

Iterated Belief Base Revision: A Dynamic Epistemic Logic Approach

Marlo Souza,¹ Álvaro Moreira,² Renata Vieira³

¹Department of Computer Science, UFBA, Salvador, Brazil

²Institute of Informatics, UFRGS, Porto Alegre, Brazil

³Polytechnic School, PUCRS, Porto Alegre, Brazil

msouza1@ufba.br, alvaro.moreira@ufrgs.br, renata.vieira@pucrs.br

Abstract

AGM's belief revision is one of the main paradigms in the study of belief change operations. In this context, belief bases (prioritised bases) have been largely used to specify the agent's belief state - whether representing the agent's 'explicit beliefs' or as a computational model for her belief state. While the connection of iterated AGM-like operations and their encoding in dynamic epistemic logics have been studied before, few works considered how well-known postulates from iterated belief revision theory can be characterised by means of belief bases and their counterpart in dynamic epistemic logic. This work investigates how priority graphs, a syntactic representation of preference relations deeply connected to prioritised bases, can be used to characterise belief change operators, focusing on well-known postulates of Iterated Belief Change. We provide syntactic representations of belief change operators in a dynamic context, as well as new negative results regarding the possibility of representing an iterated belief revision operation using transformations on priority graphs.

Introduction

One of the most influential models for belief change is the so-called AGM paradigm, named after the authors of the seminal work (Alchourrón, Gärdenfors, and Makinson 1985). Although the AGM's approach has brought profound developments for the problem of belief dynamics, influencing areas such as Computer Science, Artificial Intelligence, and Philosophy, it has not been immune to criticism (Hansson 1992; Rott 2000; Darwiche and Pearl 1997).

Particularly, Hansson (1992) criticises the use of deductively closed sets of formulas in the AGM paradigm, providing examples for which not just the meaning, but also the structure of the beliefs explicitly held by an agent may influence the change. This author proceeds to construct a different notion of belief revision, which became known as Belief Base Change, relying on the structure of the information believed by the agent.

While many studies (Nebel 1991; Williams 1994; 1995; Rott 2009) propose belief change operators based on the syntactic structure of an agent's explicit beliefs, or on the syntactic representation of an agent's belief state, few

formal connections have been established between belief change postulates and these belief base change operators - some examples of works that investigate such connection are (Hansson 1994; Fermé, Krevneris, and Reis 2008; Fermé, Garapa, and Reis 2017). The axiomatic characterisation of syntactic-based belief change, however, has been concentrated on the one-shot behaviour of these operations and, thus, are not able to clarify their iterated or dynamic behaviour.

Belief base change operators are iterated, in the sense that they result in changing the belief state of the agent (Nayak, Pagnucco, and Peppas 2003). As such, it is necessary to establish which formal properties these change operations satisfy on a dynamic sense, as studied in the literature of Iterated Belief Change (Darwiche and Pearl 1997; Nayak, Pagnucco, and Peppas 2003; Jin and Thielscher 2007).

A step into connecting belief change operators and belief base change operators has been achieved by the work of Rott (2009) and Liu (2011) which show that several belief change operations can be defined using transformations on some form of prioritised belief bases called priority graphs. Souza et al. (2016), on the other hand, provide examples of belief change operations which cannot be defined in such a way.

While these latter works focus on using syntactic representations to construct belief change operations, they do not provide a strong connection between the syntactic-based transformations provided and the formal properties of the belief change operations they represent. To our knowledge, there is no proposed characterisation in the literature of iterated belief change postulates based on syntactic representations of an agent's belief base.

In this work, we study the relationship between iterated belief change postulates and properties of belief base change operators based on the counterpart of these operators as transformations on Liu's priority graphs (2011). We obtain constraints on belief base change operators that guarantee satisfaction of some important postulates in the area of Iterated Belief Change.

We also show the limits of expressibility of belief change operators by means of changes in belief bases. These negative results provide us with a deeper understanding of the connections between belief base change and iterated belief

change, helping to delineate the limits of the correspondence between semantic and syntactic theories of dynamic belief change.

This work is structured as follows: we start discussing the background knowledge in the next section. Then, we study how iterated belief change postulates can be characterised through transformations on priority graphs. Further, we show how our results can be used to construct syntactic representations of an iterated belief change operation, and in Section **Negative Results** we present some impossibility results for this characterisation. Finally, we discuss the related literature and present some final considerations.

Background

Much work has been conducted investigating syntactic representations of the explicit beliefs of an agent (Hansson 1992; Williams 1994; Rott 2009; De Jongh and Liu 2009; Baltag, Fiutek, and Smets 2016). It is well-known that an agent's belief state - usually represented by a preference relation among worlds - can also be syntactically represented by means of orders among sentences, as well-investigated by Lafage and Lang (2005).

We start by introducing the syntactic representations we will use in this work to encode an agent's belief base. This structure, known as a priority graph, or P-graph for short, was proposed by Liu (2009; 2011) and was further developed by Van Benthem, Grossi and Liu (2014), and Souza et al (2016). We will also present their relation to preference models - a generalisation of the models used in the area of belief revision to encode an agent's belief state. With this connection, we will be able to characterise well-known postulates of iterated belief revision using operations on priority graphs, thus connecting belief base change operations and iterated belief revision.

Definition 1. (Liu 2011) Let P be a countable set of propositional symbols and $\mathcal{L}_0(P)$ the language of classical propositional sentences over the set P . A P-graph is a tuple $G = \langle \Phi, \prec \rangle$ where $\Phi \subset \mathcal{L}_0(P)$, is a set of propositional sentences and \prec is a strict partial order on Φ .

Given a non-empty set of possible worlds W and a valuation of propositional symbols over worlds of W , the order \prec of a P-graph can be lifted to an order on worlds. Such an ordering can be called a preference (or a plausibility) relation.

Definition 2. (Liu 2011) Let $G = \langle \Phi, \prec \rangle$ be a P-graph, and $v : P \rightarrow 2^W$ be a valuation function of propositions over a non-empty set W of possible worlds. The preference relation \leq_G on W induced by P-graph G is defined as follows:

$$w \leq_G w' \text{ iff } \forall \varphi \in \Phi : ((w' \models \varphi \Rightarrow w \models \varphi) \text{ or } \exists \psi \in \Phi : \psi \prec \varphi, w \models \psi, \text{ and } w' \not\models \psi)$$

More yet, if a P-graph G is finite, Van Benthem, Grossi, and Liu (2014) have shown that a preference relation \leq_G , defined as above, satisfies well-foundedness - a property deeply connected with Lewis Limit Assumption, commonly required for semantic models of an agent's belief state.

It is worthy of notice that preference relations are a common semantic representation of an agents belief state. In

fact, in Belief Revision Theory (Alchourrón, Gärdenfors, and Makinson 1985), an important semantic representation of agents belief states was proposed by Adam Grove, which is commonly known as Grove's systems of spheres (SOS) (Grove 1988), or simply as Grove's spheres.

A generalisation of Grove's spheres was given by Girard (2008), called preference models or order models, in the context of Preference Logic. Similar models have been proposed before in the context of Qualitative Decision Theory (Boutilier 1994), Non-monotonic Reasoning (Kraus, Lehmann, and Magidor 1990), among others. These models have been applied to the study of Dynamic Belief Revision by the work of Baltag and Smets (Baltag and Smets 2008), Girard and Rott (2014), Liu (2011), and Souza et al. (2016; 2017) and proved to be an expressible model for an agent's belief state. In this work, we will adopt Souza's (2016) definition of preference models - which requires the preference relation \leq to have a well-founded strict part.

Definition 3. (Souza 2016) A preference model is a tuple $M = \langle W, \leq, v \rangle$ where W is a set of possible worlds, \leq is a reflexive, transitive relation over W with a well-founded strict part, and $v : P \rightarrow 2^W$ a valuation function.

From the above definitions, it is easy to see that from a P-graph we can construct a preference model by taking the preference relation induced by such a graph.

Definition 4. Let $G = \langle \Phi, \prec \rangle$ be a P-graph and let $M = \langle W, \leq, v \rangle$ be a preference model. We say M is induced by G iff $\leq = \leq_G$.

The induction of preference models from P-graphs raises the question about the relations between these two structures. Liu (2011) shows that any finite model with a reflexive and transitive accessibility relation has an equivalent P-graph.

Theorem 5. (Liu 2011) Let $M = \langle W, R \rangle$ be a mono-modal Kripke structure. The following two statements are equivalent:

1. The relation R is reflexive and transitive;
2. There is a priority graph $G = \langle \Phi, \prec \rangle$ and a valuation v s.t. $\forall w, w' \in W. wRw' \text{ iff } w \leq_G w'$.

It is well-known from work on Iterated Belief Revision (Darwiche and Pearl 1997; Nayak, Pagnucco, and Pappas 2003) that dynamic belief change operations can be described by a transformation in the agent's belief state.

As such, we can define dynamic belief change operators using the following notion of dynamic operators on preference models, where $\mathbb{M}(\mathcal{L}_{\leq}(P))$ is the class of all preference models for a logic language $\mathcal{L}_{\leq}(P)$ (Souza, Moreira, and Vieira 2017):

Definition 6. (Souza, Moreira, and Vieira 2017) Let $\star : \mathbb{M}(\mathcal{L}_{\leq}(P)) \times \mathcal{L}_0(P) \rightarrow \mathbb{M}(\mathcal{L}_{\leq}(P))$, we say \star is a dynamic operator on preference models if for any preference model $M = \langle W, \leq, v \rangle$ and formula $\varphi \in \mathcal{L}_0$, we have that $\star(M, \varphi) = \langle W, \leq_{\star}, v \rangle$. In other words, an operation on preference models is called a dynamic operator iff it only changes the relation of preference models.

Liu et al (2011; 2014) shows that these dynamic belief changes can also be described by means of changes in the priority graphs representing the agent's belief base. In the following, $\mathbb{G}(P)$ denotes the set of all P-graphs constructed over a set P of propositional symbols.

Definition 7. We call a P-graph transformation any function $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$.

A P-graph transformation is, thus, a transformation in the agent's belief base, as represented by a priority graph.

Since P-graphs and preference models are translatable into one another, it is easy to connect P-graph transformations and dynamic operators as well.

Definition 8. Let $\star : \mathbb{M}(\mathcal{L}_{\leq}(P)) \times \mathcal{L}_0(P) \rightarrow \mathbb{M}(\mathcal{L}_{\leq}(P))$ be a dynamic operator and $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a P-graph transformation. We say \star is induced by \dagger if for any preference model $M \in \mathbb{M}(\mathcal{L}_{\leq}(P))$ and any P-graph $G \in \mathbb{G}(P)$, if M is induced by G then the preference model $\star(M, \varphi)$ is induced by the P-graph $\dagger(G, \varphi)$, where φ is any propositional formula in $\mathcal{L}_0(P)$,

Notice that not all P-graph transformations induce dynamic operators. We say that syntactically different P-graphs are equivalent if they induce the same preference model. As such, if a P-graph transformation changes equivalent P-graphs in inconsistent ways, no dynamic operator can satisfy the condition of Definition 8.

For example, consider a P-graph transformation that changes the P-graph $p \prec q$ into the P-graph $p \prec q$, and changes the P-graph $p \wedge q \prec p \wedge \neg q \prec \neg p \wedge q \prec \neg p \wedge \neg q$ into $p \wedge q \prec \neg p \wedge q \prec p \wedge \neg q \prec \neg p \wedge \neg q$. Such a transformation cannot induce any dynamic operator since the original P-graphs are equivalent, i.e., induce the same models, but the resulting P-graphs are not. As such, we define the notion of relevant graph transformation.

Definition 9. We say that a P-graph transformation \dagger is relevant if there is some dynamic operator \star that is induced by it.

Notice that, the existence of relevant P-graph transformations is guaranteed by the previous representation results of Liu (2011) and Van Benthem, Grossi, and Liu (2014) on the characterisation of some dynamic operators by means of changes in P-graphs.

Priority Changes Satisfying Postulates for Iterated Belief Changes

In this work, we investigate the connection between the properties satisfied by dynamic belief change operations, focusing on the postulates studied in the field of Iterated Belief Revision, and the properties satisfied by the P-graph transformations that encode these operations. We aim to understand better which belief change operations can or cannot be encoded this way and, thus, the differences between dynamic belief change based on semantic models and based on syntactic representations.

As such, the main results of our work can be stated as the characterisations provided in Propositions 11, 12, 13, 14, 15 and 16, as well as the negative results provided in Fact 21

and Corollary 23. An interesting application of these results will be obtained in the next Section, in which we construct a P-graph transformation to implement the operation of lexicographic revision (Nayak, Pagnucco, and Peppas 2003) and obtain, as a corollary, the harmony result proved by Van Benthem, Grossi, and Liu (2014) stating the correctness of this transformation.

The use of postulates to encode rational constraints in the way an agent must change her beliefs is a defining characteristic of the AGM approach to Belief Revision (Alchourrón, Gärdenfors, and Makinson 1985). These postulates, however, are usually defined by means of constraints on changes in the agent's belief state, thus, in our case, on preference models. We must, then, define what it means for a P-graph transformation to satisfy some postulate (or property) for belief change operators.

Definition 10. We say that a P-graph transformation \dagger satisfies a postulate \mathcal{P} if (i) \dagger is relevant and (ii) any dynamic operator \star induced by \dagger satisfies postulate \mathcal{P} .

AGM belief revision says very little about how to change one agent's beliefs repeatedly. In fact, it has been observed that the AGM approach allows some counter-intuitive behaviour in the iterated case (Darwiche and Pearl 1997). As a result, different authors have proposed additional postulates that encode rational ways to change one's beliefs in an iterated way.

Most famous among them is the work of Darwiche and Pearl (1997). They propose a set of postulates known as the DP postulates for iterated revision. Let $S = \langle W, \leq \rangle$ be a SOS and $S' = \langle W, \leq_{*\varphi} \rangle$ the result of revising the SOS S by a formula φ , the DP postulates can be stated as:

(DP-1) If $w, w' \in \llbracket \varphi \rrbracket$, then $w \leq_{*\varphi} w'$ iff $w \leq w'$

(DP-2) If $w, w' \notin \llbracket \varphi \rrbracket$, then $w \leq_{*\varphi} w'$ iff $w \leq w'$

(DP-3) If $w \in \llbracket \varphi \rrbracket$ and $w' \notin \llbracket \varphi \rrbracket$, then $w < w' \Rightarrow w <_{*\varphi} w'$

(DP-4) If $w \in \llbracket \varphi \rrbracket$ and $w' \notin \llbracket \varphi \rrbracket$, then $w \leq w' \Rightarrow w \leq_{*\varphi} w'$

We want to provide a set of constraints on P-graph transformations that guarantee that the dynamic operators induced by them satisfy these postulates. With that, we wish to study the connections between the classes of belief base revision operations and iterated belief change operators.

Let us start with DP-1. The postulate DP-1 states that for any two worlds w, w' satisfying φ , there is no reason for their relative order to change in the agent's belief state after revision. In terms of P-graphs, this means that for any formula that w' satisfies in the changed priority graph, either w must also satisfy it, or there must be a formula that is preferred to it and that w satisfies. We can ensure this property guaranteeing that the resulting P-graph is related to the original by a set of constraints in how it must be changed.

Proposition 11. Let $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a relevant P-graph transformation. If, for any P-graph $G = \langle \Phi, \prec \rangle$ and propositional formula $\varphi \in \mathcal{L}_0(P)$, the P-graph $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ satisfies the conditions below, then \dagger satisfies DP-1:

1. For all $\xi \in \Phi$, there is some $\xi' \in \Phi_{\dagger}$ s.t.

(a) $\varphi \wedge \xi \equiv \varphi \wedge \xi'$ and

- (b) $\forall \psi' \in \Phi_{\dagger}$, if $\psi' \prec_{\dagger} \xi'$ then
 $\psi' \equiv \varphi$ or
there is $\psi \in \Phi$ s.t. $\varphi \wedge \psi \equiv \varphi \wedge \psi'$ and $\psi \prec \xi$;
2. For all $\xi \in \Phi_{\dagger}$, $\xi \equiv \varphi$ or there is some $\xi' \in \Phi$ s.t.
- (a) $\varphi \wedge \xi \equiv \varphi \wedge \xi'$ and
(b) $\forall \psi' \in \Phi$, if $\psi' \prec \xi'$ then there is $\psi \in \Phi_{\dagger}$ s.t.
 $\varphi \wedge \psi \equiv \varphi \wedge \psi'$ and
 $\psi \prec_{\dagger} \xi$.

Proof. By Definition 10 we have to prove that any dynamic operator \star induced by \dagger satisfies the postulate DP-1. Let \dagger be a P-graph transformation satisfying the conditions above, \star a dynamic operator induced by \dagger , $M = \langle W, \leq, v \rangle$ a preference model, and $\varphi \in \mathcal{L}_0$ a propositional formula. Also given $M_{\star\varphi} = \star(M, \varphi) = \langle W, \leq_{\star\varphi}, v \rangle$, take a P-graph $G = \langle \Phi, \prec \rangle$ inducing M s.t. $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ induces $\star(M, \varphi)$.
 \Leftarrow :

Take $w, w' \in \llbracket \varphi \rrbracket_M$ such that $w \leq_{\star\varphi} w'$, and take $\xi \in \Phi$ such that $M, w' \models \xi$ - notice that the case in which no such formula exists is trivial, since necessarily $w \leq w'$ in such case. Then $M_{\star\varphi}, w' \models \xi$, since ξ is propositional formula. Since $\xi \in \Phi$ then, by condition 1(a), there is some $\xi' \in \Phi_{\dagger}$ s.t. $\varphi \wedge \xi \equiv \varphi \wedge \xi'$. As such, $M_{\star\varphi}, w' \models \xi'$, given that \star is induced by \dagger . As $w \leq_{\star\varphi} w'$, either:

- (i) $M_{\star\varphi}, w \models \xi'$, or
(ii) there is some $\psi' \in \Phi_{\dagger}$ s.t. $\psi' \prec_{\dagger} \xi'$, $M_{\star\varphi}, w \models \psi'$ and $M_{\star\varphi}, w' \not\models \psi'$. Hence, by condition 1 (b), there is some $\psi \in \Phi$ s.t. $\varphi \wedge \psi \equiv \varphi \wedge \psi'$ and $\psi \prec \xi$. As such $M, w \models \psi$ and $M, w' \not\models \psi$.

From (i) and (ii) we conclude that $w \leq w'$

\Rightarrow : Similar to the case before. Take $w, w' \in \llbracket \varphi \rrbracket_M$ s.t. $w \leq w'$ and $\xi \in \Phi_{\dagger}$, use condition 2 to conclude that $w \leq_{\star\varphi} w'$. \square

The postulate DP-2 describes the same information as DP-1, only restricted to those worlds that do not satisfy φ . As such, we can provide a similar characterisation.

Proposition 12. Let $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a relevant P-graph transformation. If, for any P-graph $G = \langle \Phi, \prec \rangle$ and propositional formula $\varphi \in \mathcal{L}_0(P)$, the P-graph $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ satisfies the conditions below, then \dagger satisfies DP-2 :

1. For all $\xi \in \Phi$, there is some $\xi' \in \Phi_{\dagger}$ s.t.
- (a) $\neg\varphi \wedge \xi \equiv \neg\varphi \wedge \xi'$ and
(b) $\forall \psi \in \Phi$, if $\psi \prec \xi$ then there is $\psi' \in \Phi_{\dagger}$ s.t.
i. $\neg\varphi \wedge \psi \equiv \neg\varphi \wedge \psi'$ and
ii. $\psi' \prec_{\dagger} \xi'$;
2. For all $\xi \in \Phi_{\dagger}$, $\xi \equiv \varphi$ or there is some $\xi' \in \Phi$ s.t.
- (a) $\neg\varphi \wedge \xi \equiv \neg\varphi \wedge \xi'$ and
(b) $\forall \psi \in \Phi_{\dagger}$, if $\psi \prec_{\dagger} \xi$ then
 $\psi \equiv \varphi$ or
there is $\psi' \in \Phi$ s.t. $\neg\varphi \wedge \psi \equiv \neg\varphi \wedge \psi'$ and $\psi' \prec \xi'$.

Proof. Similar to that of Proposition 11. \square

The postulate DP-3 states that for any two worlds w satisfying φ and w' not satisfying it, after revision by φ , if w was preferable to w' then it must continue to be so in the agent's belief state. In terms of P-graphs, we can guarantee this condition requiring that, if there was a formula in the original P-graph that w satisfied and w' did not, there must be a formula in the revised P-graph s.t. w satisfies and w' does not. Therefore, we can characterise DP-3.

Proposition 13. Let $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a relevant P-graph transformation. If, for any P-graph $G = \langle \Phi, \prec \rangle$ and propositional formula $\varphi \in \mathcal{L}_0(P)$, the P-graph $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ satisfies the condition below, then \dagger satisfies DP-3 :

- For all $\xi \in \Phi$ there is some $\xi' \in \Phi_{\dagger}$ s.t.
- (a) $\varphi \wedge \xi \vdash \xi'$,
(b) $\neg\varphi \wedge \xi' \vdash \xi$, and
(c) $\forall \psi' \in \Phi_{\dagger}$, if $\psi' \prec_{\dagger} \xi'$ then $\psi' \equiv \varphi$ or there is $\psi \in \Phi$ s.t. $\varphi \wedge \psi \vdash \psi'$, $\neg\varphi \wedge \psi' \vdash \psi$ and $\psi \prec \xi$.

Proof. Let \dagger be a P-graph transformation satisfying the conditions above, \star a dynamic operator induced by \dagger , $M = \langle W, \leq, v \rangle$ a preference model and $\varphi \in \mathcal{L}_0(P)$ a propositional formula. Also given $M_{\star\varphi} = \star(M, \varphi) = \langle W, \leq_{\star\varphi}, v \rangle$, take $G = \langle \Phi, \prec \rangle$ be a P-graph inducing M s.t. $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ induces $\star(M, \varphi)$.

Take $w, w' \in W$ s.t. $M, w \models \varphi$, $M, w' \not\models \varphi$ and $w < w'$. From $w < w'$ we conclude that there is some $\xi \in \Phi$ s.t. $M, w \models \xi$, $M, w' \not\models \xi$ and for all $\forall \psi \prec \xi$, if $M, w' \models \psi$ then $M, w \models \psi$. Since $\xi \in \Phi$, there is some $\xi' \in \Phi_{\dagger}$ s.t. $\varphi \wedge \xi \vdash \xi'$, $\neg\varphi \wedge \xi' \vdash \xi$. As such $M_{\star\varphi}, w \models \xi'$, $M_{\star\varphi}, w' \not\models \xi'$ and for all $\psi' \prec_{\dagger} \xi'$, if $M_{\star\varphi}, w' \models \psi'$ then there is $\psi \in \Phi$ s.t. $\varphi \wedge \psi \vdash \psi'$, $\neg\varphi \wedge \psi' \vdash \psi$ and $\psi \prec \xi$. As such $M, w' \models \psi$ and, thus, $M, w \models \psi$. Since $\varphi \wedge \psi \vdash \psi'$, then $M_{\star\varphi}, w \models \psi'$, thus $w \prec_{\dagger} w'$. \square

Along the same lines, as we did for DP-3, we can characterise postulate DP-4.

Proposition 14. Let $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a relevant P-graph transformation. If, for any P-graph $G = \langle \Phi, \prec \rangle$ and propositional formula $\varphi \in \mathcal{L}_0(P)$, the P-graph $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ satisfies the condition below, then \dagger satisfies DP-4 :

- For all $\xi \in \Phi_{\dagger}$, $\xi \equiv \varphi$ or there is some $\xi' \in \Phi$ s.t.
- (a) $\varphi \wedge \xi' \vdash \xi$,
(b) $\neg\varphi \wedge \xi \vdash \xi'$, and
(c) $\forall \psi' \in \Phi$, if $\psi' \prec \xi'$ then there is $\psi \in \Phi_{\dagger}$ s.t. $\varphi \wedge \psi' \vdash \psi$, $\neg\varphi \wedge \psi \vdash \psi'$ and $\psi \prec \xi$.

Proof. Let \dagger be a P-graph transformation satisfying the conditions above, \star a dynamic operator induced by \dagger , $M = \langle W, \leq, v \rangle$ a preference model, and $\varphi \in \mathcal{L}_0(P)$ a propositional formula. Also given $M_{\star\varphi} = \star(M, \varphi) = \langle W, \leq_{\star\varphi}, v \rangle$, take $G = \langle \Phi, \prec \rangle$ be a P-graph inducing M s.t. $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ induces $\star(M, \varphi)$.

Take $w, w' \in W$ s.t. $M, w \models \varphi$, $M, w' \not\models \varphi$, and $w \leq w'$. Take $\xi \in \Phi_{\dagger}$ s.t. $M, w' \models \xi$. There is some $\xi' \in \Phi$ s.t. $\varphi \wedge \xi' \vdash \xi$, $\neg\varphi \wedge \xi \vdash \xi'$. As such, $M, w' \models \xi'$.

Since $w \leq w'$, either (i) $M, w \models \xi'$ or (ii) there is some $\psi' \prec \xi'$ s.t. $M, w \models \psi'$ s.t. $M, w' \not\models \psi'$. From (ii), we conclude that there is some $\psi \in \Phi_{\dagger}$ s.t. $\psi \prec \xi$, $\varphi \wedge \psi' \vdash \psi$, and $\neg\varphi \wedge \psi \vdash \psi'$. Since $\neg\varphi \wedge \psi \vdash \psi'$ and $M, w' \not\models \psi'$, we conclude that $M, w' \not\models \psi$ and, thus, $M_{\star\varphi}, w' \not\models \psi$, as ψ is a propositional formula. From (i) and (ii) we conclude that $w \leq_{\star\varphi} w'$. \square

We know that the operation of lexicographic revision (Nayak, Pagnucco, and Peppas 2003) can be defined by means of transformations on P-graphs (Liu 2011). Also, Nayak et al. (2003) have shown that the operation of lexicographic revision is completely characterized by the postulates DP-1, DP-2, and the following postulate known as Recalcitrance (REC).

(REC) If $w \in \llbracket \varphi \rrbracket$ and $w' \notin \llbracket \varphi \rrbracket$, then $w <_{\star\varphi} w'$.

As such, it is expected that we can characterise REC as well. Postulate REC requires for any world satisfying φ to be preferred to those not satisfying it. This requirement can be easily guaranteed if all the minimal elements of the changed P-graph which express relevant information (i.e., not equivalent to \perp nor \top) imply φ .

Proposition 15. *Let $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a relevant P-graph transformation. If, for any P-graph $G = \langle \Phi, \prec \rangle$ and propositional formula $\varphi \in \mathcal{L}_0(P)$, the P-graph $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ satisfies the condition below, then \dagger satisfies REC:*

- For all $\xi \in \Phi_{\dagger}$, either $\xi \equiv \top$, $\xi \equiv \perp$, $\xi \vdash \varphi$ or there is some $\psi \in \Phi_{\dagger}$ s.t. $\psi \prec_{\dagger} \xi$, $\psi \not\equiv \perp$ and $\psi \vdash \varphi$;
- There is some $\xi \in \Phi$ s.t. $\xi \vdash \varphi$.

Proof. It is immediate that if a graph transformation satisfies the condition above, the induced dynamic operator must satisfy REC since any world satisfying φ in a model induced by such a graph would be preferred over any world not satisfying φ . \square

Since DP-3 and DP-4 can be characterised by means of transformations on P-graphs, it is expected that a related postulate might be as well. The postulate of Independence below, proposed by Jin and Thielscher (2007) and independently by Booth and Meyer (2006), states that a revision operation may not create arbitrary conditional beliefs in the agent's belief state.

(IND) If $w \in \llbracket \varphi \rrbracket$ and $w' \notin \llbracket \varphi \rrbracket$, then $w \leq w' \Rightarrow w <_{\star\varphi} w'$.

The postulate IND is, in fact, a stronger form of both DP-3 and DP-4. As such, we can provide the following characterisation for it, based on the characterisation of DP-3 and DP-4.

Proposition 16. *Let $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ be a relevant P-graph transformation. If, for any P-graph $G = \langle \Phi, \prec \rangle$ and propositional formula $\varphi \in \mathcal{L}_0(P)$, the P-graph $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ satisfies the condition below, then \dagger satisfies IND:*

- For all $\xi' \in \Phi_{\dagger}$, $\xi' \equiv \varphi$ or there is some $\xi \in \Phi$ s.t.
 - (a) $\varphi \wedge \xi \vdash \xi'$,

(b) $\neg\varphi \wedge \xi' \vdash \xi$, and

(c) $\forall \psi' \in \Phi_{\dagger}$, if $\psi' \prec_{\dagger} \xi'$ then there is $\psi \in \Phi$ s.t. $\varphi \wedge \psi \vdash \psi'$, $\neg\varphi \wedge \psi' \vdash \psi$ and $\psi \prec \xi$.

(d) if $\xi' \not\vdash \varphi$, there is some $\psi' \prec_{\dagger} \xi'$ s.t. $\psi' \equiv \varphi$

Proof. Let \dagger be a P-graph transformation satisfying the conditions above, \star a dynamic operator induced by \dagger , $M = \langle W, \leq, v \rangle$ a preference model and $\varphi \in \mathcal{L}_0(P)$ a propositional formula. Also given $M_{\star\varphi} = \star(M, \varphi) = \langle W, \leq_{\star\varphi}, v \rangle$, take $G = \langle \Phi, \prec \rangle$ be a P-graph inducing M s.t. $\dagger(G, \varphi) = \langle \Phi_{\dagger}, \prec_{\dagger} \rangle$ induces $\star(M, \varphi)$.

Take $w, w' \in W$ s.t. $M, w \models \varphi$, $M, w' \not\models \varphi$ and $w \leq w'$. Take $\xi \in \Phi_{\dagger}$ s.t. $M, w' \models \xi$. There is some $\xi' \in \Phi$ s.t. $\varphi \wedge \xi' \vdash \xi$, $\neg\varphi \wedge \xi \vdash \xi'$. As such, $M, w' \models \xi'$.

Since $w \leq w'$, either (i) $M, w \models \xi'$ or (ii) there is some $\psi' \prec \xi'$ s.t. $M, w \models \psi'$ and $M, w' \not\models \psi'$. As such, there is $\psi \in \Phi_{\dagger}$ s.t. $\psi \prec \xi$, $\varphi \wedge \psi' \vdash \psi$, and $\neg\varphi \wedge \psi \vdash \psi'$. Since $\neg\varphi \wedge \psi \vdash \psi'$ and $M, w' \not\models \psi'$, we conclude that $M, w' \not\models \psi$ and, thus, $M_{\star\varphi}, w' \not\models \psi$, as ψ is a propositional formula. Further, since $M, w' \models \xi'$, then $\xi' \not\vdash \varphi$. As such, there is some $\psi' \equiv \varphi$. We conclude that (iii) $M, w' \not\models \psi'$ and $M, w \models \psi'$. From (i), (ii), and (iii), we conclude that $w <_{\star\varphi} w'$. \square

Deriving P-Graph Transformations from Belief Revision Policies

The results obtained above can be used to analyse any syntax-based belief revision policy (P-graph transformations) and derive which belief revision properties (or postulates) it satisfies. On the other hand, these results can also be used to derive implementations for a belief revision operator based on P-graphs. To illustrate this last point, let us examine the case of Lexicographic Revision.

Definition 17. (Girard 2008) *Let $M = \langle W, \leq, v \rangle$ be a preference model and φ a formula of $\mathcal{L}_0(P)$. We say the preference model $M_{\uparrow\varphi} = \langle W, \leq_{\uparrow\varphi}, v \rangle$ is the result of the lexicographic revision of M by φ , where*

$$w \leq_{\uparrow\varphi} w' \text{ iff } \begin{cases} w \leq w' & \text{if } w, w' \in \llbracket \varphi \rrbracket \text{ or} \\ w \leq w' & \text{if } w, w' \notin \llbracket \varphi \rrbracket \text{ or} \\ \text{true} & \text{if } w \in \llbracket \varphi \rrbracket \text{ and } w' \notin \llbracket \varphi \rrbracket \end{cases}$$

The operation above consists of making each world satisfying φ to be strictly more preferable than those not satisfying it, while maintaining the order otherwise.

It is well-known (Nayak, Pagnucco, and Peppas 2003) that lexicographic revision is completely characterized by the postulates DP-1, DP-2 and REC. Hence, using Propositions 11, 12 and 15, we can construct a P-graph transformation that satisfies these postulates.

A simple P-graph transformation that does satisfy Propositions 11, 12, and 15 is *prefixing* which was proposed by Van Benthem (2009) based on the work of Andréka, Ryan, and Schobbens (2002).

Definition 18. *The prefixing of a P-graph $G = \langle \Phi, \prec \rangle$ by a propositional formula $\varphi \in \mathcal{L}_0(P)$ is the P-graph $; (G, \varphi) = \langle \Phi \cup \{\varphi\}, \prec_{;\varphi} \rangle$, usually denoted by $\varphi; G$, where*

$$\prec_{;\varphi} = \prec \cup \{ \langle \varphi, \psi \rangle \mid \psi \in \Phi \}$$

Observe that the resulting P-graph maintains all formulas of Φ , thus satisfying Propositions 11 and 12, and includes a formula φ (or equivalent to it) that is preferred to all formulas in Φ , thus satisfying Proposition 15. As such, the dynamic operator induced by P-graph prefixing satisfies DP-1, DP-2, and REC. Since these three postulates completely characterize lexicographic revision, we conclude the following.

Corollary 19. *Let M be a preference model induced by a P-graph G and φ a propositional formula. The preference model $M_{\uparrow\varphi}$ is induced by the P-graph $\varphi; G$.*

Negative Results

While the previous results are encouraging, Souza et al. (2016) already showed that some belief change operators cannot be defined with P-graphs. As such, it must be the case that some postulates in the area cannot be represented by means of transformations on P-graphs - or at least not in a way in which it is jointly consistent with other postulates. To prove such a result, those authors show a simple fact about P-graphs: they cannot encode all the information about the models they induce.

Fact 20. (Souza et al. 2016) *Let $G = \langle \Phi, \prec \rangle$ a P-graph and φ a propositional formula. There is no propositional formula μ_φ s.t. for every preference model $M = \langle W, \leq_G, v \rangle$ induced by G and all $w \in W$, $w \models \mu_\varphi$ iff $w \in \text{Min}_{\leq_G} \llbracket \varphi \rrbracket$.*

Fact 20 above provides us with some clues to investigate which postulates cannot be characterised through transformations on priority graphs, namely those that refer directly to the minimal worlds of a model.

One trivial example of such a postulate is the property of an iterated belief change operator to be faithful to AGM's postulates (Alchourrón, Gärdenfors, and Makinson 1985), known as postulate FAITH below.

(FAITH) If $\llbracket \varphi \rrbracket \neq \emptyset$ then $w \in \text{Min}_{\leq} \llbracket \varphi \rrbracket$ iff $w \in \text{Min}_{\leq_{*\varphi}} W$

Notice that while FAITH says something about the minimal worlds of a model, it does not characterise this set in any way. To illustrate it, it suffices to realise that lexicographic contraction satisfies FAITH - which describes the change in the agent's belief state by changing the preference of all the worlds satisfying a certain propositional formula φ . As such, if a P-graph transformation satisfies the postulates DP-1 and REC then it satisfies the postulate FAITH.

Let us then consider some belief change operators requiring a characterisation of the changes in the belief state which is completely dependent on some set of minimal worlds. To construct such an operator, let us examine the postulate of Conditional Belief Change Minimisation (CB), proposed by Boutilier (1993). This postulate states that any iterated belief revision operation must minimise changes of conditional beliefs in the belief state of the agent.

(CB) If $w, w' \notin \text{Min}_{\leq} \llbracket \varphi \rrbracket$, then $w \leq w'$ iff $w \leq_{*\varphi} w'$.

Together with FAITH, postulate CB characterises a belief change operator that is completely defined by the changes in the minimal worlds satisfying some formula φ . As such, it is fairly easy to see that no graph transformation satisfies both FAITH and CB.

Fact 21. *No relevant P-graph transformation $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ satisfies both FAITH and CB.*

Proof. We suppose that there is a relevant P-graph transformation $\dagger : \mathbb{G}(P) \times \mathcal{L}_0(P) \rightarrow \mathbb{G}(P)$ satisfying both FAITH and CB and we will derive a contradiction. Take the preference model $M = \langle \{w_1, w_2, w_3\}, \leq_1, v \rangle$ s.t. $w_1 <_1 w_2 <_1 w_3$, $M, w_1 \models \neg p \wedge q$, $M, w_2 \models p \wedge \neg q$ and $M, w_3 \models p \wedge q$. Let G be a P-graph that induces M . Since \dagger satisfies both FAITH and CB, any dynamic operator \star induced by \dagger must satisfy that $\star(M, p) = \langle \{w_1, w_2, w_3\}, \leq'_1, v \rangle$ s.t. $w_2 <'_1 w_1 <'_1 w_3$ is induced by $\dagger(G, p) = \langle \Phi', \prec' \rangle$, i.e., for any $\xi \in \Phi'$ s.t. $M, w_3 \models \xi$ either $M, w_1 \models \xi$ or there is $\psi \in \Phi'$ s.t. $\psi \prec' \xi$, $M, w_1 \models \psi$ and $M, w_3 \not\models \psi$ and there is some $\xi \in \Phi'$ s.t. $M, w_1 \models \xi$ and $M, w_3 \not\models \xi$.

Consider now the model $M' = \langle \{w_1, w_3\}, \leq_2, v \rangle$ s.t. $w_1 <_2 w_3$ and v is the same as before. Clearly, M' is induced by G as well. Since \dagger satisfies both FAITH and CB, any dynamic operator \star induced by \dagger must satisfy that $\star(M', p) = \langle \{w_1, w_3\}, \leq'_2, v \rangle$ s.t. $w_3 <'_2 w_1$ is induced by $\dagger(G, p) = \langle \Phi', \prec' \rangle$. As such, for any $\xi \in \Phi'$ s.t. $M, w_1 \models \xi$ either $M, w_3 \models \xi$ or there is $\psi \in \Phi'$ s.t. $\psi \prec' \xi$, $M, w_3 \models \psi$ and $M, w_1 \not\models \psi$ and there is some $\xi \in \Phi'$ s.t. $M, w_3 \models \xi$ and $M, w_1 \not\models \xi$. But this is a contradiction with the previous statement, since the valuation v is the same in all models. \square

It is well known, however, that Natural Revision - an iterated revision operation proposed by Boutilier (1993) - satisfies both FAITH and CB and is definable on preference models.

Definition 22. *Let $M = \langle W, \leq, v \rangle$ be a preference model and φ a formula of $\mathcal{L}_0(P)$. We say the preference model $M_{\uparrow\varphi} = \langle W, \leq_{\uparrow\varphi}, v \rangle$ is the result of the natural revision of M by φ , where*

$$w \leq_{\uparrow\varphi} w' \text{ iff } \begin{cases} w \in \text{Min}_{\leq} \llbracket \varphi \rrbracket, \text{ or} \\ w \leq w' \text{ and } w, w' \notin \text{Min}_{\leq} \llbracket \varphi \rrbracket \end{cases}$$

As such, we can conclude that Natural Revision cannot be represented as a P-graph transformation.

Corollary 23. *There is no P-graph transformation that induces Natural Revision.*

This result shows that some important belief change operations are not definable through P-graph transformations. Notice that previous examples, provided by Souza et al (2016), have all been contraction operations. What these belief change operations have in common is that their definition is intrinsically characterised by the minimal worlds satisfying a certain formula φ . In other words, these operations are only well-defined on preference models if we require that preference models be well-founded, a requirement made in Souza's (2016) definition for these models but not in Girard's (2008). As such, this result reinforces our intuition that Fact 20 is the cause of the lack of expressibility of P-graph transformations.

Notice that, as FAITH, postulate CB is not solely responsible for the impossibility of expressing Natural Revision through P-graph transformations. There is, in fact, a trivial

belief change operation that satisfies CB and is expressible by P-graph transformation: the *null change* operation.

Definition 24. Let $M = \langle W, \leq, v \rangle$ be a preference model and φ a formula of $\mathcal{L}_0(P)$. We say the preference model $M_{\circ\varphi} = \langle W, \leq, v \rangle$ is the result of the null change of M by φ

The null operation is the operation of not changing anything in the agent's belief state. It clearly satisfies postulate CB and it is trivially induced by the *null change* P-graph transformation.

Definition 25. Let $G \in \mathbb{G}(P)$ be a P-graph and $\varphi \in \mathcal{L}_0(P)$ a propositional formula. We define the null change transformation of G by φ as $\odot(G, \varphi) = G$.

Clearly the only dynamic operator induced by the null change transformation \odot is the null change operator \circ . As such, \odot satisfies postulate CB.

Related Work

The AGM approach and the vast literature based on it relies mainly on extralogical characterisation of belief change operations. The first attempt to integrate belief change operation within a logic that we are aware of is the work of Segerberg (1999), which defines Dynamic Doxastic Logic (DDL).

Similar work has focused on embedding specific belief change operations within various epistemic logics to analyse dynamic phenomena in Formal Epistemology (Van Benthem 2007; Baltag and Smets 2008). Particularly, Girard (2008) and Van Benthem (2009) propose Dynamic Preference Logic (DPL), a dynamic logic which has been used to generalise AGM-like belief change operations (Liu 2011; Souza et al. 2016). Aiming to strengthen the connection between DPL and Belief Change, Souza, Moreira, and Vieira (2017) study how well-known belief change postulates can be characterised using DPL axioms.

While these studies connect Belief Change with Epistemic Logic and provide ways to use the results from one area within the other, their approach is mainly semantic. Research on Belief Base Change, however, stemming from the work of Hansson (1992), focus on constructing belief change operators based on syntactic representations of the agent's belief state.

Searching for rich syntactic representations of agents' explicit beliefs, several authors such as Williams (1994; 1995), Rott (1991; 2009) and Benferhat et al. (2002) propose different belief base change operations. These works, however, do not explore how the belief change operations constructed over these syntactic representations are connected to the postulates of Belief Change.

The work closest to ours is that of Delgrande, Dubois, and Lang (2006) which considers the notion of iterated belief revision, as studied by Darwiche and Pearl (1997) and Nayak, Pagnucco, and Peppas (2003), as a special case of the belief change operation of *merging*. The authors use syntactic structures, similar to prioritized bases, to construct merging operations and show that they satisfy well-known iterated belief revision postulates. More so, the authors propose codifications of these postulates using the syntactic structures proposed by their work, differently than previous work.

The main drawback of their codification of postulates, in our opinion, is that they are not general enough. The proposed codifications of the postulates are obtained by translating the desired postulates, e.g., Darwiche and Pearl's (1997) iterated revision postulates, using the operation of graph prefixing to stand for revision. However, as we know, graph prefixing does not equate revision but represents a specific iterated revision policy known as lexicographic revision. As such, a more general codification of these postulates by means of syntactic representations of the agent's belief state is still an open problem.

Liu (2011) has shown that preference models can be encoded by syntactic structures known as P-graphs. Since preference models have been used to model an agent's belief state (Girard 2008; Girard and Rott 2014), Liu's priority graphs can be seen as a syntactic representation of an agent's belief state as well. More yet, Van Benthem, Grossi, and Liu (2014) and Souza (2016) have shown that this representation can be used to construct well-known belief changing operations from iterated belief change literature. As before, however, the authors do not consider how the formal properties of a belief change operator are reflected in its construction based on transformations of priority graphs.

Conclusion

This work has explored codifications of iterated belief change postulates in Dynamic Preference Logic using the syntactic representation of preference models by means of Liu's priority structures (Liu 2011) known as P-graphs. We provided conditions on P-graph transformations that enforce adherence to belief change postulates of the induced dynamic operators.

Our work can be seen as a generalisation of previous work on the integration of Belief Revision Theory and Dynamic Epistemic Logics that allows the use of the DEL framework to reason about classes of belief change operators. In some sense, this work is the complement of the characterisation of iterated belief change postulates using the proof theory of Dynamic Preference Logic (Souza, Moreira, and Vieira 2017). We point out that, since priority graphs can be seen as a syntactic form of representing evidence, our work can also be connected to the work on Evidence Logics and Explicit Knowledge (Baltag, Fiutek, and Smets 2016; Lorini 2018).

As illustrated in Section **Deriving P-Graph Transformations from Belief Revision Policies**, the characterisations proposed provide a road-map to implement belief change policies in computational systems. Besides, in Section **Negative Results** we show that a well-known iterated revision operator cannot be encoded employing graph transformations, and we indicate which policies cannot be implemented as syntactic transformations in the general case of using preference models to reason about belief change.

Acknowledgements

This study was financed in part by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES)* - Finance Code 001.

References

- 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* 50(2):510–530.
- Andréka, H.; Ryan, M.; and Schobbens, P.-Y. 2002. Operators and laws for combining preference relations. *Journal of Logic and Computation* 12(1):13–53.
- Baltag, A., and Smets, S. 2008. A qualitative theory of dynamic interactive belief revision. *Texts in logic and games* 3:9–58.
- Baltag, A.; Fiutek, V.; and Smets, S. 2016. Beliefs and evidence in justification models. In *Advances in Modal Logic*, 156–176.
- Benferhat, S.; Dubois, D.; Prade, H.; and Williams, M.-A. 2002. A practical approach to revising prioritized knowledge bases. *Studia Logica* 70(1):105–130.
- Booth, R., and Meyer, T. 2006. Admissible and restrained revision. *Journal of Artificial Intelligence Research* 26:127–11.
- Boutillier, C. 1993. Revision sequences and nested conditionals. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 519–531.
- Boutillier, C. 1994. Toward a logic for qualitative decision theory. *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning* 75–86.
- Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artificial intelligence* 89(1):1–29.
- De Jongh, D., and Liu, F. 2009. Preference, priorities and belief. In *Preference Change*, volume 42 of *Theory and Decision Library*. Springer. 85–107.
- Delgrande, J. P.; Dubois, D.; and Lang, J. 2006. Iterated revision as prioritized merging. In *Proceedings of the Tenth International Conference on the Principles of Knowledge Representation and Reasoning*, 210–220. AAAI Press.
- Fermé, E.; Garapa, M.; and Reis, M. D. L. 2017. On ensconcement and contraction. *Journal of Logic and Computation* 27(7):2011–2042.
- Fermé, E.; Krevneris, M.; and Reis, M. 2008. An axiomatic characterization of ensconcement-based contraction. *Journal of Symbolic Logic* 18(5):739–753.
- Girard, P., and Rott, H. 2014. Belief revision and dynamic logic. In *Johan van Benthem on Logic and Information Dynamics*. Springer International Publishing. 203–233.
- Girard, P. 2008. *Modal logic for belief and preference change*. Ph.D. Dissertation, Stanford University.
- Grove, A. 1988. Two modellings for theory change. *Journal of Philosophical Logic* 17(2):157–170.
- Hansson, S. O. 1992. In defense of base contraction. *Synthese* 91(3):239–245.
- Hansson, S. O. 1994. Kernel contraction. *The Journal of Symbolic Logic* 59(3):845–859.
- Jin, Y., and Thielscher, M. 2007. Iterated belief revision, revised. *Artificial Intelligence* 171(1):1–18.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial intelligence* 44(1):167–207.
- Lafage, C., and Lang, J. 2005. Propositional distances and compact preference representation. *European Journal of Operational Research* 160(3):741–761.
- Liu, F. 2011. *Reasoning about preference dynamics*, volume 354. Springer.
- Lorini, E. 2018. In praise of belief bases: Doing epistemic logic without possible worlds. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 1915–1922. AAAI Press.
- Nayak, A. C.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *Artificial Intelligence* 146(2):193–228.
- Nebel, B. 1991. Belief revision and default reasoning: Syntax-based approaches. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning*, 417–428.
- Rott, H. 1991. Two methods of constructing contractions and revisions of knowledge systems. *Journal of Philosophical Logic* 20(2):149–173.
- Rott, H. 2000. Two dogmas of belief revision. *The Journal of Philosophy* 97(9):503–522.
- Rott, H. 2009. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In Makinson, D.; Malinowski, J.; and Wansing, H., eds., *Towards Mathematical Philosophy*, volume 28 of *Trends in Logic*. Springer Netherlands. 269–296.
- Seegerberg, K. 1999. Two traditions in the logic of belief: bringing them together. In *Logic, language and reasoning*. Springer. 135–147.
- Souza, M.; Moreira, Á.; Vieira, R.; and Meyer, J.-J. C. 2016. Preference and priorities: A study based on contraction. In *Proceedings of the Fifteen International Conference on the Principles of Knowledge Representation and Reasoning*, 155–164. AAAI Press.
- Souza, M.; Moreira, Á.; and Vieira, R. 2017. Dynamic preference logic as a logic of belief change. In *Proceedings of the First International Workshop on Dynamic Logic*, 185–200. Springer.
- Souza, M. 2016. *Choices that make you change your mind: a dynamic epistemic logic approach to the semantics of BDI agent programming languages*. Ph.D. Dissertation, Institute of Informatics, Universidade Federal do Rio Grande do Sul.
- Van Benthem, J.; Grossi, D.; and Liu, F. 2014. Priority structures in deontic logic. *Theoria* 80(2):116–152.
- Van Benthem, J. 2007. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics* 17(2):129–155.
- Van Benthem, J. 2009. For better or for worse: Dynamic logics of preference. In *Preference Change*, volume 42 of *Theory and Decision Library*. Springer. 57–84.
- Williams, M.-A. 1994. On the logic of theory base change. In *Proceedings of the European Workshop on Logics in Artificial Intelligence (JELIA 94)*, 86–105. Springer.
- Williams, M.-A. 1995. Iterated theory base change: A computational model. In *Proceedings of the 14th International Joint Conference on Artificial intelligence*, 1541–1547.