multivalued dependencies are equivalent to a restricted case of conditional independence, and that both are equivalent to a generalisation of Parikh's syntax-splittings.

3 Compliance with Multivalued Dependencies

Parikh's original postulate considers only unconditional independence. However unconditional independence is a strong condition which is unrealistic 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 independencies between components of the knowledge base as a collection of multivalued dependencies. The intuitive interpretation being that a multivalued dependency $X \twoheadrightarrow Y$ captures that the only connections between knowledge over Y and knowledge over $\overline{Y}$ arise from knowledge over X. 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 $TurbineVocabulary \twoheadrightarrow RefrigeratorVocabulary$.

Once the knowledge engineer has selected a collection of multivalued dependencies which capture the conditional independence relations between different areas of knowledge being worked with, these multivalued dependencies are incorporated into the belief revision process by requiring *compliance* in the following sense:

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*

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

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, preserve 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 $\overline{BirdVocabulary}$ is preserved. In particular, our beliefs concerning the relationship between turbines and refrigerators could not be changed by any formula ϕ only referring to birds.

We refer to this as only partial compliance, for in the next section we will introduce a postulate which applies to suitable ϕ with $V(\phi) \subseteq XY$ rather than just for ϕ with $V(\phi) \subseteq Y$.

3.1 Representation via Faithful Rankings

Those belief revision operators which partially comply with a multivalued dependency can be characterised semantically by conditions on their corresponding faithful rankings.

Definition 3.2. *If $\leq_K$ is a faithful ranking for K then $\leq_K$* ***partially respects*** *$X \twoheadrightarrow Y$ if either K does not satisfy $X \twoheadrightarrow Y$ or the following conditions are satisfied:*

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

PCS2. *If $K_{\overline{Y}} \vdash \neg v$ then there exists a world $u \in [K_{\overline{Y}}]$ 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 $K_{\overline{Y}}$, 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 $K_{\overline{Y}}$ is always less plausible than some possible world u satisfying $K_{\overline{Y}}$, and furthermore such a u may be obtained from v by modifying only the variables in $\overline{Y}$.

Theorem 3.1. *If $*$ is a belief revision operator satisfying the supplementary postulates which partially complies with $X \twoheadrightarrow Y$, then there exist faithful rankings $\{\leq_K\}_K$ which partially respect $X \twoheadrightarrow Y$ such that $[K * \phi] = \min([\phi], \leq_K)$ for all K and ϕ.*

Proof. By Theorem 2.2 there exist faithful rankings $\{\leq_K\}_K$ such that $[K * \phi] = \min([\phi], \leq_K)$ for all K and ϕ. Suppose $*$ partially complies with $X \twoheadrightarrow Y$ and consider a consistent formula K. In the case K does not satisfy $X \twoheadrightarrow Y$ then $\leq_K$ partially respects $X \twoheadrightarrow Y$ in the trivial sense. Otherwise, K satisfies $X \twoheadrightarrow Y$ and we must demonstrate that $\leq_K$ satisfies (PCS1) and (PCS2).

Part 1. Suppose that (PCR) holds. In order to verify (PCS1), suppose that u and v are worlds such that $u_{XY} = v_{XY}$, $K \vdash \neg u_Y$, and $u \in [K_{\overline{Y}}]$, and $v <_K u$. Applying (PCR) it follows that

$$[(K * u_Y)_{XY}] \cap [K_{\overline{Y}}] = [K * u_Y].$$

Assume for the sake of contradiction that $u \in [K * u_Y]$. As $u \in [u_Y]$ and $u_{XY} = v_{XY}$ it follows that $v \in [u_Y]$, which means that $u \leq_K v$. However, this contradicts our assumption that $v <_K u$, so it must be the case that $u \notin [K * u_Y]$. Therefore, as $u \in [K_{\overline{Y}}]$, it follows that $u \notin [(K * u_Y)_{XY}]$, and thus $v \notin [(K * u_Y)_{XY}]$ as $u_{XY} = v_{XY}$. Theorem 2.1 implies that $K * u_Y \vdash (K * u_Y)_{XY}$, hence $v \notin [K * u_Y]$. However, as $[u_Y] \neq \emptyset$ there must exist some world $w \in [K * u_Y]$. It follows that $w_Y = u_Y$, and furthermore as $v \in [u_Y]$ yet $v \notin [K * u_Y]$ it follows that $w <_K v$ as required. Therefore, (PCS1) is satisfied.

Part 2. In order to verify (PCS2) suppose that v is a world such that $K_{\overline{Y}} \vdash \neg v$. Applying (PCR) it follows that

$$[K * v_Y] = [(K * v_Y)_{XY}] \cap [K_{\overline{Y}}].$$

By our supposition that $K_{\overline{Y}} \vdash \neg v$ it follows that $v \notin [K_{\overline{Y}}]$, and therefore $v \notin [K * v_Y]$. However, as $[v_Y] \neq \emptyset$ it follows that $[K * v_Y] \neq \emptyset$. Let $u \in [K * v_Y]$ be arbitrary, and observe that $u \in [v_Y]$ meaning $u_Y = v_Y$. As $v \in [v_Y]$ but $v \notin [K * v_Y]$ it follows then that $u <_K v$ as required. Therefore, (PCS2) holds. $\square$

Theorem 3.2. *If $\{\leq_K\}_K$ are faithful rankings which partially respect $X \twoheadrightarrow Y$, then the binary function defined by $[K * \phi] = \min([\phi], \leq_K)$ is a belief revision operator satisfying the supplementary postulates which partially complies with $X \twoheadrightarrow Y$.*

Proof. By Theorem 2.2 it follows that $*$ is a belief revision operator. Suppose K is a consistent formula such that $\leq_K$ partially respects $X \twoheadrightarrow Y$. In the case K does not satisfy $X \twoheadrightarrow Y$ there is nothing to check, so assume K satisfies $X \twoheadrightarrow Y$. This means that $\leq_K$ satisfies (PCS1) and (PCS2). Using this, we must demonstrate that $[K * \phi] = [(K*\phi)_{XY}] \cap [K_{\overline{Y}}]$ whenever $V(\phi) \subseteq Y$ and ϕ is consistent. In the case $K \wedge \phi$ is consistent then $K * \phi \equiv K \wedge \phi$, and K satisfying $X \twoheadrightarrow Y$ means $K \equiv K_{XY} \wedge K_{\overline{Y}}$ (cf. Theorem 5.3 below), hence $K * \phi \equiv K_{XY} \wedge K_{\overline{Y}} \wedge \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$. Therefore, we will assume $K \vdash \neg\phi$, in which case our proof has two parts:

Part 1. In order to show $[(K * \phi)_{XY}] \cap [K_{\overline{Y}}] \subseteq [K * \phi]$ suppose that $u \in [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$. Being that $u \in [(K * \phi)_{XY}]$ it follows that there exists $v \in [K * \phi]$ such that $u_{XY} = v_{XY}$. By our assumption that $K \vdash \neg\phi$ and the observation that $v_Y \vdash \phi$, it follows that $K \vdash \neg v_Y$. Being that $u_{XY} = v_{XY}$ it follows that $K \vdash \neg u_Y$. Assume for the sake of contradiction that $v <_K u$. It then follows from (PCS1) that there exists w with $w_Y = u_Y = v_Y$ and $w <_K v$. However, $v \in [\phi]$ and $w_Y = v_Y$ implies $w \in [\phi]$, and therefore $v \in [K * \phi]$ implies $v \leq_K w$ which contradicts our assumption that $w <_K v$. Therefore, our assumption was wrong, so it must be the case that $u \leq_K v$. This means that $u \in [\phi]$ and $v \in \min([\phi], \leq_K)$ which implies $u \in \min([\phi], \leq_K) = [K * \phi]$. Thus, as u was arbitrary, it follows that $[(K * \phi)_{XY}] \cap [K_{\overline{Y}}] \subseteq [K * \phi]$.

Part 2. In order to show $[K * \phi] \subseteq [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$ start by observing that $K * \phi \vdash (K * \phi)_{XY}$ by Theorem 2.1. Therefore, it suffices to verify that $[K * \phi] \subseteq [K_{\overline{Y}}]$. Suppose that $v \in [K * \phi]$ but assume for the sake of contradiction that $v \notin [K_{\overline{Y}}]$. It follows that $K_{\overline{Y}} \vdash \neg v$, and therefore by (PCS2) there exists a world $u \in [K_{\overline{Y}}]$ such that $u_Y = v_Y$ and $u <_K v$. Observing that $v \in [\phi]$, $V(\phi) \subseteq Y$, and $u_Y = v_Y$ it follows that $u \in [\phi]$. However, by our assumption that $v \in [K * \phi]$ this implies $v \leq_K u$ which is a contradiction as $u <_K v$. Therefore, it must be that $v \in [K_{\overline{Y}}]$. As v was arbitrary, it follows that $[K * \phi] \subseteq [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$ as required.

It follows that $K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$, showing that (PCR) holds. $\square$

3.2 Existence of Partially Compliant Operators

Parikh (1999) demonstrates the existence of a belief revision operator satisfying postulate **P** as follows: Given a knowledge base K and a formula ϕ to revise by which is inconsistent with K, first K is split as $K_Y \wedge K_{\overline{Y}}$ where Y is the smallest subset of V with $V(\phi) \subseteq Y$ and K satisfies $\emptyset \twoheadrightarrow Y$. K is then replaced by $\phi \wedge K_{\overline{Y}}$. In order to mirror this, we need to show that we can construct such an analogous Y, which we refer to here as a *section*:

Definition 3.3. *An X-**section** of ϕ is a subset $Y \subseteq V(\phi)$ disjoint from X such that ϕ satisfies $X \twoheadrightarrow Y$.*

In order to construct a smallest section, we will make use of the following properties of multivalued dependencies:

Lemma 3.1 (Abiteboul, Hull, and Vianu (1995))**.** *1. If $X \twoheadrightarrow Y$ then $X \twoheadrightarrow V - Y$;*

2. If $Y \subseteq X$ then $X \twoheadrightarrow Y$;

3. If $X \twoheadrightarrow Y$ and $Y \twoheadrightarrow X$ then $X \twoheadrightarrow Z$;

4. If $X \twoheadrightarrow Y$ then $XZ \twoheadrightarrow YZ$;

For any set of variables X we can consider the set $d_K(X)$ of Y such that K satisfies $X \twoheadrightarrow Y$, that is

$$d_K(X) := \{Y \mid K \text{ satisfies } X \twoheadrightarrow Y\}.$$

An important consequence of Lemma 3.1 is that $d_K(X)$ forms a Boolean algebra:

Corollary 3.1 (Abiteboul, Hull, and Vianu (1995))**.** *For any K and $X \subseteq V$ it follows that $d_K(X)$ is a Boolean algebra, i.e. $d_K(X)$ is closed under unions, intersections, complementation, and it contains X.*

Theorem 3.3 (Conditional Sectioning Theorem)**.** *If there is an X-section of K containing $V(\phi)$ then there is a unique smallest X-section of K containing $V(\psi)$.*

Proof. Simply take the intersection of all $Y \in d_K(X)$ such that $V(\phi) \subseteq Y$. $\square$

Theorem 3.4. *For every $X \subseteq V$ there exists a belief revision operator $*$ which satisfies the basic postulates and partially complies with every $X \twoheadrightarrow Y$ where $Y \subseteq V$ is disjoint from X.*

Proof. Construct a belief revision operator $*$ as follows. For every K and ϕ define $K * \phi$ as $K \wedge \phi$ in the case $K \wedge \phi$ is consistent. Otherwise, if there is an X-section of K containing $V(\phi)$ choose the smallest X-section Y of K containing $V(\phi)$ and define $K*\phi$ as $(K)_{\overline{Y}} \wedge \phi$. Otherwise, define $K*\phi$ as ϕ. $\square$

4 Full Compliance with Multivalued Dependencies

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 a new fact that modern wind turbines stop momentarily when an approaching bird is detected, in order to allow its safe passage, and consider how the agent may revise its knowledge base. A revision operator that partially complies with $TurbineVocabulary \twoheadrightarrow BirdVocabulary$ is not useful here, since we are revising by a formula in the language of both turbines and birds. 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. Thus, we can consider a stronger notion of compliance wherein we can revise by knowledge containing the shared variables about turbines.

Definition 4.1. *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*

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

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 4.1. *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 $X \twoheadrightarrow Y$.*

Proof. Suppose K is a consistent formula satisfying $X \twoheadrightarrow Y$, and ϕ is a consistent formula with $V(\phi) \subseteq Y$. As K and ϕ are consistent and $V(K_{\overline{Y}}) \cap V(\phi) = \emptyset$ it follows that $K_{\overline{Y}} \wedge \phi$ is consistent, and hence we may apply (CR) to write $K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$. Which is exactly what was required to show (PCR) is satisfied. Hence, $*$ partially complies with $X \twoheadrightarrow Y$. $\square$

4.1 Representation via Faithful Rankings

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. *If $\leq_K$ is a faithful ranking for K then $\leq_K$* ***fully respects*** *$X \twoheadrightarrow Y$ if either K does not satisfy $X \twoheadrightarrow Y$ or the following conditions are satisfied:*

CS1. *If $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, and $K_{\overline{Y}} \nvdash \neg u$ then $u \leq_K v$.*

CS2. *If $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, $K_{\overline{Y}} \nvdash \neg u$, and $K_{\overline{Y}} \vdash \neg v$ then $u <_K v$.*

CS3. *If $K \vdash \neg u_{XY}$, $K \vdash \neg v_{XY}$, and $K_{\overline{Y}} \nvdash \neg u_{XY}$ and $K_{\overline{Y}} \vdash \neg v_{XY}$ then there exists w with $w_{XY} = u_{XY}$ and $w <_K v$.*

Condition (CS1) states that whenever worlds u and v incompatible with K are such that $u_{XY} = v_{XY}$, and u is consistent with $K_{\overline{Y}}$, then v cannot be more plausible than u. In the case v is itself inconsistent with $K_{\overline{Y}}$, then (CS2) strengthens this to say that u is strictly more plausible than v. Finally, (CS3) ensures that whenever u is compatible with $K_{\overline{Y}}$ and v is not, then u can be modified to be strictly more plausible than v by modifying variables not in XY.

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 4.2. *If $*$ is a belief revision operator satisfying the supplementary postulates which fully complies with $X \twoheadrightarrow Y$, then there exist faithful rankings $\{\leq_K\}_K$ which fully respects $X \twoheadrightarrow Y$ such that $[K * \phi] = \min([\phi], \leq_K)$ for all K and ϕ.*

Proof. By Theorem 2.2 there exist faithful rankings $\{\leq_K\}_K$ such that $[K * \phi] = \min([\phi], \leq_K)$ for all K and ϕ. Suppose that $*$ fully complies with $X \twoheadrightarrow Y$, and consider K satisfying $X \twoheadrightarrow Y$. We must show that $\leq_K$ satisfies the conditions (CS1), (CS2), and (CS3).

Part 1. Suppose u and v are worlds such that $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, and $K_{\overline{Y}} \nvdash \neg u$. This last assumption implies that u_{XY} is consistent with $K_{\overline{Y}}$, hence we may apply the postulate (CR) to write

$$\begin{aligned}[K * u_{XY}] &= [(K * u_{XY})_{XY}] \cap [K_{\overline{Y}}] \\ &= [u_{XY}] \cap [K_{\overline{Y}}].\end{aligned}$$

As $K_{\overline{Y}} \nvdash \neg u$ it follows that $u \in [K_{\overline{Y}}]$, and tautologically $u \in [u_{XY}]$, so it follows that $u \in [K * u_{XY}]$. Hence, as $v \in [u_{XY}]$ it follows that $u \leq v$ verifying (CS1).

Part 2. In order to see (CS2) suppose further that $K_{\overline{Y}} \vdash \neg v$. In this case, $v \notin [K * u_{XY}]$ hence $u < v$ verifying (CS2).

Part 3. In order to verify (CS3) suppose u and v are worlds such that $K \vdash \neg u_{XY}$, $K \vdash \neg v_{XY}$, $K_{\overline{Y}} \nvdash u_{XY}$ and $K_{\overline{Y}} \vdash \neg v_{XY}$. Construct the formula $\phi = u_{XY} \vee v_{XY}$ and observe that v_{XY} is consistent with $K_{\overline{Y}}$ and hence ϕ

is consistent with $K_{\overline{Y}}$. However, ϕ is inconsistent with K by our hypothesis. Therefore, we may apply (CR) to write $K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$. By the success postulate, $K * \phi \vdash \phi \equiv u_{XY} \vee v_{XY}$. By (CR) we also know $K * \phi \vdash K_{\overline{Y}}$. However, we also know $K_{\overline{Y}} \vdash \neg v_{XY}$, and therefore it follows that $K * \phi \vdash u_{XY}$. Hence, choosing any $w \in [K * \phi]$ it follows that $w_{XY} = u_{XY}$ and $w \leq v$. Being that $v \notin [K * \phi]$ yet $v \in [\phi]$ it follows that $w < v$. $\square$

Theorem 4.3. *If $\{\leq_K\}_K$ are faithful rankings which fully respects $X \twoheadrightarrow Y$, then the binary function defined by $[K * \phi] = \min([\phi], \leq_K)$ is a belief revision operator satisfying the supplementary postulates which fully complies with $X \twoheadrightarrow Y$.*

Proof. By Theorem 2.2 it follows that $*$ is a belief revision operator. Suppose $\leq_K$ fully respects $X \twoheadrightarrow Y$. In the case K does not satisfy $X \twoheadrightarrow Y$ then there is nothing to verify. Assume K satisfies $X \twoheadrightarrow Y$, so that $\leq_K$ satisfies (CS1), (CS2), and (CS3). We must demonstrate that $[K * \phi] = [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$ whenever $V(\phi) \subseteq XY$ and $\phi \wedge K_{\overline{Y}}$ is consistent. In the case $K \wedge \phi$ is consistent then $K * \phi \equiv K \wedge \phi$, and K satisfying $X \twoheadrightarrow Y$ means $K \equiv K_{XY} \wedge K_{\overline{Y}}$ (cf. Theorem 5.3 below), hence $K * \phi \equiv K_{XY} \wedge K_{\overline{Y}} \wedge \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$. Therefore, we will assume $K \vdash \neg\phi$, in which case our proof has two parts:

Part 1. In order to show $[(K * \phi)_{XY}] \cap [K_{\overline{Y}}] \subseteq [K * \phi]$ suppose that $u \in [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$. Being that $u \in [(K * \phi)_{XY}]$ it follows that there exists $v \in [K * \phi]$ such that $u_{XY} = v_{XY}$. Observe that $K \vdash \neg\phi$, and furthermore $\neg\phi \vdash \neg u_{XY}$ as $u_{XY} = v_{XY}$ and $v \in [\phi]$. Hence, $K \vdash \neg u_{XY}$. However, $u \in [K_{\overline{Y}}]$ so $K_{\overline{Y}} \nvdash \neg u$. Therefore, by (CS1) it follows that $u \leq_K v$. However, $u \in [\phi]$ and $v \in \min([\phi], \leq_K)$ so it follows that $u \in \min([\phi], \leq_K) = [K * \phi]$. With u being arbitrary, it follows that $[(K * \phi)_{XY}] \cap [K_{\overline{Y}}] \subseteq [K * \phi]$ as required.

Part 2. In order to show $[K * \phi] \subseteq [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$ consider a world $v \in [K * \phi] = \min([\phi], \leq_K])$, and observe that $v \in [(K * \phi)_{XY}]$ hence it suffices to show $v \in [K_{\overline{Y}}]$. Assume for the sake of contradiction that $v \notin [K_{\overline{Y}}]$, which is to say that $K_{\overline{Y}} \vdash \neg v$. We have two cases:

1. In the case there exists a world u with $u_{XY} = v_{XY}$, and $K_{\overline{Y}} \nvdash \neg u$, argue as follows. As $u_{XY} = v_{XY}$ and $v \in [\phi]$ it follows that $u \in [\phi]$, and hence $\neg\phi \vdash \neg u_{XY}$. Observing that $K \vdash \neg\phi$ it follows that $K_{\overline{Y}} \vdash \neg u_{XY}$. As $v \notin [K_{\overline{Y}}]$ by our assumption, it follows that $K_{\overline{Y}} \vdash \neg v$. Hence, by (CS2), it follows that $u <_K v$. However, $u \in [\phi]$ and $v \in \min([\phi], \leq_K)$ so this is a contradiction.
2. In the other case, $K_{\overline{Y}} \vdash \neg v_{XY}$. Recalling that $K_{\overline{Y}}$ is consistent with ϕ, it follows that there exists a world $u \in [K_{\overline{Y}} \wedge \phi]$, for which we know that $K \vdash \neg u_{XY}$ and $K_{\overline{Y}X} \nvdash \neg u_{XY}$. However, by (CS3) we may conclude that there exists w with $w_{XY} = u_{XY}$ such that $w < v$. As ϕ is expressed over the vocabulary XY and $u \in [\phi]$ it follows that $w \in [\phi]$ and $w <_K v$. However, this contradicts $v \in \min([\phi], \leq_K)$.

In both cases, a contradiction is achieved, so our assumption that $v \notin [K_{\overline{Y}}]$ must have been false. Hence, $v \in [K_{\overline{Y}}]$ as well. With v being arbitrary, we have shown $[K * \phi] \subseteq [(K * \phi)_{XY}] \cap [K_{\overline{Y}}]$. $\square$

4.2 Existence of Fully Compliant Operators

With this representation result in hand, the next question is whether there exists a belief revision operator which fully complies with an arbitrary multivalued dependency $X \twoheadrightarrow Y$ where X need not be empty. Fortunately, the answer is affirmative:

Theorem 4.4. *If X and Y are disjoint then there exists a belief revision operator $*$ satisfying the supplementary postulates which fully complies with $X \twoheadrightarrow Y$.*

Proof. It suffices to construct a family of faithful rankings $\{\leq_K\}_K$ where each $\leq_K$ fully respects $X \twoheadrightarrow Y$, in which case the corresponding belief revision operator $*$ with $[K * \phi] = \min([\phi], \leq_K)$ will fully comply with $X \twoheadrightarrow Y$. Given K define the function $\rho_K : \Omega \to \mathbb{N}$ given by

$$\rho_K(u) := \begin{cases} 0 & \text{if } u \in [K] \\ 1 & \text{if } u \notin [K] \text{ and } K_{\overline{Y}} \nvdash \neg u \\ 2 & \text{otherwise} \end{cases}$$

The ranking $\leq_K$ is defined by setting $u \leq_K v$ if and only if $\rho_K(u) \leq \rho_K(v)$. As $\rho_K(u) = 0$ if and only if $u \in [K]$, it follows that the minimal worlds under $\leq_K$ are exactly the worlds satisfying K. Hence, $\leq_K$ is a faithful ranking for K.

In order to argue $\leq_K$ fully respects $X \twoheadrightarrow Y$ assume that K satisfies $X \twoheadrightarrow Y$, and verify (CS1), (CS2), and (CS3) as follows:

1. (CS1) Suppose that $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, and $K_{\overline{Y}} \nvdash \neg u$. It follows that $r_K(u) = 1$ and $r_K(v) \geq 1$, hence $u \leq_K v$ as required.
2. (CS2) Suppose that $u_{XY} = v_{XY}$, $K \vdash \neg u_{XY}$, $K_{\overline{Y}} \nvdash \neg u$, and $K_{\overline{Y}} \vdash \neg v$. It follows that $r_K(u) = 1$ and $r_K(v) = 2$, hence $u <_K v$ as required.
3. (CS3) Suppose that $K \vdash \neg u_{XY}$, $K \vdash \neg v_{XY}$, $K_{\overline{Y}} \nvdash \neg u_{XY}$, and $K_{\overline{Y}} \vdash \neg v_{XY}$. As a consequence of $K_{\overline{Y}} \vdash \neg v_{XY}$ it follows that $\rho_K(v) = 2$. As $K_{\overline{Y}} \nvdash \neg u_{XY}$ there exists a world w with $w_{XY} = u_{XY}$ and $w \in [K_{\overline{Y}}]$, so that $\rho_K(w) = 1$ and hence $w <_K u$.

$\square$

This leaves the open question of whether any set of multivalued dependencies can be simultaneously fully complied with by some belief revision operator.

5 Syntax Splitting and MVDs

5.1 Syntax Splitting and Conditional Independence

In this section we demonstrate that Parikh's syntax splitting generalises naturally into the framework of multivalued dependencies and conditional independence. We start by showing that syntax splitting gives rise to conditional logical independence via leveraging Craig's Interpolation Theorem (Craig 1957), which is stated as follows:

Theorem 5.1 (Craig's Interpolation Theorem)**.** *If $K \vdash \psi$ then there exists ϕ with $V(\phi) \subseteq V(K) \cap V(\psi)$ such that $K \vdash \phi$ and $\phi \vdash \psi$.*

Theorem 5.2 (The 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. Suppose u is a world, and ϕ_1 and ϕ_2 are propositional formulae with $V(\phi_1) \subseteq Y_1$ and $V(\phi_2) \subseteq Y_2$ such that $K_1 \wedge K_2 \wedge u_X \models \phi_1 \vee \phi_2$. We must demonstrate that either $K_1 \wedge K_2 \wedge u_X \models \phi_1$ or $K_1 \wedge K_2 \wedge u_X \models \phi_2$ holds.

It follows from our hypotheses that $K_1 \wedge u_X \wedge \neg\phi_1 \models \phi_2 \vee \neg K_2 \vee \neg u_X$. Applying Craig's Interpolation Theorem there exists an interpolant δ such that $V(\delta) \subseteq V(K_1 \wedge u_X \wedge \neg\phi_1) \cap V(\phi_2 \vee \neg K_2 \vee \neg u_X)$ and furthermore $K_1 \wedge u_X \wedge \neg\phi_1 \models \delta$ and $\delta \models \phi_2 \vee \neg K_2 \vee \neg u_X$.

Observing that $V(K_1 \wedge u_X \wedge \neg\phi_1) \cap V(\phi_2 \vee \neg K_2 \vee \neg u_X) \subseteq (Y_1X) \cap (Y_2X) = X$ it follows that $V(\delta) \subseteq X$. As every variable in X appears as a literal in u_X, it follows that either $u_X \vdash \delta$ or $u_X \vdash \neg\delta$. This gives two cases:

1. In the case $u_X \vdash \delta$ recall that $\delta \models \phi_2 \vee \neg K_2 \vee \neg u_X$ which means $K_2 \wedge u_X \wedge \delta \models \phi_2$ and hence $K_2 \wedge u_X \models \phi_2$.
2. In the case $u_X \vdash \neg\delta$ recall that $K_1 \wedge u_X \wedge \neg\phi_1 \models \delta$ hence $K_1 \wedge u_X \wedge \neg\phi_1 \models \perp$ and thus $K_1 \wedge u_X \models \phi_1$.

In either case, we can conclude either $K_1 \wedge K_2 \wedge u_X \models \phi_1$ or $K_1 \wedge K_2 \wedge u_X \models \phi_2$. With ϕ_1 and ϕ_2 being arbitrary, it follows that Y_1 and Y_2 are independent given X modulo $K_1 \wedge K_2$. □

The previous Theorem 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).

5.2 Relationship to Multivalued Dependencies

Our attention now turns to showing that multivalued dependencies for propositional formulae arise as a special case of Darwiche's logical conditional independence.

Theorem 5.3 (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 $K_{Y_1X} \wedge K_{Y_2X} \vdash K_{Y_1Y_2X}$ holds.*

Proof. Suppose that Y_1 and Y_2 are independent given X modulo K, and consider a world u satisfying both K_{Y_1X} and K_{Y_2X}. We must demonstrate that u satisfies $K_{Y_1Y_2X}$. Assume for the sake of contradiction that u satisfies $\neg K_{Y_1Y_2X}$ as well. Construct the formulae ϕ_1 and ϕ_2 by choosing ϕ_1 as the conjunction of literals over Y_1 satisfied by u, and ϕ_2 as the conjunction of literals over Y_2 satisfied by u. Also choose u_X to be the conjunction of literals over X satisfied by u. It follows that $K \wedge u_X \vdash K_{Y_1Y_2X} \wedge u_X$ and $K_{Y_1Y_2X} \wedge u_X \vdash \neg\phi_1 \vee \neg\phi_2$ hence $K_{Y_1Y_2X} \wedge u_X \vdash \neg\phi_1 \vee \neg\phi_2$, for otherwise there would exist a model of $K_{Y_1Y_2X} \wedge u_X$ equivalent to u on Y_1Y_2X. Being that Y_1 and Y_2 are independent given X modulo K, and u_X is X-complete, it follows that $K \wedge u_X \vdash \neg\phi_1$ or $K \wedge u_X \vdash \neg\phi_2$. However, this means that either $K_{Y_1X} \wedge u_X \vdash \neg\phi_1$ or $K_{Y_2X} \wedge u_X \vdash \neg\phi_2$ which is a contradiction as u satisfies ϕ_1, ϕ_2, u_X, and both projections of K. Therefore, u is a model of $K_{Y_1Y_2X}$ showing that $K_{Y_1X} \wedge K_{Y_2X} \vdash K_{Y_1Y_2X}$ holds.

Conversely, suppose that $K_{Y_1X} \wedge K_{Y_2X} \vdash K_{Y_1Y_2X}$ holds. Consider formulae ϕ_1 and ϕ_2 such that $V(\phi_1) \subseteq Y_1$ and $V(\phi_2) \subseteq Y_2$ along with a world u such that $K \wedge u_X \vdash \phi_1 \vee \phi_2$. We must show that either $K \wedge u_X \vdash \phi_1$ or $K \wedge u_X \vdash \phi_2$. Observe that $K_{Y_1X} \wedge K_{Y_2X} \wedge u_X \vdash u_X \vdash \phi_1 \vee \phi_2$ by the Projection Theorem, and by the Splitting Criterion Y_1 and Y_2 are independent given X modulo $K_{Y_1X} \wedge K_{Y_2X} \wedge u_X$. Thus, either $K_{Y_1X} \wedge K_{Y_2X} \wedge u_X \vdash \phi_1$ or $K_{Y_1X} \wedge K_{Y_2X} \wedge u_X \vdash \phi_2$, hence either $K \wedge u_X \vdash \phi_1$ or $K \wedge u_X \vdash \phi_2$ as required. □

It is worthwhile making two observations: as $K_{Y_1Y_2X} \vdash K_{Y_1X} \wedge K_{Y_2X}$ always holds, so this projection criterion can be rephrased as asserting independence if and only if $K_{Y_1Y_2X} \equiv K_{Y_1X} \wedge K_{Y_2X}$. Furthermore, $K_{Y_1X} \wedge K_{Y_2X}$ is effectively a splitting of $K_{Y_1Y_2X}$ which implies a converse to the Splitting Criterion.

Theorem 5.4. *If X and Y are disjoint subsets of V then a propositional theory K satisfies $X \twoheadrightarrow Y$ if and only if Y and $V - (XY)$ are independent given X modulo K.*

Proof. By the Projection Criterion and our observation it follows that Y and $V - (XY)$ are independent given X if and only if $K \equiv K_{XY} \wedge K_{\overline{Y}}$. Observe that a world $w \in [K_{XY} \wedge K_{\overline{Y}}]$ if and only if there exists $u \in [K_{XY}]$ and $v \in [K_{\overline{Y}}]$ such that $w_{XY} = u_{XY}$ and $w_{\overline{Y}} = v_{\overline{Y}}$. This is in turn equivalent to having K satisfy $X \twoheadrightarrow Y$. □

It is now clear that multivalued dependencies, logical conditional independence, and syntax splitting are different aspects of the same underlying phenomenon. As corollaries of the Splitting Criterion, we see that (PCR) and (CR) ensure that belief revision operators preserve the satisfaction of multivalued dependencies which are partially or fully complied with.

Theorem 5.5. *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 $X \twoheadrightarrow Y$.*

Proof. Observe that when writing $K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$ we have $V((K * \phi)_{XY}) \subseteq XY$ and $V(K_{\overline{Y}}) \subseteq \overline{Y}$ hence the resulting theory satisfies $X \twoheadrightarrow Y$ via the Splitting Criterion. □

Theorem 5.6. *If $*$ is a belief revision operator which fully complies with $X \twoheadrightarrow Y$, K satisfies $X \twoheadrightarrow Y$, and $V(\phi) \subseteq XY$ then $K * \phi$ satisfies $X \twoheadrightarrow Y$.*

Proof. Observe that when writing $K * \phi \equiv (K * \phi)_{XY} \wedge K_{\overline{Y}}$ we have $V((K * \phi)_{XY}) \subseteq XY$ and $V(K_{\overline{Y}}) \subseteq \overline{Y}$ hence the resulting theory satisfies $X \twoheadrightarrow Y$ via the Splitting Criterion. □

6 Discussion

6.1 Sources of Multivalued Dependencies

In our approach we consider multivalued dependencies to be specified by the knowledge engineer as part of the domain knowledge, rather than extracted automatically from the knowledge base. This avoids using possibly-spurious conditional independencies that just happen to hold. As well, we also avoid the cost of determining all potential conditional independencies prior to a revision, given that checking whether a single conditional independence holds is known to be in Π_2^p (Lang, Liberatore, and Marquis 2002).

This raises the question of how a knowledge engineer might determine appropriate multivalued dependencies This question (in the analogous case of conditional irrelevance assertions) is discussed in (Delgrande and Peppas 2018), where a number of sources are suggested: 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.

6.2 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 follows 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 $\overline{Y}$ 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.

Our results on the relationship between multivalued dependencies and syntax splitting apply as well in the unconditional setting. As an application, the postulates (P1) and (P2) from (Peppas et al. 2015) can be restated as follows:

Theorem 6.1. *Let $*$ be a belief revision operator.*

- *(P1) is equivalent to the following: if K satisfies $\emptyset \twoheadrightarrow Y$ and $V(\phi) \subseteq Y$ then $(K * \phi)_{\overline{Y}} \equiv K_{\overline{Y}}$.*
- *(P2) is equivalent to the following: if K satisfies $\emptyset \twoheadrightarrow Y$ and $V(\phi) \subseteq Y$ then $(K * \phi)_Y \equiv (K_Y * \phi)_Y$.*

6.3 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, which use the ranking conditions (CS1), (CS2), and (CS3) to factor a ranking into smaller components.

There are also a number of natural variations on our postulates which seem to merit consideration:

1. Study a "parallelised" variant of our postulates, wherein we consider revising by $\phi \wedge \psi$ with $V(\phi) \subseteq XY$ and $V(\psi) \subseteq \overline{Y}$, with our postulate saying something like $K * (\phi \wedge \psi) = (K * \phi)_{XY} \wedge (K * \psi)_{\overline{Y}}$.
2. Study a "prioritised" variant of our postulates, wherein we consider revising by ϕ with $\phi \wedge K_{\overline{Y}}$ is not necessarily consistent, with our postulate saying something like $K * \phi = K_{\overline{Y}} \circledast (K * \phi)_{XY}$ for some operator $\circledast$.
3. Study belief revision operators which fully comply with all multivalued dependencies simultaneously, and consider the analogues of (P1) and (P2) in this case which would amount to the following:

   CP1. If K satisfies $X \twoheadrightarrow Y$ with $Y \cap X = \emptyset$ and $V(\phi) \subseteq Y$ then $(K * \phi)_{\overline{Y}} \equiv K_{\overline{Y}}$.

   CP2. If K satisfies $X \twoheadrightarrow Y$ with $Y \cap X = \emptyset$ and $V(\phi) \subseteq Y$ then $(K * \phi)_Y \equiv (K_Y * \phi)_Y$.
4. Study postulates which make use of conditional independencies in the sense of Darwiche, which unlike multivalued dependencies need not partition the entire vocabulary.

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).

7 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 study of 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., and Pearl, J. 1994. Symbolic causal networks. In *AAAI*, 238–244.

Darwiche, A. 1997. A logical notion of conditional independence: Properties and applications. *Artificial Intelligence* 97(1-2):45–82.

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.

Gardenfors, P. 1990. Belief revision and relevance. In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume 1990, 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. 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 *IJCAI*, 1131–1137.

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.

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.

Lang, J.; Liberatore, P.; and Marquis, P. 2002. Conditional independence in propositional logic. *Artificial Intelligence* 141(1-2):79–121.

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

Peppas, P. 2008. Belief revision. *Foundations of Artificial Intelligence* 3:317–359.