

Parametrised Difference Revision

Pavlos Peppas^{1,2}

¹Dept of Business Administration
University of Patras
Greece

Mary-Anne Williams²

²Centre for Artificial Intelligence
School of Software, FEIT
University of Technology Sydney
Australia

Abstract

Despite the great theoretical advancements in the area of Belief Revision, there has been limited success in terms of implementations. One of the hurdles in implementing revision operators is that their specification (let alone their computation), requires substantial resources. On the other hand, implementing a specific revision operator, like Dalal's operator, would be of limited use. In a recent paper we generalised Dalal's construction defining a whole family of concrete revision operators, called *Parametrised Difference revision operators* or *PD operators* for short. This family is wide enough to cover a whole range of different applications, and at the same time it is easy to represent. In this paper we characterise axiomatically the family of PD operators, study its computational complexity, and discuss its benefits for belief revision implementations.

Introduction

The AGM framework (Alchourron *et al.* 1985), is the dominant paradigm for the study of *belief revision*. It has been studied extensively and lies on solid theoretical foundations (see (Peppas 2008) for a survey). Yet despite the success of its theoretical models, little has been done in terms of implementations of AGM belief revision operators. This is not to say that important attempts have not been made; see for example (Chou and Winslett 1991), (Williams and Sims 2000), (Beierle and Kern-Isberner 2008). None of them however has had the great impact on real-world applications that one would expect from a successful implementation of belief revision.

There are at least two major obstacles to a successful implementation of an AGM belief revision system, that can work beyond toy example. The first is the high computational complexity of the belief revision process (Eiter and Gottlob 1992); we will have more to say about this later in the paper.

The second is the large amount of information that, in principle, the user needs to provide to the system. Recall that the AGM postulates for revision specify, not one, but an entire class of revision functions. Hence, before a belief revision system can answer any queries about the result of revising a theory K by a sentence φ , denoted $K * \varphi$, the user

needs to specify the particular revision function $*$ she is interested in. There are many ways that this can be done, but in principle, they are all equivalent to specifying a family of total preorders over possible worlds; i.e., one total preorder for each theory of the object language L ((Katsuno and Mendelson 1991)). For a propositional language with n variables, there exist 2^{2^n} theories; clearly an enormous number. Even if one focuses only on a single theory K , one still needs to specify a preorder \leq over the 2^n possible worlds.

Of course there are shortcuts. For example one can request only a *partial* specification of a preorder over worlds, and fill in the remaining information automatically using some (intuitive) default rule. In this case it is important that the side-effects of the completion process are well understood, and that the formal properties of the resulting revision functions are thoroughly investigated. The other option is to avoid the requirement for preorder specification altogether, by choosing to implement only one *concrete* revision operator. The problem of course is that such a system would be rather limited in scope. Moreover, as far as concrete "off-the-shelf" AGM revision operators go, there are not all that many to choose from. Out of the few well known proposals, like (Borgida 1985), (Winslett 1988), (Satoh 1988), (Weber 1986), it is only Dalal's operator, (Dalal 1988), that satisfies all the AGM postulates for revision.

In (Peppas and Williams 2016) we introduced an entire class of concrete revision operators, all of which satisfy the full set of AGM postulates for revision; they are called *Parametrised Difference revision operators*, or *PD operators* for short. PD operators are essentially generalisations of Dalal's operator. Most importantly, each PD operator can be fully specified from a preorder over the n *propositional variables* of the object language L . In other words, a *single* preorder over the n propositional variables, suffices to generate the preorders over possible worlds associated with all 2^{2^n} theories of L . This is a *double exponential* drop on the information required from the user. Moreover in (Peppas and Williams 2016), PD operators were shown to be expressive enough to cover a wide range of belief revision scenarios, including ones on *iterated revision* (Peppas 2014).

In this paper we investigate the formal properties of PD operators. No such investigation was carried out in (Peppas and Williams 2016) as the focus in that paper was on *kinetic consistency* and *relevance*; PD operators were essentially a

by-product of that study. In particular, PD operators were defined only constructively, and were introduced mainly as a means to validate the postulates for *kinetic consistency* and *relevance*.

Herein we provide an *axiomatic characterisation* of the class of PD operators. Moreover we study the computational complexity of PD operators and show that, although more expressive than Dalal's operator, they lie at the same level of the polynomial hierarchy. Perhaps more importantly for practical applications, when confined to Horn knowledge bases, and the size of the queries are bounded by a constant, the complexity of PD operators drops to *linear time* with respect to the size of the knowledge base.

The rest of the paper is structured as follows. The next section introduces some notation and terminology, followed by a section covering the necessary background on AGM belief revision and PD operators. Then, new axioms characterising PD operators are formulated, accompanied by corresponding representation results. The next section discusses the representational cost of AGM revision, and compares the effectiveness of PD operators with previous approaches in dealing with this problem. This is followed by a section that contains our study on the computational complexity of PD operators. In the last section we provide some concluding remarks.

Preliminaries

In this article we shall be working with a propositional language L built over *finitely many* propositional variables. The finite, nonempty set of all propositional variables (also called *atoms*) is denoted by P . A *literal* is a variable in P or the negation of a variable. If l is a literal containing the variable α , then by \bar{l} we denote the literal $\neg\alpha$ if $l = \alpha$, and the literal α otherwise. The letters x, y, p, q , and z (possibly with subscripts and/or superscripts) will always represent literals. The letters A, B, C, D , and E (possible with subscripts and/or superscripts) will always represent sets of literals. For a set of literals A , we define \bar{A} to be the set $\bar{A} = \{\bar{q} : q \in A\}$. We will sometimes abuse notation and treat a set of literals A as a sentence, namely the conjunction of all its literals, leaving it to the context to resolve any ambiguity; thus for example, in " $A \subseteq B$ ", A is a set of literals whereas in " $\neg A$ ", A is a sentence.

The set of all interpretations over P is denoted \mathcal{M} . Interpretations will also be called possible worlds and are identified with the set of literals they satisfy. Like with any set of literals, a possible world r may sometimes be treated as a sentence (as for example in the expression " $r \vee r'$ "). Possible worlds, or simply worlds, will be denoted with the letters w, r, u , possibly with subscripts and/or superscripts.

Arbitrary sentences of L will be denoted by the Greek letters φ, ψ (possibly with subscripts and/or superscripts). For a set of sentences Γ of L , by $Cn(\Gamma)$ we denote the set of all logical consequences of Γ , i.e., $Cn(\Gamma) = \{\varphi \in L : \Gamma \models \varphi\}$. A theory K of L is any set of sentences of L closed under \models , i.e., $K = Cn(K)$. We shall use the letters K, H , and T to denote theories of L . The set of all consistent theories is denoted by \mathcal{K} . A theory K is complete iff for all sentences

$\varphi \in L$, $\varphi \in K$ or $\neg\varphi \in K$.

For a set of sentences Γ of L , $[\Gamma]$ denotes the set of all possible worlds that satisfy Γ . Often we shall use the notation $[\varphi]$ for a sentence $\varphi \in L$, as an abbreviation of $[\{\varphi\}]$. For a theory K and a set of sentences Γ of L , we shall denote by $K + \Gamma$ the closure under \models of $K \cup \Gamma$, i.e., $K + \Gamma = Cn(K \cup \Gamma)$. For a sentence $\varphi \in L$ we shall often write $K + \varphi$ as an abbreviation of $K + \{\varphi\}$.¹

The AGM Framework

In the AGM framework, belief revision is modelled as a function $*$ mapping a theory K and a sentence φ , to a theory $K * \varphi$, representing the result of revising K by φ . Alchourrón, Gärdenfors and Makinson have introduced a set of eight postulates, numbered $(K * 1) - (K * 8)$, that ought to be satisfied by any *rational* revision function. These postulates are now known as the *AGM postulates for revision*, and the functions that satisfy these postulates are known as *AGM revision functions* (or simply *revision functions*).²

It turns out that any AGM revision function can be constructed with the use of a set of total preorders over possible worlds; one total preorder \leq_K for each theory K . Recall that a total preorder \leq_K over \mathcal{M} is any binary relation in \mathcal{M} that is reflexive and transitive, and such that for all $w, w' \in \mathcal{M}$, $w \leq_K w'$ or $w' \leq_K w$. As usual, $<_K$ denotes the strict part of \leq_K . Moreover, we shall write $w \approx_K w'$ iff $w \leq_K w'$ and $w' \leq_K w$.

A total preorder \leq_K is said to be *faithful to K* iff for all $w, w' \in [\mathcal{M}]$, (i) if $w \in [K]$ then $w \leq_K w'$, and, (ii) if $w \in [K]$ and $w' \notin [K]$ then $w <_K w'$.

Given a faithful preorder \leq_K for each theory K , one can construct a revision function $*$ by means of the following condition, (Katsuno and Mendelzon 1991):³

$$(\leq^*) \quad [K * \varphi] = \min([\varphi], \leq_K).$$

In the above definition, $\min(S, \leq_K)$ is the set of *minimal* elements of the set S with respect to \leq_K ; i.e., $\min(S, \leq_K) = \{w \in S : \text{for all } w' \in S, \text{ if } w' \leq_K w, \text{ then } w \leq_K w'\}$. Hence according to (\leq^*) , $K * \varphi$ is defined as the theory satisfied precisely by the \leq_K -minimal worlds in $[\varphi]$.

Katsuno and Mendelzon have shown that the functions induced from faithful preorders via (\leq^*) are exactly those satisfying the AGM postulates for revision.

For ease of presentation, in the rest of the paper we shall focus only on revision of *consistent* theories by *consistent* sentences. Hence from now on, unless explicitly stated otherwise, we assume that the initial belief set K is a consistent theory, and that the epistemic input φ is a consistent sentence.

¹In this paper we essentially follow the notation and terminology that is typically used in the Belief Revision literature; some parts of it are also borrowed from (Peppas and Williams 2016).

²Due to space limitations the AGM postulates have been omitted. See (Gärdenfors 1988) or (Peppas 2008) for details.

³We note that in fact Katsuno and Mendelzon assign faithful preorders to *sentences* rather than to *theories*. However this difference is inconsequential under the assumptions of this paper.

Parametrised Difference Operators

(Dalal 1988) provides a very natural way of defining the preorder \leq_K associated to a theory K . We note that \leq_K is meant to encode the comparative plausibility of possible worlds: the closer a world is to the beginning of the preorder the more plausible it is. Dalal defines plausibility in terms of a notion of *difference* between worlds.

In particular, for any two worlds $w, r \in \mathcal{M}$, the difference between w and r , denoted $Diff(w, r)$, is defined to be the set of propositional variables over which the two worlds disagree; i.e., $Diff(w, r) = \{q \in P : w \models q \text{ and } r \models \bar{q}\} \cup \{q \in P : r \models q \text{ and } w \models \bar{q}\}$. The preorder \sqsubseteq_K that Dalal assigns to a consistent theory K is defined as follows: for all $r, r' \in \mathcal{M}$, $r \sqsubseteq_K r'$ iff there is a $w \in [K]$ such that for all $w' \in [K]$, $|Diff(w, r)| \leq |Diff(w', r')|$. Dalal's operator, which we denote \square , is defined as the revision function induced from $\{\sqsubseteq_K\}_{K \in \mathcal{K}}$.

An example of Dalal's preorder for a language L built from only three variables a, b, c , assigned to the theory $K = Cn(\{a, b, c\})$, is given below:

$$abc \sqsubseteq_K \begin{array}{c} a\bar{b}\bar{c} \\ \bar{a}bc \end{array} \sqsubseteq_K \begin{array}{c} \bar{a}\bar{b}\bar{c} \\ \bar{a}b\bar{c} \end{array} \sqsubseteq_K \bar{a}\bar{b}c$$

In the above example, the plausibility of a world r is determined by the number of propositional variables over which r differs from the initial world abc . As noted in (Peppas and Williams 2016), "a silent assumption in Dalal's approach is that *all variables have the same epistemic value*; hence for example, a change in the variable a is assumed to be as plausible (or implausible) as a change in variable b ." This is clearly a severe restriction that limits considerably the range of applicability of Dalal's operator. Hence, a generalisation of Dalal's approach was considered in (Peppas and Williams 2016), where propositional variables are allowed to have different epistemic values.

Suppose for example that for a certain application, the atoms a and b have greater epistemic value than the atom c , and consequently a change in a or b is less plausible than a change in c . This can be encoded by a total preorder \leq over the variables a, b, c as follows: $c \triangleleft a, c \triangleleft b, a \leq b$, and $b \leq a$.⁴ Given \leq we can *refine* Dalal's preorder to take into account the difference in epistemic value between a, b , and c :

$$abc \sqsubseteq_K^{\leq} a\bar{b}\bar{c} \sqsubseteq_K^{\leq} \begin{array}{c} \bar{a}bc \\ \bar{a}\bar{b}c \end{array} \sqsubseteq_K^{\leq} \begin{array}{c} \bar{a}\bar{b}\bar{c} \\ \bar{a}b\bar{c} \end{array} \sqsubseteq_K^{\leq} \bar{a}\bar{b}c \sqsubseteq_K^{\leq} \bar{a}\bar{b}c$$

In the example above the ranking of possible worlds takes place in two stages. The first stage is identical to Dalal's ranking: each world r is ranked according to the number of switches in propositional variables that are necessary to turn the initial world abc into r . At the second stage the ranking is further refined to take into account the different epistemic value of the propositional variables that have been

⁴As usual \triangleleft denotes the strict part of \leq ; moreover the obvious relationships $a \leq a, b \leq b$, and $c \leq c$ have been omitted for the sake of readability.

switched. In particular, for any two worlds r, r' that require the same number of switches from abc (i.e., $|Diff(abc, r)| = |Diff(abc, r')|$), r is more plausible than r' iff $Diff(abc, r)$ *lexicographically precedes* $Diff(abc, r')$ with respect to \leq . Thus for example, $a\bar{b}\bar{c} \sqsubseteq_K^{\leq} \bar{a}bc$ because $c \triangleleft a$ (despite the fact that both worlds are one switch away from abc).

The example above illustrates the basic idea in generalising Dalal's approach. The formal definition of PD preorders, presented in (Peppas and Williams 2016), is given below.

Let \leq be any total preorder over P (the set of propositional variables). For a set $S \subseteq P$ and a variable $q \in P$, by S_q we denote the set $S_q = \{p \in S : p \leq q\}$. The definition of \leq can now be extended to *sets* of propositional variables. In particular, for any two sets $S, S' \subseteq P$, $S \leq S'$ iff one of the following three conditions holds:

- $|S| < |S'|$.
- $|S| = |S'|$, and for all $q \in P$, $|S_q| = |S'_q|$.
- $|S| = |S'|$, and for some $q \in P$, $|S_q| > |S'_q|$, and for all $p \triangleleft q$, $|S_p| = |S'_p|$.

In the above definition, condition (b) states that S and S' are lexicographically indistinguishable with respect to \leq , whereas (c) states that S lexicographically precedes S' (wrt \leq). It is not hard to verify that (the extended) \leq is a total preorder over 2^P .

The intended reading of \leq , defined over sets of variables, is the same as before: $S \leq S'$ means that S' as a whole is at least as important than S (as a whole). Therefore, if we had to choose between changing all variables in S or changing all variables in S' , we will pick the former.

Based on this reading, the *PD preorder* \sqsubseteq_K^{\leq} over \mathcal{M} , induced from \leq at a theory K , is defined as follows: $r \sqsubseteq_K^{\leq} r'$ iff there is a $w \in [K]$ such that for all $w' \in [K]$, $Diff(w, r) \leq Diff(w', r')$.

From the results in (Peppas and Williams 2016), it follows immediately that \sqsubseteq_K^{\leq} is a total preorder which moreover is faithful to K .

Notice that according to this definition, a *single* preorder \leq over P suffices to determine the preorders assigned to *all* consistent theories K . Hence a preorder \leq generates a family of PD preorders $\{\sqsubseteq_K^{\leq}\}_{K \in \mathcal{K}}$ which in turn define a revision function $*$. A revision function so constructed is called a *Parametrised Difference revision operator* or a *PD operator* for short.

To illustrate the use of PD operators in encoding belief revision scenarios, consider the following example from (Peppas and Williams 2016):⁵ a circuit consists of two adders and one multiplier. The variables a_1, a_2 , and m represent the facts that "adder1 is working", "adder2 is working", and "the multiplier is working" respectively. Initially we believe that the circuit is working properly. Moreover we know that multipliers are less reliable than adders. Hence, if we observe that there is a malfunction in the circuit, it is plausible to assume that the multiplier (rather than one of the adders) is not working properly.

⁵Which in turn is modified version of an earlier example in (Darwiche and Pearl 1997).

This scenario can easily be encoded with a PD operator. In particular, consider the PD operator $*$ induced from the following preorder \leq on the propositional variables a_1, a_2, m : $m < a_1, m < a_2, a_1 \leq a_2$, and $a_2 \leq a_1$ (in addition, \leq includes all pairs that follow from reflexivity and transitivity). It is not hard to verify that with this preorder, the revision of $Cn(\{a_1, a_2, m\})$ by $\neg a_1 \vee \neg a_2 \vee \neg m$ leads us to $Cn(\{a_1, a_2, \neg m\})$ as desired.⁶

There are more indicative examples in (Peppas and Williams 2016), including some on iterated revision, that can be readily encoded with PD operators. These examples suggest that PD operators have a wide range of applicability, which combined with their low representational cost, make PD operators an important class of AGM revision functions.

It should be noted that the idea of generalising Dalal's notion of distance between worlds, by differentiating between atoms has been used in the *Belief Merging* literature for quite some time. In particular, preorders on a *weighted Hamming distance* are quite similar to PD preorders. A weighted Hamming distance assigns a *numerical* value (i.e., a *weight*) to each variable of the language. The distance between two possible worlds is then defined as the sum of the weights of all variables over which the two worlds differ (see for example, (Konieczny *et al.* 2004)). These numerical weights assigned to variables can be thought of as the quantitative analog of the preorder \leq over variables used in the construction of a PD preorder. There is a major difference however between preorders induced from weighted Hamming distances and PD preorders: with the former it is possible for three worlds w, r, r' , to be such that r' is closer to w than r , even though r differs from w in fewer variables than r' (i.e., $|Diff(w, r)| < |Diff(w, r')|$);⁷ this can never be the case with PD preorders.⁸

Axioms for PD Operators

PD operators were only defined constructively in (Peppas and Williams 2016). In this section we introduce eight new axioms which together with the original AGM postulates, characterise precisely the family of PD operators.

We note from the outset that our new axioms are *not* on a par with the AGM postulates. In fact the two have a totally different purpose. AGM postulates encode general principles of rational belief change. Our new axioms on the other hand, are simply formal properties that characterise a certain class of AGM revision functions (namely those induced from PD preorders), thus providing insight to their behaviour.

Formulating the new axioms was not trivial. The task is complicated by the fact that a PD operator $*$ is constructed

⁶On the other hand, Dalal spreads the blame equally to all three components of the circuit; i.e., the Dalal-revision of $Cn(\{a_1, a_2, m\})$ by $\neg a_1 \vee \neg a_2 \vee \neg m$ yields the theory $Cn(\bar{a}_1 a_2 m \vee a_1 \bar{a}_2 m \vee a_1 a_2 \bar{m})$.

⁷This can happen for example if one of the variables in $Diff(w, r)$ has a weight that is greater than the sum of all weights in $Diff(w, r')$.

⁸We thank the anonymous reviewer for pointing out previous work on weighted Hamming distances and the similarity of their induced preorders to PD preorders.

from a preorder \leq over P in two stages; first \leq induces $\{\subseteq_K^{\leq}\}_{K \in \mathcal{K}}$, which in turn induces $*$. Thus, metaphorically speaking, $*$ is *two steps away* from its generator \leq . That makes it harder to devise constraints on $*$ that would project correctly, at a two-steps distance, to \leq .

For the sake of readability we shall introduce the new axioms in four stages. At each stage we provide representation results that highlight the role of the new axioms in the overall characterisation of PD operators.

We recall that throughout this paper, x, y, p, q, z denote literals, A, B, C, D, E denote nonempty consistent sets of literals, φ, ψ denote consistent sentences, and K, H, T denote consistent theories. Moreover, we shall often use concatenation as an abbreviation for conjunction; thus for example AB is an abbreviation of $A \wedge B$, and Ap is an abbreviation of $A \wedge p$.

For nonempty sets of literals A, B , we define $A \leq_K B$ iff $A, B \subseteq K$ and $\neg(\bar{A}) \notin K * (\bar{A} \vee \bar{B})$. Intuitively, $A \leq_K B$ holds whenever, starting from the belief set K (which contains both A and B), it is at least as costly to change (the values of) all literals in B than it is to change all literals in A . We define $A <_K B$ as $A \leq_K B$ and $B \not\leq_K A$ (or equivalently, $A, B \subseteq K$ and $\neg\bar{B} \in K * (\bar{A} \vee \bar{B})$). Finally, for literals p, q , we define $p \leq_K q$ and $p <_K q$ to be abbreviations of $\{p\} \leq_K \{q\}$, and $\{p\} <_K \{q\}$ respectively.

The Special Case of Consistent Complete Theories

Let us start by assuming that the initial belief set K is a consistent *complete* theory. This assumption will allow us to arrive quickly at preliminary representation results that will be instrumental in establishing the general results of the next subsection. Most proofs of the results reported in this paper are omitted due to space limitations; however all missing proofs can be found at the web page of the first author.

Our first axiom says that if one needs to reverse all literals in A or all literals in B , then revision never picks the larger set; in other words, the more literals one needs to reverse during revision, the more costly it is:

$$(D1) \text{ If } A \leq_K B, \text{ then } |A| \leq |B|.$$

(D1) alone suffices to characterise an interesting superclass of PD operators. In particular, consider the following constraint on a faithful preorder \leq_K assigned to K :⁹

$$(H) \text{ If } |Diff(K, r)| < |Diff(K, r')|, \text{ then } r <_K r'.$$

We shall call a preorder \leq_K over \mathcal{M} satisfying (H), a *Hamming preorder*.¹⁰

Theorem 1 *Let K be a consistent complete theory, $*$ an AGM revision function and \leq_K the faithful preorder that $*$ assigns to K . Then \leq_K is a Hamming preorder iff $*$ satisfies (D1) at K .*

⁹Recall that in this subsection, K is assumed to be a consistent complete theory.

¹⁰We note that Hamming preorders are similar, but not quite the same as the preorders induced from weighted Hamming distances discussed earlier.

All PD preorders are Hamming preorders, but not the other way around. Let us take a closer look at the difference between the two.

Given the initial world K , all remaining worlds can be partitioned according to the number of atoms in which they differ from K . In both PD and Hamming preorders, the worlds that differ from K in one atom, precede those that differ from K in two atoms, which precede those that differ from K in three atoms, etc. On the other hand, the relative order of the worlds that belong to the *same* partition is quite different in Hamming and PD preorders: in Hamming preorders the ordering with a partition is arbitrary, whereas in PD preorders it is highly regulated. More precisely, in a PD preorder, the way that the worlds in the first partition are ordered, *fully determines* the ordering of the worlds in all subsequent partitions. In other words, if two PD preorders, faithful to K , agree on the ordering of worlds that differ from K on one atom, then the two preorders are identical. This observation has been the basis for formulating the extra axioms required for PD revisions:

(D2) If $A \preceq_K B$, $p \preceq_K q$, and $q \notin B$, then $Ap \preceq_K Bq$.

Axiom (D2) essentially says that if switching the literals in A is at least as easy as switching the literals in B , and switching p is at least as easy as switching q , then switching A and p together is at least as easy as switching B and q together (provided that q is not already in B).

(D3) If $A \preceq_K B$, $p <_K q$, and $q \notin B$, then $Ap <_K Bq$.

Axiom (D3) is essentially the strict version of (D2). Like in (D2), we assume that reversing A is at least as easy as reversing B , but this time we assume that reversing p is strictly easier than reversing q . In this case, says (D2), reversing A and p together is strictly easier than reversing B and q together (provided that $q \notin B$).

(D4) If $A <_K B$, $p \in K$, $q \notin B$, and for all $z \in B$, $z \preceq_K q$, then $Ap <_K Bq$.

Axiom (D4) is based on a similar intuition as (D2) and (D3), but deals with a different case. Suppose that reversing A is strictly easier than reversing B . Moreover assume that reversing the literal q is at least as hard as reversing any literal z in B . Then, says (D4), for *any* literal $p \in K$, changing A and p together is strictly easier than changing B and q together (provided that q is not already in B).

Theorem 2 *Let K be a consistent complete theory, $*$ an AGM revision function and \preceq_K the faithful preorder that $*$ assigns to K . If \preceq_K is a PD preorder then $*$ satisfies (D1) - (D4) at K .*

Proof. Assume that \preceq_K is a PD preorder. Then there exists a preorder \preceq over P , such that the preorder \preceq_K generated from \preceq , is identical to \preceq_K .

To proceed with the proof we first need to introduce some more notation. For any variable $q \in P$, by q_K we denote q itself if $q \in K$, and the literal $\neg q$ otherwise. Clearly, since K is complete, $q_K \in K$ for all $q \in P$. For a set of variables $A \subseteq P$, by A_K we denote the set $A_K = \{q_K : q \in A\}$.

Next we show that for all $p, q \in P$, $p \preceq q$ iff $p_K \notin K * (\overline{p_K} \vee \overline{q_K})$. Consider any $p, q \in P$ such that $p \preceq q$ and

suppose towards contradiction that $p_K \in K * (\overline{p_K} \vee \overline{q_K})$. From the latter we derive that there is a $\overline{q_K}$ -world, call it r , such that $r \sqsubseteq_K^{\preceq} r'$ for all $r' \in [\overline{p_K}]$. Define r'' to be the world that agrees with K on all literals except q . Then clearly, $Diff(K, r'') = \{q\} \subseteq Diff(K, r)$. Hence $Diff(K, r'') \preceq Diff(K, r)$ and consequently, $r'' \sqsubseteq_K^{\preceq} r$. Next define u to be the world that agrees with K on all literals except p . Thus $Diff(K, u) = \{p\}$. Given that $u \models \overline{p_K}$, we derive that $r \sqsubseteq_K^{\preceq} u$, and consequently, $r'' \sqsubseteq_K^{\preceq} u$. Therefore $Diff(K, r'') \triangleleft Diff(K, u)$, which leads us to $q \triangleleft p$ contradicting our initial assumption $p \preceq q$. Hence we have shown that if $p \preceq q$ then $p \notin K * (\overline{p_K} \vee \overline{q_K})$.

For the converse, suppose that $p, q \in P$ are such that $p_K \notin K * (\overline{p_K} \vee \overline{q_K})$. Then there is a $\overline{p_K}$ -world, call it r , such that $r \sqsubseteq_K^{\preceq} r'$ for all $r' \in [\overline{q_K}]$. Let r'' be the world that agrees with K on all literals except p . Then $Diff(K, r'') = \{p\} \subseteq Diff(K, r)$. Consequently, $r'' \sqsubseteq_K^{\preceq} r$. Define u to be the world that agrees with K on all literals except q . Thus $Diff(K, u) = \{q\}$. Given that $u \models \overline{q_K}$, we derive that $r \sqsubseteq_K^{\preceq} u$, and consequently, $r'' \sqsubseteq_K^{\preceq} u$. This again entails that $p \preceq q$ as desired. Hence we have shown that for all $p, q \in P$, $p \preceq q$ iff $p \notin K * (\overline{p_K} \vee \overline{q_K})$.

We can now proceed to show the validity of the postulates (D1) - (D4).

For (D1), let $A, B \subseteq P$ be such that $\neg \overline{A_K} \notin K * (\overline{A_K} \vee \overline{B_K})$. We will show that $|A| \leq |B|$. Assume on the contrary that $|B| < |A|$. Call r the world that differs from K only over the variable in B . Then clearly, $Diff(K, r) = \{B\}$. Moreover, for any $\overline{A_K}$ -world r' , $A \subseteq Diff(K, r')$. Therefore, from $|B| < |A|$ we derive that for any $\overline{A_K}$ -world r' , $|Diff(K, r)| < |Diff(K, r')|$, and consequently, $Diff(K, r) \triangleleft Diff(K, r')$. This again entails $r \sqsubseteq_K^{\preceq} r'$, for all $r' \in [\overline{A_K}]$, and consequently, $\neg \overline{A_K} \in K * (\overline{A_K} \vee \overline{B_K})$. Contradiction.

For (D2), consider any $p, q \in P$ and $A, B \subseteq P$ such that $q \notin B$. Assume that $\neg \overline{A_K} \notin K * (\overline{A_K} \vee \overline{B_K})$ and $p_K \notin K * (\overline{p_K} \vee \overline{q_K})$. We will show that $\neg \overline{A_K \overline{p_K}} \notin K * (\overline{A_K \overline{p_K}} \vee \overline{B_K \overline{q_K}})$. From $p_K \notin K * (\overline{p_K} \vee \overline{q_K})$ it follows that $p \preceq q$. Moreover from $\neg \overline{A_K} \notin K * (\overline{A_K} \vee \overline{B_K})$ we derive that there is a $\overline{A_K}$ -world, call it r , such that $r \sqsubseteq_K^{\preceq} r'$, for all $r' \in [\overline{A_K}] \cup [\overline{B_K}]$. Define r'' to be the world that differs from K only over the variables in A . Clearly then, since $r \in [\overline{A_K}]$, we derive that $Diff(K, r'') = \{A\} \subseteq Diff(K, r)$. Consequently, $r'' \sqsubseteq_K^{\preceq} r$. Next define u to be the world that differs from K only over the variables in B . Then, $Diff(K, u) = B$ and $u \models \overline{B_K}$. Hence, $r \sqsubseteq_K^{\preceq} u$, and consequently $r'' \sqsubseteq_K^{\preceq} u$. This again entails that $A \preceq B$. From this, $p \preceq q$, and $q \notin B$, it is not hard to derive that $A \cup \{p\} \preceq B \cup \{q\}$. Next define w to be the world that differ from K only over the variables in $A \cup \{p\}$. Clearly then, $w \models \overline{A_K \overline{p_K}}$ and $Diff(K, w) = A \cup \{p\}$. Moreover observe that for any $\overline{B_K \overline{q_K}}$ -world w' , $B \cup \{q\} \subseteq Diff(K, w')$ and therefore $B \cup \{q\} \preceq Diff(K, w')$. Since $A \cup \{p\} \preceq B \cup \{q\}$ we then derive that $w \sqsubseteq_K^{\preceq} w'$ for all $w' \in [\overline{B_K \overline{q_K}}]$. This again entails $\neg \overline{A_K \overline{p_K}} \notin K * (\overline{A_K \overline{p_K}} \vee \overline{B_K \overline{q_K}})$ as desired.

For (D3), consider any $p, q \in P$ and $A, B \subseteq P$ such that $q \notin B$. Assume that $\neg \overline{A_K} \notin K * (\overline{A_K} \vee \overline{B_K})$ and $q_K \in K * (\overline{p_K} \vee \overline{q_K})$. We will show that $\neg \overline{B_K \overline{q_K}} \in K * (\overline{A_K \overline{p_K}} \vee \overline{B_K \overline{q_K}})$.

Firstly observe that from $\overline{q_K} \in K * (\overline{p_K} \vee \overline{q_K})$ we derive that $p \triangleleft q$. Moreover from $\overline{\neg A_K} \notin K * (\overline{A_K} \vee \overline{B_K})$ we derive that there is a $\overline{A_K}$ -world, call it r , such that $r \sqsubseteq_K^{\leq} r'$, for all $r' \in [\overline{A_K}] \cup [\overline{B_K}]$. Define r'' to be the world that differs from K only over the variables in A . Clearly then, since $r \in [\overline{A_K}]$, it follows that $\text{Diff}(K, r'') = \{A\} \subseteq \text{Diff}(K, r)$. Consequently, $r'' \sqsubseteq_K^{\leq} r$. Next define u to be the world that differs from K only over the variables in B . Then, $\text{Diff}(K, u) = B$ and $u \in [\overline{B_K}]$. Hence, $r \sqsubseteq_K^{\leq} u$, and consequently $r'' \sqsubseteq_K^{\leq} u$. This again entails that $A \leq B$, which in turn, when combined with $p \triangleleft q$, and $q \notin B$, leads to $A \cup \{p\} \triangleleft B \cup \{q\}$. Next define w to be the world that differ from K only over the variables in $A \cup \{p\}$. Clearly then, $w \in [\overline{A_K \overline{p_K}}]$ and $\text{Diff}(K, w) = A \cup \{p\}$. Moreover observe that for any $\overline{B_K \overline{q_K}}$ -world w' , $B \cup \{q\} \subseteq \text{Diff}(K, w')$ and therefore $B \cup \{q\} \leq \text{Diff}(K, w')$. Since $A \cup \{p\} \triangleleft B \cup \{q\}$ we then derive that $w \sqsubseteq_K^{\leq} w'$ for all $w' \in [\overline{B_K \overline{q_K}}]$. This again entails $\overline{\neg B_K \overline{q_K}} \in K * (\overline{A_K \overline{p_K}} \vee \overline{B_K \overline{q_K}})$ as desired.

Finally for (D4), consider any $p, q \in P$ and $A, B \subseteq P$ such that $q \notin B$. Assume that $\overline{\neg B_K} \in K * (\overline{A_K} \vee \overline{B_K})$ and for all $z \in B$, $z_K \notin K * (\overline{z_K} \vee \overline{q_K})$. We will show that $\overline{\neg B_K \overline{q_K}} \in K * (\overline{A_K \overline{p_K}} \vee \overline{B_K \overline{q_K}})$. First observe that from $z_K \notin K * (\overline{z_K} \vee \overline{q_K})$ for all $z \in B$, we derive that $z \leq q$ for all $z \in B$. In other words, q is \leq -maximal in $B \cup \{q\}$. Moreover from $\overline{\neg B_K} \in K * (\overline{A_K} \vee \overline{B_K})$ we derive that there is a $\overline{A_K}$ -world, call it r , such that $r \sqsubseteq_K^{\leq} r'$, for all $r' \in [\overline{B_K}]$. Define r'' to be the world that differs from K only over the variables in A . Clearly then, since $r \in [\overline{A_K}]$, it follows that $\text{Diff}(K, r'') = \{A\} \subseteq \text{Diff}(K, r)$. Consequently, $r'' \sqsubseteq_K^{\leq} r$. Next define u to be the world that differs from K only over the variables in B . Then, $\text{Diff}(K, u) = B$ and $u \in [\overline{B_K}]$. Hence, $r \sqsubseteq_K^{\leq} u$, and consequently $r'' \sqsubseteq_K^{\leq} u$. This again entails that $A \triangleleft B$. Then because q is \leq -maximal in $B \cup \{q\}$, it is not hard to verify that $A \cup \{p\} \triangleleft B \cup \{q\}$, for any variable $p \in P$. Define w to be the world that differ from K only over the variables in $A \cup \{p\}$. Clearly, $w \in [\overline{A_K \overline{p_K}}]$ and $\text{Diff}(K, w) = A \cup \{p\}$. Moreover observe that for any $\overline{B_K \overline{q_K}}$ -world w' , $B \cup \{q\} \subseteq \text{Diff}(K, w')$ and therefore $B \cup \{q\} \leq \text{Diff}(K, w')$. Since $A \cup \{p\} \triangleleft B \cup \{q\}$ we then derive that $w \sqsubseteq_K^{\leq} w'$ for all $w' \in [\overline{B_K \overline{q_K}}]$. This again entails $\overline{\neg B_K \overline{q_K}} \in K * (\overline{A_K \overline{p_K}} \vee \overline{B_K \overline{q_K}})$. \square

The converse of Theorem 2 is also true:

Theorem 3 *Let K be a consistent complete theory, $*$ an AGM revision function and \leq_K the faithful preorder that $*$ assigns to K . If $*$ satisfies (D1) - (D4) at K then \leq_K is a PD preorder.*

According to Theorem 3, if the revision function $*$ satisfies (D1) - (D4) at K , then there exists a preorder \leq over P , such that \sqsubseteq_K^{\leq} is identical to the preorder that $*$ assigns to K . Clearly, if $*$ also satisfies (D1) - (D4) at some other theory H , then Theorem 3 entails that the preorder that $*$ assigns to H can also be induced from some preorder \leq' over atoms. Notice however, that \leq and \leq' are not necessarily the same. To ensure this we need the axiom (D5) below:

(D5) If $p \leq_K q$, $x \in \{p, \overline{p}\}$, $y \in \{q, \overline{q}\}$ and $x, y \in H$, then $x \leq_H y$.

Axiom (D5) says that if for a given theory K it is at least as easy to reverse p than it is to reverse q , then this relationship is preserved for any other theory H and any other two literals x, y that share the same atoms with p and q respectively; for example if $p \leq_K q$ and $\neg p, q \in H$, then $\neg p \leq_H q$.

It can be shown that the addition of (D5) to (D1) - (D4) suffices to characterise the AGM revision functions that assign PD preorders to every consistent complete theory, all of which are generated from the *same* preorder \leq over P .

The General Case

Having characterised PD revision for the special case of consistent complete theories, we now turn to the general case where the initial belief set K is an arbitrary consistent theory.

The intuition behind the next axiom originates from the strong connection that exists in PD operators, between the revision of K by φ , and the revision of φ by K (which in turn is due to the fact that Diff is *symmetric*; i.e., $\text{Diff}(w, r) = \text{Diff}(r, w)$). Clearly there is a slight abuse of notation here that needs to be explained before we proceed any further. Since we are working with a propositional language built from *finitely many* propositional variables, for any theory K there exists a sentence ψ such that $K = \text{Cn}(\psi)$. Thus $\varphi * K$ is just an abbreviation for $\text{Cn}(\varphi) * \psi$; likewise, $w * \varphi$ is an abbreviation for $\text{Cn}(w) * \varphi$. With these clarifications we introduce our next axiom:

$$(D6) \quad K * \varphi = \bigcap_{w \in [\varphi * K]} w * \varphi.$$

It is worth noting (without proof due to space limitations) that a consequence of (D6) is the following intuitive property:¹¹

$$(KC) \quad \text{If } \neg\varphi \notin K * (\varphi \vee \psi) \text{ and } \neg\varphi \notin H * (\varphi \vee \psi), \text{ then } \neg\varphi \notin (K \cap H) * (\varphi \vee \psi).$$

Condition (KC) is quite easy to understand. It says that if bringing about φ is at least as easy as bringing about ψ , regardless of whether one starts at K or at H , then this is also true when one starts at their intersection $K \cap H$.

Like (KC), the next axiom associates the revision policies at two different theories K and H , with revisions at their intersection. However it deals with a more complex case. Suppose that φ is easier to bring about than ψ when starting at K , but it is the other way around when starting at H . What should happen when one starts at $K \cap H$? Axiom (D7) deals with this case. Yet instead of arbitrary sentence φ, ψ , (D7) refers only to sets of literals (φ is replaced by \overline{AB} and ψ by \overline{CD}):

$$(D7) \quad \text{If } A \leq_K E, B \subseteq K, \neg(\overline{AB}) \notin K * (\overline{AB} \vee \overline{CD}), C, D \subseteq H, \text{ and } L_C = L_E \text{ then } \neg(\overline{AB}) \notin (K \cap H) * (\overline{AB} \vee \overline{CD}).$$

Let us consider the intuition behind (D7). Assume that the antecedent of (D7) is true. Then, starting from K , it is at least as easy to bring about \overline{AB} as it is to bring about

¹¹The name of this condition comes from the fact that it is a stronger version of the first axiom for *kinetic consistency* reported in (Peppas and Williams 2016).

\overline{CD} .¹² If this is also the case when starting from H , then (KC) tells that the same is true when starting from $K \cap H$; i.e., $\neg(\overline{AB}) \notin (K \cap H) * (\overline{AB} \vee \overline{CD})$. What happens though if, starting from H , it is easier to reach \overline{CD} than to reach \overline{AB} . In that case, presumably that the rational thing to do is to compare the cost of the K -to- \overline{AB} transition with the cost of the H -to- \overline{CD} transition, and pick the outcome with the lowest transition cost. To do so, we will first map the H -to- \overline{CD} transition to a transition starting from K that has the same cost. In particular, observe that since $C, D \subseteq H$, the cost of the H -to- \overline{CD} transition is the cost of reversing the literals in C . Define E to be the literals in K that share the same variables with C ($L_E = L_C$). Then the cost of the H -to- \overline{CD} transition is the same as the cost of reversing E at K .¹³ If that cost happens to be at least as large as the cost of reversing A at K , then the H -to- \overline{CD} transition costs at least as much as the K -to- \overline{AB} transition. Hence, says (D7), the transition to \overline{AB} is at least as good as the transition to \overline{CD} , when the starting point can be either K or H : i.e., $\neg(\overline{AB}) \notin (K \cap H) * (\overline{AB} \vee \overline{CD})$.

The last axiom is essentially the strict version of (D7):

$$(D8) \text{ If } A <_K E, B \subseteq K, \neg(\overline{CD}) \in K * (\overline{AB} \vee \overline{CD}), C, D \subseteq H, \text{ and } L_C = L_E \text{ then } \neg(\overline{CD}) \in (K \cap H) * (\overline{AB} \vee \overline{CD}).$$

Observe that all axioms, including (D1) - (D5) originally introduced for complete theories, can be applied to any consistent theory. Theorem 4 below shows that in fact they are all satisfied by PD operators:

Theorem 4 *Let $*$ be an AGM revision function. If $*$ is a PD operator then it satisfies (D1) - (D8).*

The converse of Theorem 4 is also true:

Theorem 5 *Let $*$ be an AGM revision function. If $*$ satisfies (D1) - (D8) then $*$ is a PD operator.*

Implementations and Representational Cost

An implementation of AGM belief revision would presumably answer queries of the form “does ψ hold after the revision by φ ?”. These queries will be assessed against a background *knowledge base* B ¹⁴ and a revision policy associated with B . Revision policies can be modelled in different ways, however they are typically encoded as preorders \leq either over possible worlds (faithful preorders), or over sentences (epistemic entrenchments), or sets of sentences (remainders). The problem is that the size of these preorders is, in general, *exponential* to the number of atoms in the object language. This high *representational cost* is one of the main obstacles in the development of real-world belief revision applications.

¹²This follows from $\neg(\overline{AB}) \notin K * (\overline{AB} \vee \overline{CD})$.

¹³The silent assumption here is that the cost of changing a variable from positive to negative is the same as changing it from negative to positive. Hence, two sets of literals defined over the same variables, have exactly the same reversal cost. This property is a central feature of PD operators.

¹⁴A knowledge base B is a finite set of sentences representing the initial belief set K ; i.e., $K = Cn(B)$.

PD operators provide a very efficient solution to this problem: a single preorder \leq over the atoms of the object language L (hence *linear* in size to the number of atoms) suffices to determine the revision policy of *every* theory (or knowledge base) of L .

Observe that an added benefit of having a single preorder \leq generating the revision policy for all theories, is that we thus also solve the problem of *iterated revision* with no extra representational cost: the revision policy at $K * \varphi$ is fully determined by \leq , in the same way it is determined for any other theory.

Previous work on computational approaches to AGM revision, called *belief base revision schemes* in (Nebel 1998)), are primarily syntax-based. We briefly review two of the most influential such approaches and compare them to PD operators.¹⁵

The first one is based on the notion of *ensconcement* introduced in (Williams 1994).

Formally, an ensconcement \leq related to a belief base B is defined as a preorder over the elements of B that satisfies the following two constraints for all $\varphi \in B$:¹⁶

- (i) If $\models \varphi$, then $\{\psi \in B : \varphi < \psi\} \models \varphi$.
- (ii) $\models \varphi$ iff $\psi \leq \varphi$ for all $\psi \in B$.

Intuitively, an ensconcement can be thought of as a succinct representation of an epistemic entrenchment. Indeed, it was shown in (Williams 1994) that any ensconcement \leq over B can be extended to an epistemic entrenchment related to $Cn(B)$. Moreover it is possible to answer queries about the revision of $Cn(B)$, working directly with the ensconcement \leq , rather than the induced epistemic entrenchment. This addresses the problem of the representational cost, since the size of an ensconcement is linear to the size of the knowledge base B .

A second influential belief base revision scheme, called *linear belief base revision*, was introduced in (Nebel 1994). This approach partitions a knowledge base B into priority classes B_1, \dots, B_n . To revise B by a sentence φ , one removes an entire priority class B_i if (one or more of) its sentences are responsible for a contradiction with φ , and none of the lower priority classes can be blamed for the contradiction. It was shown in (Nebel 1994) that this procedure induces revision functions that satisfy all the AGM postulates for revision. Moreover, this approach also deals with the problem of the representational cost since any partition of B is linear to the size of B .

Comparing ensconcement-based revision and linear belief base revision with PD operators, one can immediately identify two advantages for the latter. Firstly the size of a knowledge base is typically much larger than the number of atoms, and therefore PD revision has (in principle) a lower representational cost. Secondly, and more importantly, PD revision has an embedded solution to the iterated revision

¹⁵See (Rott 2009) for more approaches that are likewise based on prioritised (belief) bases.

¹⁶The symbol \leq that we use herein for ensconcement will be used in the next section to compare numbers; any ambiguity with this slight abuse of notation is resolved by the context.

problem (at no extra representational cost). This is missing from both ensconcement-based revision and linear belief base revision: in both cases, new preorders need to be provided explicitly after each revision step (clearly a prohibitive requirement for real-world applications).

On the other hand, the formal results in (Nebel 1994), (Williams 1994) seem to suggest that ensconcement-based revision and linear belief base revision have a greater *range of applicability* than PD operators. In particular, it has been shown that both these approaches can encode *any* AGM revision function. In contrast, PD operators is a proper subclass of AGM revision functions (namely the subclass satisfying (D1) - (D8)). However a careful reading of the results in (Nebel 1994), (Williams 1994) reveals a somewhat different picture.

It is true that any AGM revision function can be generated from prioritised knowledge bases with the method described by Nebel. But only if the belief base B is allowed to *vary* according to the desired revision policy. More precisely, given a knowledge base B , and an AGM revision function $*$, it could well be the case that no prioritisation of B produces the same results as $*$ at $Cn(B)$. All that the results in (Nebel 1994) tell us is that, in that case there exists some *other* belief base B' , that is logically equivalent to B , for which such a prioritisation can be found.

Yet, we argue, that a knowledge base B ought to be independent from the revision policy employed. Adding to B (logically) redundant sentences, just to address the technical requirements of a certain revision representation method, is not in our view an elegant way to increase the range of applicability.

The results in (Williams 1994) also require a varying knowledge base, and therefore the same comments apply.

Complexity of PD Operators

We now turn to the computational complexity of PD operators. First we need to turn the computation of a PD operator into a decision problem.

We define a *PD revision instance* (or *PDR instance* for short) to be a tuple $\langle P, R, K, \varphi, \psi \rangle$ where,

- P is a nonempty set of propositional variables.
- R is a function from $P \mapsto [1..|P|]$, represented as a set of ordered pair (p, i) where $p \in P$ and $1 \leq i \leq |P|$.
- K is a consistent set of clauses over the variables in P .
- φ is a consistent set of clauses over the variables in P .
- ψ is a consistent set of clauses over the variables in P .

A PDR instance $Q = \langle P, R, K, \varphi, \psi \rangle$ represents a specific belief revision scenario. In particular, P represents the set of propositional variables over which beliefs are expressed, K represents the (base of) the current belief set, φ is the sentence by which K is revised, and ψ is the sentence we wish to test at the revised state (see below). The function R is used to represent a preorder \leq over the variables in P ; in particular, for any $p, q \in P$, $p \leq q$ iff $R(p) \leq R(q)$. Clearly \leq generates PD preorders, which in turn define a PD revision operator $*$. The decision problem associated with the

PDR instance Q , which we call the *PD revision problem*, is whether $Cn(K) * \varphi \models \psi$.

Observe that if $R(p) = 1$ for all $p \in P$, then $*$ reduces to Dalal's operator \square . In (Eiter and Gottlob 1992), it was shown that deciding if $Cn(K) \square \varphi \models \psi$ is $P^{NP[O(\log n)]}$ -complete (see their Theorem 6.9).¹⁷ Hence we immediately derive the following result.

Theorem 6 *The PD revision problem is $P^{NP[O(\log n)]}$ -hard.*

An upper bound to the computational complexity of the PD revision problem is given by the following theorem:

Theorem 7 *The PD revision problem belongs to $P^{NP[O(\sqrt{n} \log n)]}$.*

Proof. Let $Q = \langle P, R, K, \varphi, \psi \rangle$ be a PDR instance and let $*$ be the PD revision function associated with Q . We prove membership in the class $P^{NP[O(\sqrt{n} \log n)]}$ by outlining an algorithm that decides $Cn(K) * \varphi \models \psi$ with $O(\sqrt{n} \log n)$ calls to an NP oracle, where $n = |P|$.

The algorithm has three phases. In the first phase we compute the smallest number k in the set $\{|Diff(w, r)| : w \in [K] \text{ and } r \in [\varphi]\}$. Observe that $k \leq n$. Hence we can use binary search to determine k in $\log n$ steps, where at each step the question of whether $k \leq j$ (i.e., whether there exist $w \in [K]$ and $r \in [\varphi]$ such that $|Diff(w, r)| \leq j$), is decided with a call to the NP oracle.¹⁸

Before proceeding with the second phase of the algorithm we need some further notation and terminology.

Let P_1, P_2, \dots, P_m be the equivalence classes induced from \leq (alias R); i.e., the P_i 's are nonempty, pairwise disjoint sets, such that their union equals P , and moreover, for any $p, q \in P$, $p \leq q$ iff $p \in P_i$, $q \in P_j$ and $i \leq j$. Clearly, P_1, \dots, P_m , can be computed in deterministic polynomial time.

For a set of propositional variables $S \subseteq P$, define the *profile of S* to be the tuple $\langle j_1, \dots, j_m \rangle$, where $j_1 = |S \cap P_1|$, \dots , $j_m = |S \cap P_m|$; that is, the profile of S is the number of elements that S shares with each equivalence class P_1, \dots, P_m . A crucial observation, is that all \leq -minimal elements in $\{|Diff(w, r)| : w \in [K] \text{ and } r \in [\varphi]\}$, have the same profile. We call this profile, the *minimal profile* wrt K and φ and we shall denote it by $\langle y_1, \dots, y_m \rangle$. At the second phase our algorithm computes the numbers y_1, \dots, y_m .

Define $x_1 = |P_1|$, \dots , $x_m = |P_m|$. Hence, $x_i > 0$ and $\sum_{i=1}^m x_i = n$. Moreover, $0 \leq y_i \leq x_i$, for all $1 \leq i \leq m$; also it is not hard to see that $\sum_{i=1}^m y_i = k$.

The second phase of the algorithm starts with the computation of y_1 . Observe that y_1 is the maximal size of $Diff(w, r) \cap P_1$ under the constraints that $w \in [K]$, $r \in [\varphi]$ and $|Diff(w, r)| = k$. Hence y_1 can be computed with binary search in $\log x_1$ steps, where at each step the question

¹⁷We recall that a decision problem Π belongs to the class P^{NP} if it can be solved in polynomial time by a deterministic Turing machine M with an NP oracle. If in addition M can solve any instance of Π of length n , with no more than $g(n)$ calls to its NP oracle, we say that Π belongs to $P^{NP[g(n)]}$ – see the review in (Eiter and Gottlob 1992) for more details. For an excellent text on NP-completeness refer to (Garey and Johnson 1979).

¹⁸This first phase is identical to the first phase of the algorithm described in proof of Theorem 6.9 in (Eiter and Gottlob 1992).

whether $y_1 \geq j$ (i.e., whether there exist $w \in [K]$ and $r \in [\varphi]$ such that $|Diff(w, r)| = k$ and $|Diff(w, r) \cap P_1| \geq j$), is decided with a call to the NP oracle.

Now the rest of the y_i 's can be computed based on the following observation: y_{i+1} is the maximal size of $Diff(w, r) \cap P_{j+1}$, under the constraints that $w \in [K]$, $r \in [\varphi]$, $|Diff(w, r)| = k$, and $|Diff(w, r) \cap P_j| = y_j$ for all $1 \leq j \leq i$. Hence y_{i+1} can be computed with binary search in $\log x_{i+1}$ steps, where at each step the question whether $y_{i+1} \geq j$ is decided with a call to the NP oracle.

The whole minimal profile $\langle y_1, \dots, y_m \rangle$ can then be computed in polynomial time with $\log x_1 + \dots + \log x_m$ calls to an NP oracle. Given that $\sum_{i=1}^m x_i = n$, from the inequality of arithmetic and geometric means we derive that $\log x_1 + \dots + \log x_m \leq \frac{1}{2} \sqrt{n} \log n$. Hence in the first two phases our algorithm makes at most $O(\sqrt{n} \log n)$ calls to the NP oracle.

The third phase involves only one extra call to the oracle. In particular, the algorithm tests, with the aid of the NP oracle, whether there are worlds $w \in [K]$ and $r \in [\varphi]$ such that $Diff(w, r)$ has profile $\langle y_1, \dots, y_m \rangle$ and moreover $r \models \neg\psi$. If the answer is positive, the algorithm returns “no” to the original question “ $K * \varphi \models \psi?$ ”; otherwise it returns “yes”.

□

Theorems 6, 7 show that the PD revision problem belongs to the second level of the polynomial hierarchy. This is the same level where (the computation of) Dalal's operator belongs. Hence the added expressivity of PD operators doesn't have any drastic effects in time complexity.

We conclude this section by considering the restriction of the PD revision problem to Horn logic.

In particular, let $Q = \langle P, R, K, \varphi, \psi \rangle$ be a PDR instance such that all clauses in K , φ and ψ are Horn clauses. Moreover assume that the length of φ , denoted $||\varphi||$, is bounded by a constant k ; i.e., $||\varphi|| \leq k$. We shall call such a PDR instance, a *bounded Horn PDR instance* and the associated decision problem the *bounded Horn PDR problem*. We note that by bounding $||\varphi||$, $Diff(w, r)$ contains at most k variables, for all $w \in [K]$ and all worlds r that are \sqsubseteq_K -minimal in $[\varphi]$.¹⁹ Hence the following result can be obtained:

Theorem 8 *Let $\langle P, R, K, \varphi, \psi \rangle$ be a bounded Horn PDR instance and let $*$ be the revising function induced by it at K . Deciding if $Cn(K) * \varphi \models \psi$ can be computed in $O(||K|| \cdot ||\psi||)$ time.*

In the theorem above $||K||$ denotes the length of the knowledge base K and $||\psi||$ denotes the length of ψ . Notice that if $||\psi||$ is also bounded by a constant, then deciding if $Cn(K) * \varphi \models \psi$ can be done in time *linear* to the length of the knowledge base K .

Conclusion

PD operators is an important family of concrete AGM revision operators, essentially a generalisation of Dalal's operator, introduced in (Peppas and Williams 2016). The importance of PD operators is due to their low representational cost (any PD operator can be constructed from a preorder

over atoms) and their ability to cover a wide range of belief revision scenarios.

In this paper we have provided an axiomatic characterisation of PD operators. Moreover we have studied the computational complexity of PD operators showing that they lie at the same level of the polynomial hierarchy as Dalal's operator (despite the extra expressivity). In the special case of a Horn knowledge base the computation of bounded Horn queries, can be performed in linear time to the size of the knowledge base.

Observe that the preorder \leq over atoms that defines a PD operator is not effected by new evidence. An interesting direction for future work would be to consider the possibility of changing \leq depending on the new information that is received. This ought to be done with no or very little additional representational and computational cost, for otherwise the benefits of using PD operators would be cancelled.

Acknowledgements

The authors are grateful to Theofanis Aravanis and the anonymous reviewers for valuable comments on this work.

References

- C. Alchourron, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet functions for contraction and revision. *Journal of Symbolic Logic*, 50:510–530, 1985.
- C. Beierle and G. Kern-Isberner. A verified asml implementation of belief revision. *Lecture Notes in Computer Science*, 5238:98–111, 2008.
- A. Borgida. Language features and flexible handling of exceptions in information systems. *ACM Transactions on Database Systems (TODS)*, 10(4):563–603, 1985.
- T. Chou and M. Winslett. The implementation of a model-based belief revision system. *ACM SIGART Bulletin*, 2(3):28–34, 1991.
- M. Dalal. Investigations into theory of knowledge base revision: Preliminary report. In *Proceedings of 7th National Conference of the American Association for Artificial Intelligence (AAAI'88)*, pages 475–479, 1988.
- A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
- T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57:227–270, 1992.
- P. Gärdenfors. *Knowledge in Flux*. MIT press, 1988.
- M. Garey and D.S. Johnson. *Computers and Intractability – A Guide to the Theory of NP-Completeness*. W.H. Freeman, 1979.
- H. Katsuno and A. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.
- S. Konieczny, J. Lang, and P. Marquis. DA² merging operators. *Artificial Intelligence*, 157:49–79, 2004.
- B. Nebel. Base revision operations and schemes: Semantics, representation, and complexity. In *Proceedings of the 11th*

¹⁹This follows from an observation in (Eiter and Gottlob 1992).

European Conference on Artificial Intelligence, pages 341–345. Wiley, 1994.

B. Nebel. How hard is it to revise a belief base? In D. Dubois and H. Prade, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 3: Belief Change*, pages 77–145. Kluwer Academic, 1998.

P. Peppas and M.-A. Williams. Kinetic consistency and relevance in belief revision. In *Proceedings of the 15th European Conference on Logics in Artificial Intelligence (JELIA 2016)*, 2016.

P. Peppas. Belief revision. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 317–359. Elsevier Science, 2008.

P. Peppas. A panorama of iterated revision. In *David Makinson on Classical Methods for Non-Classical Problems - Outstanding Contributions to Logic 71-94*, pages 71–94. Springer, 2014.

H. Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. *Towards mathematical philosophy. Trends in logic*, 28:269–296, 2009.

K. Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 455–462, Tokyo, 1988. Springer-Verlag.

A. Weber. Updating propositional formulas. *Proc. First Conference on Expert Database Systems*, pages 487–500, 1986.

M.-A. Williams and A. Sims. Saten: An object-oriented web-based revision and extraction engine. In *International Workshop on Nonmonotonic Reasoning (NMR-2000)*. *Online Computer Science Abstract*, <http://arxiv.org/abs/cs.AI/0003059/>, 2000.

M.-A. Williams. On the logic of theory base change. In D. Pearce and L. Pereira, editors, *European Conference on Logics in Artificial Intelligence (JELIA'94)*, pages 86–105. Springer-Verlag, 1994.

M. Winslett. Reasoning about action using a possible models approach. In *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence (AAAI'88)*, pages 89–93, 1988.