Properties of Parametrized-Difference Revision

Theofanis Aravanis

Department of Business Administration
School of Economics & Business
University of Patras, Greece
taravanis@upatras.gr

Abstract

Parametrized-difference (PD) revision is a special type of rational belief revision, based on a fixed ranking over the atoms of the underlying language, with a plethora of appealing characteristics. The process of PD revision is encoded in the so-called PD operators, which essentially constitute a particular family of rational revision functions. In this article, we identify some new properties of PD revision. Specifically, we demonstrate how this type of belief change is tightly connected with selective revision, another type of revision according to which the new information is partially accepted in the revised state of belief. Furthermore, we show that PD operators respect a relevance-sensitive postulate, which introduces dependencies between revisions carried out on different (overlapping) belief sets.

1 Introduction

Belief revision (or, simply, *revision*) is the process by which a rational agent modifies her/his beliefs, in the light of new information [8, 11]. A prominent and versatile approach which formalizes belief revision is that proposed by Alchourrón, Gärdenfors and Makinson [1], known as the *AGM paradigm*. The AGM paradigm characterizes any *rational* revision operator, named *AGM revision function*, which, essentially, is a (binary) function that maps a belief set (theory) and an epistemic input (a sentence that represents new information) to a (new) revised belief set.

A family of well-behaved concrete AGM revision functions, called *parametrized-difference* (PD) operators, was recently introduced by Peppas and Williams [12, 13]. Each PD operator is *uniquely* defined by means of a *single* total preorder over the atoms of the underlying language, hence, it is compactly-specified. PD operators, also, have an embedded solution to the *iterated-revision* problem, and are expressive enough to cover a wide range of revision-scenarios, features that make them an ideal candidate for real-world implementations.

In this article, we identify some new interesting properties of PD revision. Specifically, we first demonstrate how PD revision is strongly connected with *selective revision*, a type of belief change according to which the new information is *partially accepted* in the revised belief set [7]. Already in [12, 4], it was shown that PD operators respect Parikh's *relevance-sensitive* axiom [10], an intuitive principle that supplements the AGM postulates for revision in dealing with *relevant* change. Herein, we show that PD operators, also, respect another relevance-sensitive postulate, which introduces dependencies between revisions carried out on different (overlapping) belief sets.

The article is structured as follows. The next section fixes the required formal preliminaries. Sections 3 and 4 introduce the AGM paradigm and PD revision, respectively. Thereafter, Section 5 discusses PD revision with respect to selective revision, and Section 6 points out some relevance-sensitive properties of PD revision. A brief conclusion closes the paper.

2 Formal Preliminaries

Herein, we work with a propositional language $\mathcal{L}$, built over a *finite*, non-empty set $\mathcal{P}$ of atoms (propositional variables), using the standard Boolean connectives $\wedge$ (conjunction), $\vee$ (disjunction), $\rightarrow$ (implication), $\leftrightarrow$ (equivalence), $\neg$ (negation), and governed by *classical propositional logic*. The classical consequence relation is denoted by $\models$.

A sentence φ of $\mathcal{L}$ is *contingent* iff $\not\models \varphi$ and $\not\models \neg\varphi$. For a set of sentences Γ of $\mathcal{L}$, $Cn(\Gamma)$ denotes the set of all logical consequences of Γ; i.e., $Cn(\Gamma) = \{\varphi \in \mathcal{L} : \Gamma \models \varphi\}$. An agent's belief corpus shall be modelled by a *theory*, also referred to as a *belief set*. A theory K is any deductively closed set of sentences of $\mathcal{L}$; i.e., $K = Cn(K)$. The set of all *consistent* theories is denoted by $\mathbb{K}$. For a theory K and a sentence φ of $\mathcal{L}$, we define $K + \varphi = Cn(K \cup \{\varphi\})$.

A *literal* is an atom $p \in \mathcal{P}$ or its negation. For a set of literals Q, $|Q|$ denotes the cardinality of Q. A *possible world* (or, simply, *world*) r is a consistent set of literals, such that, for any atom $p \in \mathcal{P}$, either $p \in r$ or $\neg p \in r$. The set of all possible worlds is denoted by $\mathbb{M}$. For a sentence (set of sentences) φ of $\mathcal{L}$, $[\varphi]$ is the set of worlds at which φ is true.

Let Q be a subset of $\mathcal{P}$. We denote by $\mathcal{L}^Q$ the *sublanguage* of $\mathcal{L}$ defined over Q, using the standard Boolean connectives. For a sentence x of $\mathcal{L}$, $\mathcal{P}_x$ denotes the (*unique*) *minimal* subset of $\mathcal{P}$, through which a sentence that is logically equivalent to x can be formulated. If x is inconsistent or a tautology, we take $\mathcal{P}_x$ to be the empty set. Then, we define $\mathcal{L}_x$ and $\overline{\mathcal{L}_x}$ to be the propositional (sub)languages defined over $\mathcal{P}_x$ and $\mathcal{P} - \mathcal{P}_x$, respectively, using the standard Boolean connectives. Let φ be a contingent sentence of $\mathcal{L}$. For a world $r \in \mathbb{M}$, r_φ and $r_{\overline{\varphi}}$ denote the *restrictions* of r to $\mathcal{L}_\varphi$ and $\overline{\mathcal{L}_\varphi}$, respectively; i.e., $r_\varphi = r \cap \mathcal{L}_\varphi$ and $r_{\overline{\varphi}} = r \cap \overline{\mathcal{L}_\varphi}$. For a set of worlds V, V_φ and $V_{\overline{\varphi}}$ denote the sets of (restricted) worlds resulting from the restriction of all V-worlds to $\mathcal{L}_\varphi$ and $\overline{\mathcal{L}_\varphi}$, respectively; i.e., $V_\varphi = \{r_\varphi : r \in V\}$ and $V_{\overline{\varphi}} = \{r_{\overline{\varphi}} : r \in V\}$.

A *preorder* over a set V is any reflexive, transitive binary relation in V. A preorder $\preceq$ is called *total* iff, for all $r, r' \in V$, $r \preceq r'$ or $r' \preceq r$. Also, $min(V, \preceq)$ denotes the set of all $\preceq$-minimal elements of V; i.e., $min(V, \preceq) = \{r \in V : \text{for all } r' \in V, \text{ if } r' \preceq r, \text{ then } r \preceq r'\}$.

3 The AGM Paradigm

Within the AGM paradigm, the process of belief revision is modelled as a (binary) function $*$ mapping a theory K and a sentence φ to a revised (new) theory $K * \varphi$. *Rational* revision functions, the so-called *AGM revision functions*, are those constrained by a set of eight postulates, called *AGM postulates for revision*, listed below [8, 11].

(**K** $*$ **1**) $K * \varphi$ is a theory of $\mathcal{L}$.

(**K** $*$ **2**) $\varphi \in K * \varphi$.

(**K** $*$ **3**) $K * \varphi \subseteq K + \varphi$.

(**K** $*$ **4**) If $\neg\varphi \notin K$, then $K + \varphi \subseteq K * \varphi$.

(**K** $*$ **5**) $K * \varphi$ is inconsistent iff φ is inconsistent.

(**K** $*$ **6**) If $Cn(\{\varphi\}) = Cn(\{\psi\})$, then $K * \varphi = K * \psi$.

(**K** $*$ **7**) $K * (\varphi \wedge \psi) \subseteq (K * \varphi) + \psi$.

(**K** $*$ **8**) If $\neg\psi \notin K * \varphi$, then $(K * \varphi) + \psi \subseteq K * (\varphi \wedge \psi)$.

Katsuno and Mendelzon proved that the revision functions that satisfy postulates $(K * 1)$–

$(K*8)$ are precisely those that are induced by means of a special type of total preorders over all possible worlds, called *faithful preorders* [9].

Definition 1 (Faithful Preorder, [9])**.** *A total preorder $\preceq_K$ over $\mathbb{M}$ is faithful to a theory K iff the $\preceq_K$-minimal worlds are those satisfying K; i.e., $min(\mathbb{M}, \preceq_K) = [K]$.*

Intuitively, $r \preceq_K r'$ holds when r is at least as *plausible* (relative to K) as r'.

Definition 2 (Faithful Assignment, [9])**.** *A faithful assignment is a function that maps each theory K of $\mathcal{L}$ to a total preorder $\preceq_K$ over $\mathbb{M}$, that is faithful to K.*

The following representation theorem precisely characterizes the class of AGM revision functions, in terms of faithful preorders.

Theorem 1 ([9])**.** *A revision function $*$ satisfies $(K*1)$–$(K*8)$ iff there exists a faithful assignment that maps each theory K to a total preorder $\preceq_K$ over $\mathbb{M}$, such that, for any $\varphi \in \mathcal{L}$:*

$$\textbf{(F*)} \qquad [K * \varphi] = min([\varphi], \preceq_K).$$

For ease of presentation, we shall consider, herein, only the principal case of *consistent* belief sets and *contingent* epistemic input.

4 Parametrized-Difference Revision

Peppas and Williams [12, 13], recently, introduced a *proper sub-class* of concrete AGM revision functions, well-suited for real-world implementations, called *parametrized-difference* (PD) operators. PD operators are a generalization of the *Hamming-based* Dalal's revision operator [6], as each such operator is specified by a $\trianglelefteq$-parametrization of Dalal's construction, where $\trianglelefteq$ denotes a fixed total preorder over all atoms of $\mathcal{P}$, which encodes their (prior) relative epistemic value; the more epistemic entrenched (and, thus, more resistant to change) an atom is, the higher it appears in $\trianglelefteq$. In this section, we briefly review PD operators; for details on their definition, the reader is referred to [12, 13, 4, 5].

Definition 3 (Difference between Worlds)**.** *The difference between two (possibly restricted) worlds w, r, denoted by Diff (w, r), is the set of atoms over which w and r disagree. In symbols,*

$$\mathit{Diff}\,(w, r) = \Big((w - r) \cup (r - w)\Big) \cap \mathcal{P}.$$

For a set of atoms $\mathcal{S}$ and an atom q, we define $\mathcal{S}_q = \{p \in \mathcal{S} : p \trianglelefteq q\}$. Definition 4 extends, then, the total preorder $\trianglelefteq$ to *sets* of atoms.

Definition 4 (Total Preorder over Sets of Atoms, [12])**.** *For any two sets of atoms $\mathcal{S}$, $\mathcal{S}'$, $\mathcal{S} \trianglelefteq \mathcal{S}'$ iff one of the next three conditions holds ($\triangleleft$ denotes the strict part of $\trianglelefteq$):*

- **(i)** $|\mathcal{S}| < |\mathcal{S}'|$.
- **(ii)** $|\mathcal{S}| = |\mathcal{S}'|$, *and for all* $q \in \mathcal{P}$, $|\mathcal{S}_q| = |\mathcal{S}'_q|$.
- **(iii)** $|\mathcal{S}| = |\mathcal{S}'|$, *and for some* $q \in \mathcal{P}$, $|\mathcal{S}_q| > |\mathcal{S}'_q|$, *and for all* $p \triangleleft q$, $|\mathcal{S}_p| = |\mathcal{S}'_p|$.

Definition 5 (PD Operator, [12])**.** *Let $\trianglelefteq$ be a total preorder over $\mathcal{P}$. A PD operator is the revision function induced, via condition (F$*$), from the family of PD preorders $\{\sqsubseteq_T^{\trianglelefteq}\}_{T\in\mathbb{K}}$, where each PD preorder $\sqsubseteq_K^{\trianglelefteq}$ is uniquely defined, for any $r, r' \in \mathbb{M}$, by means of condition (PD) below.*

$$\textbf{(PD)} \quad r \sqsubseteq_K^{\trianglelefteq} r' \quad \textit{iff} \quad \text{there is a } w \in [K] \text{, such that, for all } w' \in [K], \; \mathit{Diff}\,(w, r) \trianglelefteq \mathit{Diff}\,(w', r').$$

Definition 5 implies that there is a *one-to-one correspondence* between the total preorders over atoms and the PD operators. Note, lastly, that, when $\trianglelefteq = \mathcal{P} \times \mathcal{P}$ (i.e., all atoms have *equal* epistemic value), the PD preorder $\sqsubseteq_K^{\trianglelefteq}$ is defined, for any $r, r' \in \mathbb{M}$, as follows: $r \sqsubseteq_K^{\trianglelefteq} r'$ iff there is a $w \in [K]$, such that, for all $w' \in [K]$, $|\mathit{Diff}\,(w, r)| \leqslant |\mathit{Diff}\,(w', r')|$. In this case, the family $\{\sqsubseteq_K^{\trianglelefteq}\}_{K\in\mathbb{T}}$ produces Dalal's operator [6].

5 PD and Selective Revision

In this section, we demonstrate how a total preorder $\trianglelefteq$ over the atoms of the language —which, essentially, constitutes the generative unit of PD revision— can be utilized for implementing *selective revision*, a special type of belief revision according to which *only a part* of an epistemic input is accepted in the revised state of belief [7]. Recall that, in the standard AGM paradigm, the new information is *always* accepted (due to postulate $(K * 2)$). This, however, is a rather unrealistic assumption, since real-world rational agents do not always receive information from reliable sources. The following scenario, borrowed from [7], is illustrative.

Example 1 ([7])**.** *You return back from work and your son tells you, as soon as you see him: "A dinosaur has broken grandma's vase in the living-room". You, probably, accept the information that the vase has been broken, and reject the part of the information that refers to the dinosaur.*

A plausible way for such information filtering would be by taking into account the relative plausibility of the "building blocks" of sentences of the language. As these "building blocks", essentially, are the atoms of the language, and the relative plausibility of the atoms is encoded in a total preorder $\trianglelefteq$, it turns out that PD revision provides a means for *information filtering.* To see this, suppose that, during revision, the following *filtering-rule* is applied:

> *"If the epistemic input is a conjunction of atoms, then only the strictly most $\trianglelefteq$-plausible atoms should be accepted in the revised belief set".*

On that premise, if $a \wedge b$ is an epistemic input (where a, b are atoms), and, moreover, we have that $a \triangleleft b$, then only the atom b should be accepted in the revised state of belief.

6 Relevance-Sensitive Properties of PD Revision

Parikh pointed out that the AGM postulates for revision are liberal in their treatment of *relevance* [10]. To remedy this weakness, he proposed an additional axiom that supplements postulates $(K*1)$–$(K*8)$, named axiom (P), according to which the revision of a theory K that can be divided in two *syntax-disjoint* compartments by an epistemic input φ that is syntax-related only to the first compartment of K should *not affect* the second compartment of K. In a subsequent work [14], two interpretations of Parikh's axiom were identified, namely, its *weak* and *strong* version.[1] Already in [12], it was shown that PD operators respect the weak version

[1] The *semantic* properties of Parikh's axiom were investigated, in detail, in [2, 3].

of axiom (P), whereas, in [4], it was shown that PD operators, also, respect the strong version of (P).

In this section, we point out further interesting *relevance-sensitive* properties of PD revision. For establishing our results, we shall use the already known fact that PD operators respect the following postulate (R), which states that the $\overline{\mathcal{L}_\varphi}$-part of the revised theory $K * \varphi$ contains *at least* every sentence of the $\overline{\mathcal{L}_\varphi}$-part of the initial theory K.

$$\textbf{(R)} \quad K \cap \overline{\mathcal{L}_\varphi} \subseteq (K * \varphi) \cap \overline{\mathcal{L}_\varphi}.$$

Postulate (R) —which is equivalent to $[K * \varphi]_{\overline{\varphi}} \subseteq [K]_{\overline{\varphi}}$ in the realm of possible worlds— implies some interesting properties of PD revision, encoded in Lemma 1. To present this lemma, let us first introduce the required notation. For an arbitrary theory K and a sentence φ, such that $\mathcal{L}_\varphi \subset \mathcal{L}$, we denote by $[\varphi]^K$ the set of φ-worlds whose $\overline{\mathcal{L}_\varphi}$-part agrees with the $\overline{\mathcal{L}_\varphi}$-part of some K-world; i.e., $[\varphi]^K = \Big\{ r \in [\varphi] : r_{\overline{\varphi}} \in [K]_{\overline{\varphi}} \Big\}$. By definition, it holds that $[\varphi]^K \subseteq [\varphi]$.

Lemma 1. *Let $\trianglelefteq$ be a total preorder over atoms, and let $*$ be a PD operator induced from the family of PD preorders $\{\sqsubseteq^{\trianglelefteq}_T\}_{T \in \mathbb{K}}$. Moreover, let K be a theory, and let φ be a sentence, such that $\mathcal{L}_\varphi \subset \mathcal{L}$. Then, the following identity is true:*

$$[K * \varphi] = min\big([\varphi]^K, \sqsubseteq^{\trianglelefteq}_{K_{IC}}\big)$$

$$= \Big\{ u \in [\varphi]^K : \exists\ w' \in [K] \ \textit{s.t.}\ \textit{Diff}\,(w', u) \in min\Big(\big\{ \textit{Diff}\,(w, r) : w \in [K] \ \textit{and}\ r \in [\varphi]^K \big\}, \trianglelefteq \Big) \Big\}$$

$$= \Big\{ u \in [\varphi]^K : \exists\ w' \in [K] \ \textit{s.t.}\ \textit{Diff}\,(w', u) \in min\Big(\big\{ \textit{Diff}\,(w_\varphi, r_\varphi) : w \in [K] \ \textit{and}\ r \in [\varphi]^K \big\}, \trianglelefteq \Big) \Big\}$$

Proof. The first equality follows directly from condition (F$*$), from which we have that $[K * \varphi] = min([\varphi], \sqsubseteq^{\trianglelefteq}_K)$, and postulate (R), which entails that $[K * \varphi] \subseteq [\varphi]^K \subseteq [\varphi]$. The second equality follows from condition (PD). The last equality follows from the fact that $\big\{ w_{\overline{\varphi}} : w \in [K] \big\} = \big\{ r_{\overline{\varphi}} : r \in [\varphi]^K \big\}$, which is implied by the definition of the set of worlds $[\varphi]^K$. ■

Lemma 1, essentially, says that the specification of the $\sqsubseteq^{\trianglelefteq}_K$-minimal φ-worlds, through the differences between worlds of Definition 3, does *not* require the $\overline{\mathcal{L}_\varphi}$-part of the involved worlds.

Against this background, we will show that PD operators, also, respect the following *relevance-sensitive* postulate (C).

$$\textbf{(C)} \quad \text{If } K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi, \text{ then } (K * \varphi) \cap \mathcal{L}_\varphi = (H * \varphi) \cap \mathcal{L}_\varphi.$$

Postulate (C) makes an association between the revision policies of two *different* (overlapping) theories. In particular, it states that, if two theories K and H share the same $\mathcal{L}_\varphi$-part, then the revised theories $K * \varphi$ and $H * \varphi$ should, also, share the same $\mathcal{L}_\varphi$-part. Therefore, any beliefs of K, H that are outside the sublanguage $\mathcal{L}_\varphi$ *do not affect* the way that the $\mathcal{L}_\varphi$-parts of K, H are modified — stated otherwise, the *context* of the $\mathcal{L}_\varphi$-parts of K, H does not affect the modification of the $\mathcal{L}_\varphi$-parts of K, H themselves.

It is noteworthy that, since $K \cap \mathcal{L}_\varphi = Cn\big(K \cap \mathcal{L}_\varphi\big) \cap \mathcal{L}_\varphi$, the following identity is, immediately, derived from postulate (C):

$$(K * \varphi) \cap \mathcal{L}_\varphi = \Big(Cn\big(K \cap \mathcal{L}_\varphi\big) * \varphi \Big) \cap \mathcal{L}_\varphi.$$

Theorem 2 proves that PD operators respect postulate (C).

Theorem 2. *PD operators satisfy postulate (C).*

Proof. Let $\trianglelefteq$ be a total preorder over atoms, and let $*$ be a PD operator induced from the family of PD preorders $\{\sqsubseteq_T^{\trianglelefteq}\}_{T \in \mathbb{K}}$. Moreover, let K, H be two theories, and φ be a sentence of $\mathcal{L}$, such that $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$. If φ is consistent with both K and H, or $\mathcal{L}_\varphi = \mathcal{L}$, then (C) trivially holds. Assume, therefore, that φ contradicts K, H, and $\mathcal{L}_\varphi \subset \mathcal{L}$. Given that $[K]_\varphi = [H]_\varphi$ (as $K \cap \mathcal{L}_\varphi = H \cap \mathcal{L}_\varphi$), $[\varphi]^K_\varphi = [\varphi]^H_\varphi$, $\{w_{\overline{\varphi}} : w \in [K]\} = \{r_{\overline{\varphi}} : r \in [\varphi]^K\}$ and $\{w_{\overline{\varphi}} : w \in [H]\} = \{r_{\overline{\varphi}} : r \in [\varphi]^H\}$, we derive that $min\Big(\big\{ Diff\,(w_\varphi, r_\varphi) : w \in [K]\,,\, r \in [\varphi]^K \big\}, \trianglelefteq\Big) = min\Big(\big\{ Diff\,(w_\varphi, r_\varphi) : w \in [H]\,,\, r \in [\varphi]^H \big\}, \trianglelefteq\Big)$. Then, it is not hard to verify that Lemma 1 entails $[K * \varphi]_\varphi = [H * \varphi]_\varphi$; thus, $(K * \varphi) \cap \mathcal{L}_\varphi = (H * \varphi) \cap \mathcal{L}_\varphi$, as desired. ■

7 Conclusion

Parametrized-difference (PD) revision constitutes a well-behaved type of belief revision, which is perfectly-suited for real-world implementations. In this work, we identified some new interesting properties of this type of revision. Specifically, we demonstrated how PD revision is tightly connected with selective revision. Furthermore, we pointed out that PD operators respect a relevance-sensitive postulate, which introduces dependencies between revisions carried out on different (overlapping) belief sets.

Acknowledgments

This research is co-financed by Greece and the European Union (European Social Fund – ESF) through the Operational Programme "Human Resources Development, Education and Lifelong Learning" in the context of the project "Reinforcement of Postdoctoral Researchers – 2nd Cycle" (MIS-5033021), implemented by the State Scholarships Foundation (IKY).

I am grateful to Pavlos Peppas and Yannis Stamatiou for many fruitful discussions on this work, as well as to the anonymous referees for their valuable input.

References

[1] Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.

[2] Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams. Full characterization of Parikh's relevance-sensitive axiom for belief revision. *Journal of Artificial Intelligence Research*, 66:765–792, 2019.

[3] Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams. A study of possible-worlds semantics of relevance-sensitive belief revision. *Journal of Logic and Computation*, 30:1125–1142, 2020.

[4] Theofanis Aravanis, Pavlos Peppas, and Mary-Anne Williams. An investigation of parametrized difference revision operators. *Annals of Mathematics and Artificial Intelligence*, 89:7–28, 2021.

[5] Theofanis I. Aravanis. Strengthening parametrized-difference belief revision. *Bulletin of the Hellenic Mathematical Society*, 64:69–91, 2020.

[6] Mukesh Dalal. Investigations into theory of knowledge base revision: Preliminary report. In *Proceedings of 7th National Conference of the American Association for Artificial Intelligence (AAAI 1988)*, pages 475–479, 1988.

[7] Eduardo Fermé and Sven O. Hansson. Selective revision. *Studia Logica*, 63:331–342, 1999.

[8] Peter Gärdenfors. *Knowledge in Flux – Modeling the Dynamics of Epistemic States.* MIT Press, Cambridge, Massachusetts, 1988.

[9] Hirofumi Katsuno and Alberto Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

[10] Rohit Parikh. Beliefs, belief revision, and splitting languages. In Lawrence S. Moss, Jonathan Ginzburg, and Maarten de Rijke, editors, *Logic, Language and Computation*, volume 2, pages 266–278. CSLI Publications, 1999.

[11] Pavlos Peppas. Belief revision. In Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, 2008.

[12] Pavlos Peppas and Mary-Anne Williams. Kinetic consistency and relevance in belief revision. In *Proceedings of the 15th European Conference on Logics in Artificial Intelligence (JELIA 2016)*, pages 401–414. Springer International Publishing, 2016.

[13] Pavlos Peppas and Mary-Anne Williams. Parametrised difference revision. In *Proceedings of the 16th International Conference on Principles of Knowledge Representation and Reasoning (KR 2018)*, pages 277–286, 2018.

[14] Pavlos Peppas, Mary-Anne Williams, Samir Chopra, and Norman Foo. Relevance in belief revision. *Artificial Intelligence*, 229:126–138, December 2015.