

Modeling Belief Change on Epistemic States

Jianbing Ma, Weiru Liu

School of Electronics, Electrical Engineering and Computer Science,
Queen's University Belfast, Belfast BT7 1NN, UK
{jma03, w.liu}@qub.ac.uk

Abstract

Belief revision always results in trusting new evidence, so it may admit an unreliable one and discard a more confident one. We therefore use belief change instead of belief revision to remedy this weakness. By introducing epistemic states, we take into account of the strength of evidence that influences the change of belief. In this paper, we present a set of postulates to characterize belief change by epistemic states and establish representation theorems to characterize those postulates. We show that from an epistemic state, a corresponding ordinal conditional function by Spohn can be derived and the result of combining two epistemic states is thus reduced to the result from combining two corresponding ordinal conditional functions proposed by Laverny and Lang. Furthermore, when reduced to the belief revision situation, we prove that our results induce all the Darwiche and Pearl's postulates.

Introduction

Belief change depicts the process that an agent revises its beliefs when new information/knowledge is received. Often, new information is to some extent conflicting with the agent's current beliefs. Therefore, belief change shall formalize a process as how a new set of beliefs can be obtained based on both the current beliefs and new information. Belief change has gained considerable attention in philosophy and artificial intelligence. Belief revision is a very important subfield of belief change which always believes in new information and thus revises the current beliefs to accommodate new evidence to reach a consistent set of beliefs. Most studies on belief revision are based on the AGM postulates (Alchourrón, Gärdenfors & Makinson 1985).

The AGM postulates formulated in the propositional setting in (Katsuno & Mendelzon 1991), denoted as R1-R6, characterize what a revision operator shall comply. The R1 postulate, also called *success postulate*, requires that the revision result of a belief set K by a proposition μ (new information) should always maintain μ being believed. This postulate has been questioned (e.g., (Boutilier, Friedman, & Halpern 1998; Hansson 1999; Booth, Meyer & Wong 2006), etc), because it is often undesirable in situations where an agent's observation is imprecise or uncertain. This issue was

also mentioned (though not explicitly pointed out) by Darwiche and Pearl (Darwiche & Pearl 1997). They argued that to remedy this, the strength of new evidence should be introduced in the revision process, i.e., if information about the strength of evidence is available, the evidence (including the prior beliefs) with a stronger strength should overrule the one with a weaker strength. This issue was very briefly discussed in the *Future Work Section* in (Darwiche & Pearl 1997). However, from this brief discussion it is not clear how a revision operator that incorporates a strength value m could be defined.

As a landmark research on transition from belief revision to revision on epistemic states, in addition to a set of modified AGM postulates (denoted as R*1-R*6) on epistemic revision operators, four additional postulates (C1-C4) were presented to formalize *iterated* revision operators (Darwiche & Pearl 1997). More work on the advantage of using epistemic states in belief revision can be found in (Benferhat *et al* 2000; Booth & Meyer 2006; Meyer 2000; Jin & Thielscher 2007; Nayak, Pagnucco, & Peppas 2003).

To overcome the weakness of R1, in this paper, we give a formal definition of epistemic state and then present a set of postulates, denoted as B0-B6, to characterize operators on *iterated belief change* by epistemic states (the phrase "belief change" is to emphasize that the success postulate is no longer required, as opposed to belief revision). In this belief change framework, both the prior beliefs and new evidence are represented by epistemic states. We also provide representation theorems for our postulates.

An interesting phenomena in the research on epistemic revision is that almost all the papers on this topic use Spohn's ordinal conditional function (OCF) (Spohn 1988) or its variants as illustrative examples. In this paper, we examine the ordinal conditional function and its combination rule (Laverny & Lang 2005) in our belief change framework. We prove that from an epistemic state, a corresponding ordinal conditional function can be derived and we also prove that the result of combining two epistemic states is equivalent to the result from combining two corresponding OCFs. This finding is important since it provides a justification for the combination rule of ordinal conditional functions proposed in (Laverny & Lang 2005), which is the most notable combination rule for OCFs so far.

Furthermore, when reduced to the belief revision situation

where new evidence must be accepted in the revised belief set, we prove that our result can induce Darwiche and Pearl's (DP's) belief revision postulates (Darwiche & Pearl 1997).

The rest of the paper is organized as follows. Section 2 provides preliminaries and our motivation on belief change. In Section 3, we provide formal definitions for epistemic states and study iterated belief change on epistemic states as well as a justification for the combination rule proposed for OCFs. In Section 4, we prove that our postulates can induce all the iterated belief revision postulates. Finally, we discuss future work and conclude the paper in Section 5.

Motivation

Preliminaries: We consider a propositional language $\mathcal{L}$ defined on a finite set $\mathcal{A}$ of propositional atoms. A proposition ϕ is constructed by propositional atoms with logic connections $\neg, \wedge$ in the standard way. An interpretation ω (or possible world) is a function mapping $\mathcal{A}$ onto $\{0, 1\}$. The set of all ω s on $\mathcal{A}$ is denoted as W . Function ω can be extended to any proposition in $\mathcal{L}$ in the usual way, $\omega : \mathcal{L} \rightarrow \{0, 1\}$. An interpretation ω is a model of (or satisfies) ϕ iff $\omega(\phi) = 1$, denoted as $\omega \models \phi$. We use $Mod(\phi)$ to denote the set of models for ϕ .

A pre-order $\leq$ defined on any set A is a reflexive and transitive relation over $A \times A$. $\leq$ is total iff for all elements $a, b \in A$, either $a \leq b$ or $b \leq a$ holds. Conventionally, a strict order $<$ and an indifferent relation $=$ can be induced by $\leq$ such that $\forall a, b \in A$, $a < b$ iff $a \leq b$ but $b \not\leq a$, and $a = b$ iff $a \leq b$ and $b \leq a$.

Motivations: We use an example to demonstrate the motivation of belief change on epistemic states.

Example 1 (Derived from Example 17 in (Darwiche & Pearl 1997)) *We face a murder trial with two main suspects, John and Mary. Initially, it appears that the murder was committed by one person, hence, our belief can be characterized as: $(John \wedge \neg Mary) \vee (\neg John \wedge Mary)$. As the trial unfolds, however, we receive a very reliable testimony incriminating John, followed by another reliable testimony incriminating Mary. At this point, how can we judge these two pieces of evidence in relation to the one-person theory? If we do not strongly believe in the one-person theory, we should believe that both suspects took part in the murder; whilst if we believe in the one-person theory more strongly than the testimonies, then based on belief revision, we are forced to believe that Mary is the murderer no matter how compelling the evidence against John is (because the evidence incriminating Mary comes later). This is counterintuitive in two accounts: first, whether we should dismiss the testimony against John should depend on how strongly we believe in it compared with how strongly we believe in the one-person theory; second, whether we should dismiss the testimony against John should also depend on how strongly we believe in it compared with how strongly we believe the testimony against Mary (if one-person theory is to be held).*

This example is very interesting. First, it shows that without providing the strength of evidence, any belief revision postulates could lead us to the wrong way, i.e., John can be either a murderer or innocent, any postulate favoring the

prior belief (i.e. the one-person theory) may let a potential murderer escape (*John*), and any postulate favoring the testimony may convict an innocent person (*John*). Second, it shows that the underlying assumption of belief revision, i.e. that the most recent evidence has the highest priority, has a major drawback: even if the testimony against *John* is more compelling than that to *Mary*, *Mary* is still the murderer if the one-person theory is believed. Therefore, who is the murderer somehow depends on which evidence arrives last.

The assumption of giving priority to the most recent evidence is also questioned in (Delgrande, Dubois, & Lang 2006). To get around this assumption, iterated belief revision is taken as a prioritized merging where a set of evidence is prioritized according to their reliability (strength) rather than the *order* that these pieces of evidence are received. The revised (or the merged) result is a consistent belief set such that when a more reliable piece of evidence is inconsistent with a less reliable piece of evidence, the reliable evidence should be preserved in the revised belief set.

Now we examine Example 1 again by incorporating the strengths of evidence. Suppose that the prior knowledge, *one-person theory*, and two pieces of evidence, *John is the murderer*, and *Mary is the murderer* (no matter in which order the evidence is collected), are available and have evidence strengths α , β and γ respectively. With a rational belief change process, we should find the real murderer(s) according to those strengths. For example, if one-person theory is to be kept (i.e., $\alpha > \max(\beta, \gamma)$), then the murderer is *John* if $\beta > \gamma$; the murderer is *Mary* if $\beta < \gamma$; and we do not know who the murderer is if $\beta = \gamma$. This solution is obviously more intuitive than the result obtained from iterated belief revision where *Mary* has to be the murderer under one-person theory regardless how strong the evidence against *John* is.

From the analysis above, we get the following

1. Epistemic states should be used to represent both the prior beliefs and new evidence to resolve such problems.
2. It should be the strength of evidence, not the order that the evidence is collected, determines the outcome of belief change.

Darwiche and Pearl also realized this issue and concluded that a natural way to resolve this is to allow *the outcome of belief change depends on the strength of evidence triggering the change*. In the Future Work Section (Darwiche & Pearl 1997), Darwiche and Pearl introduced the notions of *evidence strength* and *degree of acceptance*. In particular, they stated that a proposition μ is accepted by an epistemic state Φ to degree m if it takes a piece of evidence $\neg\mu$ with strength m to retract μ from Φ . Formally, they gave the following definition.

Definition 1 (Darwiche & Pearl 1997) *Proposition μ is accepted by an epistemic state Φ to degree m (written $\Phi \models_m \mu$) precisely when*

$$\Phi \not\models \neg\mu; \Phi \circ_m \neg\mu \models \neg\mu; \text{ and } \Phi \circ_m \neg\mu \not\models \mu.$$

Here $\circ_m$ is a revision operator incorporating value m .

However, how to design such operators and how to manage a sequence of revision operators $\circ_{m_1}, \circ_{m_2}, \dots, \circ_{m_n}$ remain to be investigated.

Belief Change by Epistemic States

Our approach on belief change based on epistemic states is inspired by Darwiche and Pearl's idea in Definition 1. That is we want to describe the belief change process that can truly reflect the strengths of prior beliefs and evidence. Below we first define epistemic states and then give the corresponding postulates to characterize belief change.

Epistemic states

Ordinal conditional function (Spohn 1988) is commonly regarded as a form of epistemic state.

An **ordinal conditional function**, also known as a **ranking function** or a **kappa function**, commonly denoted as κ , is a function from a set of possible worlds to the set of ordinal numbers with its **belief set** defined as $Bel(\kappa) = \varphi$ where $Mod(\varphi) = \{w | \kappa(w) = 0\}$. Value $\kappa(w)$ is understood as the degree of **disbelief** of world w . So the smaller the $\kappa(w)$ value, the more plausible the world is. The ranking value of a proposition μ is defined as:

$$\kappa(\mu) = \min_{w \models \mu} \kappa(w).$$

The combination of two ordinal conditional functions κ_1 and κ_2 is defined in (Lavervy & Lang 2005) as

$$(\kappa_1 \oplus \kappa_2)(w) = \kappa_1(w) + \kappa_2(w) - \min_{w \in W} (\kappa_1(w) + \kappa_2(w)) \quad (1)$$

It is applicable only when $\min_{w \in W} (\kappa_1(w) + \kappa_2(w)) < +\infty$.

We can see that there is a re-normalization step in the combination of OCFs to make the minimal worlds have kappa value 0. However, when modeling the belief change process, we want to solely concentrate on the nature of the *changing* process, and ignore the re-normalization. So we do not simply use OCFs as our epistemic states but define epistemic states as follows.

Definition 2 An **epistemic state** Φ is a mapping from W to $Z \cup \{-\infty, \infty\}$ where Z is the set of integers.

Obviously, this definition follows the spirits of OCF and the epistemic state defined in (Meyer 2000) (in that paper it is defined as a mapping from W to the set of ordinals, and such a definition was also implied in (Williams 1994)).

Definition 3 Let Φ be an epistemic state, the **belief set** of Φ , denoted as $Bel(\Phi)$, is defined as $Bel(\Phi) = \psi$ where $Mod(\psi) = \min(W, \leq_\Phi)$. Here $\leq_\Phi$ is a total pre-order relation on W such that $w_1 \leq_\Phi w_2$ iff $\Phi(w_1) \geq \Phi(w_2)$.

Here we can see that the belief set derived from an OCF or an epistemic state defined in (Meyer 2000) is the same as that in Definition 3, i.e., the belief set has all the most plausible worlds as its models. In addition, ignoring re-normalization will not affect the belief set of an epistemic state.

In both an OCF and Meyer's epistemic state, a possible world with a lower value is more plausible than one with a higher value whilst in our definition, it is the opposite. Another major difference is that the range of our definition of epistemic states is $Z \cup \{-\infty, \infty\}$ instead of ordinals. These two differences will enable us to avoid the re-normalization.

Φ can be extended to proposition formulae.

Definition 4 (Extension of epistemic state) Let Φ be an epistemic state, then Φ can be extended to any propositional formula μ such that $\Phi(\mu) = \max_{w \models \mu} (\Phi(w))$.

We denote $f_\Phi(\mu) = \Phi(\mu) - \Phi(\neg\mu)$ as the *strength of preference* on μ which is interpreted as the relative preference of μ over $\neg\mu$. The notion of strength of preference is not new. In (Vickers 2001), it stated "The strength of preference for a proposition X over a proposition Y is the expectation (based on an agent's probabilistic beliefs) that a world in X is better than a world in Y ." Our notion follows a similar explanation. For distinction, we call $\Phi(\mu)$ the weight of μ .

Now we consider a special case of epistemic state.

Definition 5 An epistemic state Φ is called a **simple epistemic state** iff $\exists \mu$ such that

$$\Phi(w) = \begin{cases} m & \text{for } w \models \mu, \\ 0 & \text{for } w \not\models \mu. \end{cases}$$

Here m is an integer and we simply write Φ as (μ, m) .

Simple epistemic states are introduced to make the representation of postulates in the next subsection simpler.

For a simple epistemic state $\Phi = (\mu, m)$, we have $Bel(\Phi) = \mu$ if $m > 0$; $Bel(\Phi) = \neg\mu$ if $m < 0$, and $Bel(\Phi) = W$ if $m = 0$. It also shows that if $m = 0$, this simple epistemic state is totally ignorant.

In the next section, for illustration and simplicity, we will first consider simple epistemic states to model evidence and then extend to general cases that use any epistemic states to model evidence.

Postulates

Following (Katsuno & Mendelzon 1991) and (Darwiche & Pearl 1997), we also use the notation $form(w_1, w_2, \dots)$ to denote a proposition μ which has $w_1, w_2, \dots$ as its models, that is, $Mod(\mu) = \{w_1, w_2, \dots\}$. By abuse of notations, we also use $form(A)$ to denote a proposition μ such that $Mod(\mu) = A$. We also assume that when an epistemic state Φ is embedded in a propositional formula, it stands for $Bel(\Phi)$, e.g. $\Phi \wedge \psi$ means $Bel(\Phi) \wedge \psi$; $\Phi \models \mu$ means $Bel(\Phi) \models \mu$ and $Mod(\Phi) = Mod(Bel(\Phi))$, etc.

A belief change operator associates two epistemic states to a resulting one satisfying certain postulates. Now we investigate what postulates a belief change operator shall comply in order to make the belief change process rational. Since the main idea of belief change is to allow strengths of the current beliefs and evidence to determine the outcome, we propose the following seven postulates (with their explanations) to characterize belief change.

B0 $\Phi \circ (\mu, 0) = \Phi$ for any μ .

Explanation: If an agent obtains indifferent information, then its beliefs shall not be changed.

B1 If $f_\Phi(\mu) > m$, then $\Phi \circ (\neg\mu, m) \models \mu$; If $f_\Phi(\mu) < m$, then $\Phi \circ (\neg\mu, m) \models \neg\mu$; If $f_\Phi(\mu) = m$, then $\Phi \circ (\neg\mu, m) \not\models \mu$ and $\Phi \circ (\neg\mu, m) \not\models \neg\mu$.

Explanation: This intuitively follows Definition 1, which shows that it should be the strengths of agent's current beliefs and new evidence that determine the outcome of belief change. When $f_\Phi(\mu) = m$ and new evidence gives

$(\neg\mu, m)$, then the prior beliefs prefer μ with strength m and the evidence prefers $\neg\mu$ with the same strength, so neither μ nor $\neg\mu$ should be believed.

B2 If $\Phi \wedge \mu$ is satisfiable, then $\Phi \circ (\mu, m) \equiv \Phi \wedge \mu$ if $m > 0$.

Explanation: If new evidence is consistent with an agent's current beliefs, then the agent incorporates the new evidence into its beliefs.

B3 If μ is satisfiable, then $\Phi \circ (\mu, m)$ is an epistemic state.

Explanation: If new evidence is consistent, then a new epistemic state should be obtained after belief change.

B4 If $\Phi_1(w_1) = \Phi_2(w_2)$ and $(\mu_1, m)(w_1) = (\mu_2, n)(w_2)$, then $(\Phi_1 \circ (\mu_1, m))(w_1) = (\Phi_2 \circ (\mu_2, n))(w_2)$.

Explanation: If two agents hold the same (prior) weight on two possible worlds w_1 and w_2 , and two new pieces of evidence they receive also have the same weight on these two worlds respectively, then both agents shall still have the same weight on the two worlds after changing their prior beliefs with new evidence.

B5 $\Phi \circ (\mu, m) \circ (\mu, n) = \Phi \circ (\mu, m + n)$.

Explanation: The strength of preference is reinforced when multiple pieces of evidence supporting it are received. This postulate requires that these pieces of evidence are independently obtained. A typical scenario of this postulate is in a rumor spreading process. Cumulative rumors usually destroy people's current beliefs.

B6 $\Phi \circ (\mu, m) \circ (\psi, n) = \Phi \circ (\psi, n) \circ (\mu, m)$.

Explanation: The order of evidence (received) shall not influence the outcome of belief change. This and **B1** together determine the main difference between a belief change process and a belief revision process.

Now we give the following representation theorem for these postulates¹.

Theorem 1 A belief change operator $\circ$ satisfies postulates **B0-B6** precisely when

$$\forall w, (\Phi \circ (\mu, m))(w) = \Phi(w) + (\mu, m)(w) \quad (2)$$

This theorem describes the belief change process when the evidence is expressed as a simple epistemic state. From this theorem, we get that if an epistemic state Φ is such that $\Phi(w) = m_i$ iff $w \models \mu_i$, $1 \leq i \leq n$, then we have

$$\forall w, \Phi(w) = ((\mu_1, m_1) \circ \dots \circ (\mu_n, m_n))(w).$$

Now we can extend to the case that the evidence can be expressed as any epistemic state. To accommodate this extension, we only need to naturally extend postulate B6 to B6* as follows.

B6* $\Phi_1 \circ \Phi_2 \circ \Phi_3 = \Phi_1 \circ \Phi_3 \circ \Phi_2$.

Then we have the following representation theorem.

Theorem 2 A belief change operator $\circ$ satisfies postulates **B0-B5** and **B6*** precisely when

$$\forall w, (\Phi \circ \Phi')(w) = \Phi(w) + \Phi'(w). \quad (3)$$

¹ All proofs can be found in a longer version: <http://www.cs.qub.ac.uk/~W.Liu/flairs09-24.pdf>.

This theorem shows that with our definition of epistemic states, only *pointwise addition operation* can serve as the role that the outcome of belief change is triggered by the strength of evidence.

Example 2 (Example 1 Revisited) Let the prior epistemic state about the initial one-person theory be Φ such that

$$\Phi(\text{John}, \neg\text{Mary}) = \alpha, \Phi(\neg\text{John}, \text{Mary}) = \alpha, \alpha > 0$$

$$\Phi(\text{John}, \text{Mary}) = 0, \Phi(\neg\text{John}, \neg\text{Mary}) = -\infty$$

Let the two testimonies incriminating John and Mary be represented by two simple epistemic states as (John, β) and (Mary, γ) where $\beta, \gamma > 0$. Let $\circ$ be a belief change operator on epistemic state Φ satisfying postulates **B0-B6**, then applying $\circ$ to these three epistemic states, we have the final epistemic state Φ_{JM} as

$$\Phi_{JM}(\text{John}, \neg\text{Mary}) = \alpha + \beta, \Phi_{JM}(\neg\text{John}, \text{Mary}) = \alpha + \gamma,$$

$$\Phi_{JM}(\text{John}, \text{Mary}) = \beta + \gamma, \Phi_{JM}(\neg\text{John}, \neg\text{Mary}) = -\infty$$

If the one-person theory is the most reliable evidence, i.e. $\alpha > \max(\beta, \gamma)$, then we have both

$$\Phi_{JM}(\text{John}, \neg\text{Mary}) > \Phi_{JM}(\text{John}, \text{Mary})$$

$$\Phi_{JM}(\neg\text{John}, \text{Mary}) > \Phi_{JM}(\text{John}, \text{Mary})$$

which show that the murderer is one of them which is intuitively what we want to get. Furthermore, who exactly committed the crime is based on the strengths of the evidence β vs. γ . When $\beta > \gamma$, i.e., when $\Phi_{JM}(\text{John}, \neg\text{Mary}) > \Phi_{JM}(\neg\text{John}, \text{Mary})$, then John is the murderer; when $\gamma > \beta$, i.e., when $\Phi_{JM}(\text{John}, \neg\text{Mary}) < \Phi_{JM}(\neg\text{John}, \text{Mary})$, then Mary is the murderer. An interesting situation is when $\beta = \gamma$, we cannot decide who is the murderer and this is intuitively correct also, because the evidence is not against one over the other.

On the other hand, if $\alpha < \min(\beta, \gamma)$, then the result of belief change $(\text{John} \wedge \text{Mary})$ suggests that both John and Mary are murderers which is also intuitively explainable.

This example shows that belief change operators satisfying postulates **B0-B6** do allow the strengths of evidence play the essential role in determining the outcome of belief change.

A Justification on OCF Combination

Below we prove that our definition of epistemic state and the belief change rule $(\Phi \circ \Phi')(w) = \Phi(w) + \Phi'(w)$ can induce the ordinal conditional function and its combination method defined by Equation 1, respectively. Thus, our postulates justify the rationale for the combination of ordinal conditional functions using Equation 1.

Definition 6 Let Φ be an epistemic state defined in Definition 2. We define $\kappa_\Phi : \mathcal{L} \rightarrow \mathbb{Z}$ as a corresponding function for Φ such that

$$\kappa_\Phi(\mu) = \max_{w \in W} (\Phi(w)) - \Phi(\mu). \quad (4)$$

We have the following immediate result.

Proposition 1 Let Φ be an epistemic state and κ_Φ be its corresponding function based on Definition 6, then κ_Φ is an ordinal conditional function.

The following theorem shows that the result of changing an epistemic state Φ with another epistemic state Φ' is equivalent to the result of combining the two corresponding functions derived from Φ and Φ' respectively.

Theorem 3 Let Φ and Φ' be two epistemic states and $\Phi \circ \Phi'$ be the resulting epistemic state after belief change. Let κ_Φ , $\kappa_{\Phi'}$, and $\kappa_{\Phi \circ \Phi'}$ be their corresponding functions respectively, then we have $\forall w, \kappa_{\Phi \circ \Phi'}(w) = (\kappa_\Phi \oplus \kappa_{\Phi'})(w)$.

The figure below illustrates Theorem 3 intuitively.

$$\begin{array}{ccc} (\Phi \circ \Phi')(w) & = & \Phi(w) \oplus \Phi'(w) \\ \downarrow & & \downarrow \quad \downarrow \\ \kappa_{\Phi \circ \Phi'}(w) & = & (\kappa_\Phi \oplus \kappa_{\Phi'})(w) \end{array}$$

Figure 1. Illustration of Theorem 3.

Belief Change vs. Belief Revision

DP's postulates on iterated belief revision

To demonstrate the inadequacy of iterated belief revision, in (Darwiche & Pearl 1997) Darwiche and Pearl deployed a set of examples to show how counterintuitive results can be obtained if AGM postulates are to be followed. They recommended that to ensure the rational preservation of conditional beliefs² during (iterated) belief revision, a revision process shall be carried out on epistemic states rather than on their belief sets. With this intention, epistemic states are used to represent an agent's original beliefs and new evidence is taken as a propositional formula. Correspondingly, they modified the AGM postulates (R1-R6) to obtain a set of revised postulates (R*1-R*6) for iterated epistemic revision. Let $\circ_r$ be a revision operator, the revised postulates are

- R*1** $\Psi \circ_r \mu$ implies μ .
- R*2** If $\Psi \wedge \mu$ is satisfiable, then $\Psi \circ_r \mu \equiv \Psi \wedge \mu$.
- R*3** If μ is satisfiable, then $\Psi \circ_r \mu$ is also satisfiable.
- R*4** If $\Psi_1 = \Psi_2$ and $\mu_1 \equiv \mu_2$, then $\Psi_1 \circ_r \mu_1 \equiv \Psi_2 \circ_r \mu_2$.
- R*5** $(\Psi \circ_r \mu) \wedge \phi$ implies $\Psi \circ_r (\mu \wedge \phi)$.
- R*6** If $(\Psi \circ_r \mu) \wedge \phi$ is satisfiable, then $\Psi \circ_r (\mu \wedge \phi)$ implies $(\Psi \circ_r \mu) \wedge \phi$.

In the above postulates, Ψ (or Ψ_1, Ψ_2) stands for an epistemic state and μ and ϕ are propositional formulae. $\Psi \circ_r \mu$ is an epistemic state after revising Ψ with revision operator $\circ_r$ by μ . When an epistemic state (e.g., Ψ) is embedded in a propositional formula, it is used to stand for its belief set³ (e.g., $Bel(\Psi)$) not an epistemic state for simplification purpose. For example, $\Psi \wedge \phi$ means $Bel(\Psi) \wedge \phi$.

²Let ψ and α be two propositional formulae and let $\circ_r$ be a belief revision operator, then the revision of ψ by α is a new propositional formula and is denoted as $\psi \circ_r \alpha$. $\beta|\alpha$ is called a conditional belief of ψ if $\psi \circ_r \alpha \models \beta$ (Darwiche & Pearl 1997).

³In (Darwiche & Pearl 1997), it demonstrates that each epistemic state has an associated belief set which characterizes the set of propositions that the agent is committed at any given time. In this section (only), the phrase epistemic states follows DP's setting.

To regulate iterated epistemic revision to preserve conditional beliefs, Darwiche and Pearl gave the following four additional postulates which are for four disjoint types of conditional beliefs.

- C1** If $\alpha \models \mu$, then $(\Psi \circ_r \mu) \circ_r \alpha \equiv \Psi \circ_r \alpha$.
- C2** If $\alpha \models \neg\mu$, then $(\Psi \circ_r \mu) \circ_r \alpha \equiv \Psi \circ_r \alpha$.
- C3** If $\Psi \circ_r \alpha \models \mu$, then $(\Psi \circ_r \mu) \circ_r \alpha \models \mu$.
- C4** If $\Psi \circ_r \alpha \not\models \neg\mu$, then $(\Psi \circ_r \mu) \circ_r \alpha \not\models \neg\mu$.

$\Psi \circ_r \alpha \models \beta$ here stands for $Bel(\Psi \circ_r \alpha) \models \beta$.

Two representation theorems are given to characterize these two sets of postulates. But first, we introduce the definition of faithful assignment.

Definition 7 (Darwiche & Pearl 1997) Let W be the set of all worlds (interpretations) of a propositional language $\mathcal{L}$ and suppose that the belief set of any epistemic state belongs to $\mathcal{L}$. A function that maps each epistemic state Φ to a total pre-order $\leq_\Phi^r$ on worlds W is said to be a faithful assignment if and only if:

1. $w_1, w_2 \models \Phi$ only if $w_1 =_\Phi^r w_2$.
2. $w_1 \models \Phi$ and $w_2 \not\models \Phi$ only if $w_1 <_\Phi^r w_2$.
3. $\Phi = \Psi$ only if $\leq_\Phi^r = \leq_\Psi^r$.

Theorem 4 (Darwiche & Pearl 1997) A revision operator $\circ_r$ satisfies postulates **R*1-R*6** precisely when there exists a faithful assignment that maps each epistemic state Φ to a total pre-order $\leq_\Phi^r$ such that:

$$Mod(\Phi \circ_r \mu) = \min(Mod(\mu), \leq_\Phi^r).$$

This representation theorem shows that the revised belief is determined by μ and the total pre-order associated with Φ .

Theorem 5 (Darwiche & Pearl 1997) Suppose that a revision operator $\circ_r$ satisfies postulates **R*1-R*6**. The operator satisfies postulates **C1-C4** iff the operator and its corresponding faithful assignment satisfy:

- CR1** If $w_1 \models \mu$ and $w_2 \models \mu$, then $w_1 \leq_\Phi^r w_2$ iff $w_1 \leq_{\Phi \circ_r \mu}^r w_2$.
- CR2** If $w_1 \models \neg\mu$ and $w_2 \models \neg\mu$, then $w_1 \leq_\Phi^r w_2$ iff $w_1 \leq_{\Phi \circ_r \mu}^r w_2$.
- CR3** If $w_1 \models \mu$ and $w_2 \models \neg\mu$, then $w_1 <_\Phi^r w_2$ only if $w_1 <_{\Phi \circ_r \mu}^r w_2$.
- CR4** If $w_1 \models \mu$ and $w_2 \models \neg\mu$, then $w_1 \leq_\Phi^r w_2$ only if $w_1 \leq_{\Phi \circ_r \mu}^r w_2$.

This representation theorem shows that an epistemic revision operator $\circ_r$ satisfies postulate Ci iff condition CRi is satisfied, $1 \leq i \leq 4$.

Belief change versus belief revision

In this section, we want to prove that when reducing to the belief revision situation, our result should derive all the belief revision postulates including Darwiche and Pearl's belief revision postulates R*1-R*6 and C1-C4 (Darwiche & Pearl 1997). Because the essential difference between belief change and belief revision is that belief revision takes the most recent evidence as the most reliable one. From our view of belief change, belief revision always assigns a reasonably larger strength of preference to the most recent evidence. For instance, given a prior epistemic state Φ , we can create evidence (μ, m^*) such that $m^* = \max_{w \in W}(\Phi(w)) - \min_{w \in W}(\Phi(w)) + 1$, then the strength of new evidence μ is stronger than the strength of preference on any subset A

of 2^W . To simulate an iterated belief revision process with evidence sequence $\mu_1, \dots, \mu_n$ with a belief change operator, we only need to let $m_1 = m^*$, and $m_i = 2^{i-1} * m_1$, $1 < i \leq n$ to ensure that new evidence always has a stronger strength than previous pieces of evidence. For convenience, we name (μ, m) **revision evidence** if $m \geq m^*$.

Theorem 6 *Let Φ be an epistemic state and (μ, m) be revision evidence, then a belief change operator $\circ$ satisfying postulates **B0-B6** guarantees the existence of a faithful assignment which maps each epistemic state Φ to a total preorder $\preceq_\Phi$ such that:*

$$Mod(\Phi \circ (\mu, m)) = \min(Mod(\mu), \preceq_\Phi)$$

This theorem, accompanied with Theorem 4, shows that the belief set obtained from belief change on an epistemic state Φ with revision evidence (μ, m) is equivalent to the belief set obtained from belief revision on Φ with formula μ . That is $Bel(\Phi \circ (\mu, m)) = Bel(\Phi \circ_r \mu)$.

Theorem 7 *Let Φ be an epistemic state and (μ, m) be revision evidence, a belief change operator $\circ$ satisfies postulates **B0-B6** will lead to the following:*

CR1* *If $w_1 \models \mu$ and $w_2 \models \mu$, then $w_1 \preceq_\Phi w_2$ iff $w_1 \preceq_{\Phi \circ (\mu, m)} w_2$.*

CR2* *If $w_1 \models \neg\mu$ and $w_2 \models \neg\mu$, then $w_1 \preceq_\Phi w_2$ iff $w_1 \preceq_{\Phi \circ (\mu, m)} w_2$.*

CR3* *If $w_1 \models \mu$ and $w_2 \models \neg\mu$, then $w_1 \prec_\Phi w_2$ only if $w_1 \prec_{\Phi \circ (\mu, m)} w_2$.*

CR4* *If $w_1 \models \mu$ and $w_2 \models \neg\mu$, then $w_1 \preceq_\Phi w_2$ only if $w_1 \preceq_{\Phi \circ (\mu, m)} w_2$.*

This theorem shows our belief change postulates lead to Darwiche-Pearl's iterated belief revision postulates C1-C4.

Conclusion

In this paper, we presented a framework of belief change by epistemic states, in which prior beliefs and new evidence are both represented by epistemic states. We proposed a set of postulates to characterize belief change operators. We also provided representation theorem for our postulates. In addition, we show that these postulates can be seen as a justification of the combination of two OCFs.

When reducing to the iterated belief revision situation by Darwiche and Pearl (where new evidence is a propositional formula), our postulates induce all of DP's postulates.

Belief revision requires that the most recent evidence is the most reliable one. When it is not the case, belief revision cannot be applied, as studied in (Boutilier, Friedman, & Halpern 1998; Hansson 1999; Booth, Meyer & Wong 2006; Delgrande, Dubois, & Lang 2006). However, these approaches cannot achieve the goal that "the outcome of belief change depends on the strength of evidence triggering the change" (Darwiche & Pearl 1997). Our approach, inspired by this observation, solves this problem by deploying epistemic states to model evidence and their strengths, so that the strengths of evidence trigger belief change.

With Theorem 2, we find that a belief change operator is in fact a merging operator. Many merging operators on epistemic states were discussed in (Meyer 2000). Here we provide a full characterization of one of them, although with a different definition of epistemic states. Our belief change is also closely related to non-prioritized belief revision (cf. (Hansson 1999) for an overview), however, strength does not play an important role in non-prioritized belief revision.

Postulate **B4** is somehow a rather strong postulate which restricts the combination of epistemic states to be point-wise combination. Further research on how to weaken this postulate would be interesting and worthwhile.

References

- C E Alchourrón, P Gärdenfors, and D Makinson. On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision. *Jour. Symb. Log.*, 50, 510-530, 1985.
- S Benferhat, S Konieczny, O Papini, and R P Pérez. Iterated Revision by Epistemic States: Axioms, Semantics and Syntax. *Proc. of 14th ECAI*, 13-17, 2000.
- C Boutilier, N Friedman, and J Y Halpern. Belief Revision with Unreliable Observations. *Proc. of 15th AAAI*, 127-134, 1998.
- R Booth, T Meyer and K S Wong. A Bad Day Surfing Is Better Than a Good Day Working: How to Revise a Total Preorder. *Proc. of KR'06*, 2006.
- R Booth and T Meyer. Admissible and Restrained Revision. *JAIR*, 26, 127-151, 2006.
- J Delgrande, D Dubois, and J Lang. Iterated Revision as Prioritized Merging. *Proc. of 10th KR*, 210-220, 2006.
- A Darwiche and J Pearl. On the Logic of Iterated Belief Revision. *Artif. Intel.*, 89, 1-29, 1997.
- Y Jin and M Thielscher. Iterated Belief Revision, Revised. *Artif. Intel.*, 171, 1-18, 2007.
- S O Hansson. A Survey of Non-prioritized Belief Revision. *Erkenntnis*, 50, 1999.
- H Katsuno and A O Mendelzon. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intel.*, 52, 263-294, 1991.
- N Laverny and J Lang. From Knowledge-based Programs to Graded Belief Based Programs, Part I: On-line Reasoning. *Synthese*, 147(2), 277-321, 2005.
- T Meyer. Merging Epistemic States *Proc. of PRICAI 2000*, 286-296, 2000.
- A C Nayak, M Pagnucco, and P Peppas. Dynamic Belief Revision Operators. *Artif. Intel.*, 146, 193-228, 2003.
- W Spohn. Ordinal Conditional Functions: A Dynamic Theory of Epistemic States. In W. Harper and B. Skyrms (Eds.), *Causation in Decision, Belief Change, and Statistics*, 2, 105-134, 1988 by Kluwer Academic Publishers.
- J M Vickers. Values and Beliefs about Them. www.ihpst.univ-paris1.fr/_sources/jvic_valeurs_en.pdf, 2001.
- M.A. Williams. Transmutations of Knowledge Systems. *Proc. of 4th KR*, 619-629, 1994.