Artificial Intelligence 229 (2015) 126–138

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

Relevance in belief revision ☆

Pavlos Peppas [a,b,*], Mary-Anne Williams [b], Samir Chopra [c], Norman Foo [d]

[a] *Dept of Business Administration, University of Patras, Patras 265 00, Greece*
[b] *Centre for Quantum Computation and Intelligent Systems, FEIT, University of Technology, Sydney, NSW 2007, Australia*
[c] *Dept of Philosophy, Graduate Center and Brooklyn College of the City University of New York, USA*
[d] *School of Computer Science and Engineering, University of New South Wales, NSW 2052, Australia*

ARTICLE INFO

Article history:
Received 13 August 2014
Received in revised form 12 August 2015
Accepted 15 August 2015
Available online 20 August 2015

Keywords:
Belief revision
Possible-world semantics
Knowledge representation

ABSTRACT

Possible-world semantics are provided for Parikh's relevance-sensitive axiom for belief revision, known as axiom (P). Loosely speaking, axiom (P) states that if a belief set K can be divided into two disjoint compartments, and the new information ϕ relates only to the first compartment, then the second compartment should not be effected by the revision of K by ϕ. Using the well-known connection between AGM revision functions and preorders on possible worlds as our starting point, we formulate additional constraints on such preorders that characterise precisely Parikh's axiom (P). Interestingly, the additional constraints essentially generalise a criterion of plausibility between possible worlds that predates axiom (P). A by-product of our study is the identification of two possible readings of Parikh's axiom (P), which we call the *strong* and the *weak* versions of the axiom. Regarding specific operators, we show that Dalal's belief revision operator satisfies both weak and strong (P), and it is therefore relevance-sensitive.

1. Introduction

Research in the area of *Belief Revision* examines the dynamics of belief sets. Much of the work in the field is based on the classical work of Alchourrón, Gärdenfors and Makinson, [1], that has given rise to a formal framework for studying this process, commonly referred to as the *AGM paradigm*, [15]. Within the AGM paradigm there are two constituents that are of particular interest for this article. The first is the set of rationality postulates for belief revision, known as the *AGM postulates* [1]. The second is a special kind of preorders on possible worlds, called *faithful preorders*, through which belief revision functions can be constructed. Katsuno and Mendelzon have shown, [9], that the revision functions constructed from faithful preorders are precisely those satisfying the AGM postulates.

The AGM postulates have been widely accepted as *sound* constraints for rational revision functions. There are however credible doubts about their *completeness*. For example, it is well-known that the AGM postulates do not suffice to address the case of *iterated revision* (see [16] for a recent survey).

Another weakness identified more recently by Parikh is that the AGM postulates for revision are rather liberal in their treatment of the notion of *relevance*. More precisely, Parikh argues that during belief revision a rational agent does not change her entire belief corpus, but only the portion of it that is relevant to the new information. This intuition of *local*

☆ This article is an extension and elaboration of previous work published in [12].

* Corresponding author at: Dept of Business Administration, University of Patras, Patras 265 00, Greece.
E-mail addresses: pavlos@upatras.gr (P. Peppas), Mary-Anne.Williams@uts.edu.au (M.-A. Williams), schopra@brooklyn.cuny.edu (S. Chopra), norman@cse.unsw.edu.au (N. Foo).

http://dx.doi.org/10.1016/j.artint.2015.08.007

change, Parikh claims, is not fully captured by the AGM postulates. To remedy this shortcoming, Parikh introduced an additional axiom named (P). Loosely speaking, axiom (P) says that when new information ϕ is received, only part of the initial belief set K will be affected, namely the part that shares common propositional variables with the minimal language of ϕ. Parikh's approach is also known as the *language splitting model*.

Although axiom (P) is just a first step towards capturing the role of relevance in belief revision, Parikh's work has received considerable attention since the publication of [11] (see for example [3–5,10]). Yet, despite all the research on axiom (P), no semantics for it have been formulated. Perhaps more importantly, it is not even known if axiom (P) is consistent with the AGM postulates. This is the gap that the present article aims to fill. We examine new constraints on faithful preorders and, building on Katsuno and Mendelzon's result, we prove that in the presence of the AGM postulates, axiom (P) is necessary and sufficient with respect to these new semantic constraints. Moreover we prove that, contrary to what Parikh conjectured in [11], axiom (P) is consistent with the full set of AGM postulates for revision.

What is particularly pleasing about our result is that the new constraints on faithful preorders are in fact not new at all; they essentially generalise a very natural condition that predates axiom (P) and has been motivated independently by other authors, like Borgida [2], Dalal [6], Satoh [17], and Winslett [18].

In course of formulating semantics for axiom (P) we observe that there are in fact two possible readings of this axiom, which we call the *strong* and the *weak* versions of (P). We present both these versions and show that strong (P) brings with it a new feature in the picture of classical AGM revision: it makes associations between the revision policies of *different* theories.

Using our semantic characterisation, we prove that Dalal's revision operator satisfies strong (P).[1] An immediate consequence of this result is that strong (P) is consistent with the full set of the AGM postulates for revision.

The outline of the paper is as follows. We first present some background material on the AGM paradigm and the language splitting model (first three sections). The crucial axiom (P) is then examined in greater detail to fully flesh out its possible readings (Section 5). We then proceed with the formulation of semantics for weak (P) and strong (P). For ease of exposition and clarity, we start by focusing on the special case of "opinionated" agents, that is, agents whose belief set is a consistent complete theory (Section 6). In Section 7 we consider the general case of incomplete theories. Section 8 examines Dalal's operator against axiom (P). The last section contains some concluding remarks.

2. Formal preliminaries

Throughout this paper we work with a finite set of propositional variables P. We define L to be the propositional language generated from P (using the standard Boolean connectives $\wedge, \vee, \rightarrow, \leftrightarrow, \neg$ and the special symbol $\perp$) and governed by classical propositional logic. For any subset of $Q \subseteq P$, by $Q^{\pm}$ we denote the set containing all literals induced by Q, i.e. $Q^{\pm} = Q \cup \{\neg p : p \in Q\}$.

A sentence $\phi \in L$ is *contingent* iff $\nvDash \phi$ and $\nvDash \neg\phi$. For a set of sentences Γ of L, we denote by $Cn(\Gamma)$ the set of all logical consequences of Γ, i.e., $Cn(\Gamma) = \{\phi \in L : \Gamma \models \phi\}$. We shall often write $Cn(\phi_1, \phi_2, \ldots, \phi_n)$, for sentences $\phi_1, \phi_2, \ldots, \phi_n$, as an abbreviation of $Cn(\{\phi_1, \phi_2, \ldots, \phi_n\})$. For any two sentences ϕ, ψ we shall write $\phi \equiv \psi$ iff $Cn(\phi) = Cn(\psi)$.

A theory K of L is any set of sentences of L closed under $\models$, i.e., $K = Cn(K)$. We denote the set of all consistent theories of L by $\mathcal{T}$. We define a *possible world* r (or simply a *world*), to be a consistent set of literals such that for any propositional variables $p \in P$, either $p \in r$ or $\neg p \in r$. We will often identify a world r with the conjunction of its literals, leaving it to the context to resolve any ambiguity (for example in "$\neg r$", r is a sentence, whereas in "$r \cap \{x\}$", r is a set of literals). We denote the set of all possible worlds by $\mathcal{M}$. For a set of sentences Γ of L, $[\Gamma]$ denotes the set of all possible worlds that entail Γ; i.e. $[\Gamma] = \{r \in \mathcal{M} : r \models \Gamma\}$. Often we use the notation $[\phi]$ for a sentence $\phi \in L$, as an abbreviation of $[\{\phi\}]$. For a theory K and a set of sentences Γ of L, we denote by $K + \Gamma$ the closure under $\models$ of $K \cup \Gamma$, i.e., $K + \Gamma = Cn(K \cup \Gamma)$. For a sentence $\phi \in L$ we often write $K + \phi$ as an abbreviation of $K + \{\phi\}$. For a set of worlds $V \subseteq \mathcal{M}$, by $t(V)$ we denote the set of sentences satisfied by all worlds in V; i.e. $t(V) = \{\phi \in L : r \models \phi$, for all $r \in V\}$. If $V = \varnothing$, then we define $t(V) = L$. It is not hard to see that $t(V)$ is always a theory.

In the course of this paper, we often consider *sublanguages* of L. Let Q be a subset of the set of variables in P. By L^Q we denote the sublanguage of L defined over Q. In the limiting case where Q is empty, we take L^Q to be the language generated by $\perp$ and the Boolean connectives. For a sentence χ of L, by L_χ we denote the *minimal* sublanguage of L within which χ can be expressed (i.e., L_χ contains a sentence that is logically equivalent to χ, and moreover no *proper* sublanguage of L_χ contains such a sentence). If χ is inconsistent, we take L_χ to be $L^{\emptyset}$.[2] Moreover, by P_χ we denote the propositional variables in the minimal language of χ.

By $\overline{L_\chi}$ we denote the language L^{P-P_χ}. Notice that since we have assumed that P is finite, for any theory K there is a sentence $\chi \in L$ such that $K = Cn(\chi)$. Hence we can define L_K, P_K, and $\overline{L_K}$, as L_χ, P_χ, and $\overline{L_\chi}$ respectively, for a $\chi \in L$ such that $K = Cn(\chi)$ (the choice of χ is irrelevant; any χ that is logically equivalent to K gives the same results).

---

[1] As shown in [9], Dalal's operator is the only one among a number of popular alternatives – like, Satoh's operator [17], Winslett's operator [18], Borgida's operator [2], etc. – that satisfies all the AGM postulates.

[2] It is not hard to verify that for every χ, L_χ is unique – see [11] for details.