RESPECTING RELEVANCE IN BELIEF CHANGE

DAVID MAKINSON AND GEORGE KOUROUSIAS

Abstract. In this paper dedicated to Carlos Alchourrón, we review an issue that emerged only after his death in 1996, but would have been of great interest to him: To what extent do the formal operations of AGM belief change respect criteria of *relevance*? A natural (but also debateable) criterion was proposed in 1999 by Rohit Parikh, who observed that the AGM model does not always respect it. We discuss the pros and cons of this criterion, and explain how the AGM account may be refined, if we so desire, so that it is always respected.

1. Introduction

Named after the trio Carlos Alchourrón, Peter Gärdenfors and David Makinson who introduced it in their paper of 1985, the AGM account of belief change offers a model of the logic of belief revision. Like many models, it is rather idealized. For example, it considers the properties of 'one-shot' belief change, with scant attention to further conditions that might be appropriate for iterated change. It is formulated in terms of classical propositional logic, rather than some richer language such as that of first-order logic. It identifies belief states with sets of formulae, rather than with more complex items. Finally, it makes use of only deductive consequence relations between propositions, disregarding any relations of uncertain inference.

These are major idealizations, and it is easy to dismiss the entire enterprise as too simple for practical use. Yet it has shown itself to be surprisingly robust. Since 1985 many researchers have put their minds to relaxing some of its constraints, enriching its apparatus and extending its scope, but these investigations have usually taken the simple basic structure as their point of departure. It has turned out that some of the enrichment problems are more recalcitrant than one might imagine – witness, for example, the lack on consensus on the many different accounts of iterated revision.

The basic AGM approach thus remains a starting point for fresh journeys, and a platform for novel constructions. In this paper we will see how it may be refined in order to respect an interesting criterion of relevance due to Parikh 1999. The refinement leads us to a new concept of classical logic, which is significant quite independently of the context of belief change: that of parallel interpolation.

Our purpose is to highlight the general ideas, and so we omit all proofs. These may be found in the more technical paper Kourousias and Makinson (to appear). However, we will need to assume some familiarity with classical propositional logic and also with AGM theory, in particular the AGM notions of partial meet contraction and revision.

The notation is essentially the same as that of the AGM 1985 paper itself. Formulae of classical propositional logic are indicated by lower-case Latin letters, sets of formulae by upper-case ones. Classical consequence is written as Cn when seen

as an operation and as $\vdash$ when taken as a relation. Classical equivalence as a relation is written as $\dashv\vdash$. The only differences with the notation of the 1985 paper are that we use an asterisk for the revision operation and a plain subtraction sign for contraction, in accord with current conventions.

2. Revision via Contraction

Suppose we begin with a belief set K, and wish to introduce a proposition x that is inconsistent with K, in such a manner that the resulting belief set contains x, but is nevertheless consistent if x itself is. This process is known as *revision*, and its subtlety arises from the fact that unlike simple closure under consequence, it can have more than one outcome. It is usually analysed into two steps.

- First, reduce K enough that it is no longer inconsistent with x (i.e. no longer implies $\neg x$), but without throwing away more than is strictly necessary for the task. This step is known as *contraction*, and can also have more than one possible outcome.
- The second step is to take the result $K - (\neg x)$ of the contraction, which we write without brackets as $K - \neg x$, add the input proposition x to get $(K - \neg x) \cup \{x\}$, and then close under classical consequence, getting $Cn((K - \neg x) \cup \{x\})$, written more briefly as $(K - \neg x) + x$. This second step, known as *expansion*, is evidently fully determinate, and quite unproblematic in so far as one is happy with classical consequence as a formalization of logical implication.

Revision is identified with the result of carrying out the two steps in that order: $K * x$ is defined as $(K - \neg x) + x$, i.e. as $Cn((K - \neg x) \cup \{x\}))$. If one is working with bases rather than sets closed under consequence, the second step is even simpler – we omit the application of Cn.

The analysis of revision into two steps, of contraction followed by expansion, is due to the philosopher Isaac Levi 1980, and the definition $K * x = (K - \neg x) + x$ is accordingly known as the Levi identity. The essential contribution of Alchourrón, Gärdenfors and Makinson in 1985 was to throw light on the underlying process of contraction. They introduced certain regularity conditions (often referred to as the AGM postulates) that contraction may plausibly be taken to satisfy despite its indeterminacy, and showed that those conditions are characterized by a specific kind of construction using intersections of maximal non-implying subsets, known as *partial meet* contraction. The Levi identity then permits us to pass from contraction to revision in a straightforward manner, on both the syntactic and semantic levels.

In what follows we will also work with contraction, drawing corresponding results for revision as corollaries.

3. Parikh's Criterion for Relevance in Belief Change

In 1999, Rohit Parikh observed that changes carried out using the AGM model may fail a natural criterion of relevance. They may discard more than they should, by eliminating from K items that are, in this sense, irrelevant to the inconsistency of K with the formula being introduced or discarded.

Parikh formulated the observation in terms of revision, but it may equally well be made in terms of contraction, as we will here. Moreover, it can be made irrespective of whether we are taking belief sets as sets already closed under classical consequence, or as arbitrary sets of formulae (belief bases) that need not be closed. Let

K be any set of formulae of classical propositional logic. Suppose $K \dashv\vdash \bigcup\{B_i\}_{i\in I}$ where for any two distinct $i,j \in I$, no elementary letter occurs in both some formula in B_i and some formula in B_j. That is, writing $E(B_i)$ for the set of all elementary letters occurring in formulae in B_i, suppose that the $E(B_i)$ are pairwise disjoint. Let x be a formula that we wish to contract from K. Following Parikh 1999, we say that a formula a is *irrelevant to* x (*modulo the representation of* K *as* $\bigcup\{B_i\}_{i\in I}$) iff there is no E_i that contains both some letter occurring in x and some (possibly different) letter occurring in a. Finally, we say that a is irrelevant to x (modulo K itself) iff there is *some* such representation of K, modulo which a is irrelevant to x.

It should be noted that this notion of irrelevance has nothing to do with so-called 'relevance logics', which are certain subsystems of classical logic that are too weak for contradictions to imply all formulae or for arbitrary formulae to imply tautologies. Parikh's definition of irrelevance is formulated in terms of classical logic alone.

Note also that we are really working with a three-place relation, of a being irrelevant to x modulo a third term. This third term is in the first place a representation of the set K of formulae as a letter-disjoint family; then, via existential quantification over those families, K itself. We will see shortly that the latter step can be reformulated using Parikh's 'finest splitting theorem'.

4. How AGM Contraction Can Fail Parikh's Criterion

As Parikh observed, an AGM contraction (indeed, we add, even a maxichoice AGM contraction) $K - x$ can eliminate formulae that are irrelevant to x modulo K. We give a quite trivial example.

Let p,q be two distinct elementary letters, and put $K = Cn(p, q)$. Then there is an AGM maxichoice contraction that puts $K - p$ to be $Cn(p \leftrightarrow q)$, thus eliminating not only p but also q from K. However, the letter q is irrelevant to p modulo K. This is because the representation of K by $\{p, q\}$ puts $E_1 = \{p\}$, $E_2 = \{q\}$, and neither of these two sets contains *both of* the letters p and q.

The example is robust in the sense that it goes through even if we work with belief bases rather than belief sets already closed under consequence. Put $K_0 = \{p \leftrightarrow q, q\}$, so that $Cn(K_0) = K$ above. Then one of the AGM maxichoice base contractions puts $K_0 - p$ to be $\{p \leftrightarrow q\}$, which eliminates q. However, the letter q is irrelevant to p modulo K_0 because there is *another* representation of K_0 as $\{p, q\}$, which puts $E_1 = \{p\}$, $E_2 = \{q\}$, and neither of these two sets contains both of the letters p and q.

5. Should Parikh's Criterion be Respected?

Of course, the question arises whether the elimination of irrelevant formulae as defined by Parikh is really undesirable. Is the failure of AGM to satisfy the criterion really a shortcoming? In the authors' view, this question does not have a categorical answer. Violation of this kind of relevance will be undesirable in some contexts but may be perfectly acceptable in others, depending on the epistemic policy guiding the contraction.

Consider the same example, where we are contracting the letter p from the closed belief set $K = Cn(p, q)$, or from its base $K_0 = \{p \leftrightarrow q, q\}$. If we are working with the base, we may perhaps regard it as supplying us with epistemic information, namely that its elements are particularly important items that deserve to be protected more

than other items not appearing in it. But even so, the base gives us no information to discriminate *between* its elements. In the example $K_0 = \{p \leftrightarrow q, q\}$, we are not told explicitly that the elementary letter q deserves protection more than the biconditional $p \leftrightarrow q$. We may wish to allow the possibility that the latter is more deeply entrenched, less vulnerable, than the former. In that case, when we discard p we will jettison the letter q and keep the biconditional, regardless of the fact that q is irrelevant to p modulo K in the sense that Parikh has defined.

On the other hand, there may be occasions in which we wish to treat elementary letters systematically as the only carriers of epistemic significance. In the authors' view, this policy is difficult to justify in theoretical terms, but it sometimes appears to be adopted for reasons of computational convenience in contexts of artificial intelligence. In this situation, regardless of whether we are working with belief bases or closed belief sets, we would want to preserve relevance in the sense of Parikh when contracting.

In summary: AGM contraction, even when maxichoice, can eliminate formulae that are irrelevant to the formula being discarded modulo the belief set undergoing contraction, in the sense that Parikh has defined. This is not necessarily undesirable: in some contexts it may be just what we want to allow. But in some others we may wish to prevent it.

6. How to Make AGM Contraction Respect Parikh's Criterion

So the question arises: Is it possible to refine the AGM operations so as to guarantee that relevance is respected?

Parikh and collaborators have approached this problem from a postulational perspective (see Chopra et al 2000, Peppas et al 2004): what conditions should be added to the AGM postulates to ensure that relevance is always respected? In our view, it is more perspicuous to approach it semantically: in what way may the operation of partial meet contraction be tweaked so as to ensure respect?

To answer this question, we make use of a tool that Parikh has also supplied: his 'finest splitting theorem'. Consider again any set K of formulae of classical propositional logic. Let $\{E_i\}_{i \in I}$ be any partition of the set E of all elementary letters of the language. We say that $\{E_i\}_{i \in I}$ is a *splitting of* K iff there is a family $\{B_i\}_{i \in I}$ of sets of formulae with each $E(B_i) \subseteq E_i$ and $K \dashv\vdash \bigcup\{B_i\}_{i \in I}$. Following customary terminology, we can say that one partition is at least as fine as another iff every cell of the latter is the union of cells of the former; equivalently, iff the equivalence relation associated with the former is a sub-relation of that associated with the latter.

In his 1999 paper Parikh showed that in the finite case (i.e. the case where the language has only finitely many elementary letters), every set K of formulae has a unique finest splitting. It is also possible to prove the result for the infinite case. Note that with this finest splitting theorem in hand one can streamline the definition of irrelevance itself, getting rid of the existential quantification. A formula a is irrelevant to x (modulo K) iff in *the unique finest splitting* $\{E_i\}_{i \in I}$ of K there is no E_i that contains both some letter occurring in x and some (possibly different) letter occurring in a. This is, in fact, the way that Parikh originally defined it.

How can the finest splitting theorem be used to make AGM contraction respect relevance? By not applying the operation directly to K itself, but rather to a representation of it in terms of its finest splitting. In other words, given a set K

of formulae and a formula x that we wish to discard from it, we first consider its unique finest splitting $\{E_i\}_{i\in I}$ and a set $K' = \bigcup\{B_i\}_{i\in I}$ where each $E(B_i) \subseteq E_i$ and $K \dashv\vdash \bigcup\{B_i\}_{i\in I}$. We then perform an AGM contraction on K', obtaining $K' - x$.

The first of these two steps may be thought of as a preliminary 'massaging' of K, to get it into a normal form K', which we call the *finest form* of K. Strictly speaking, it is not unique: there may be many such families $\{B_i\}_{i\in I}$, but they will all have the same letter-set family $\{E_i\}_{i\in I}$, namely the unique finest splitting of K, and that is all that matters.

Note that even when K is closed under classical consequence, i.e. $K = Cn(K)$, neither the individual sets B_i nor their union $K' = \bigcup\{B_i\}_{i\in I}$ will in general be so. Contraction on a closed set is thus reconstrued as contraction of a canonical base for it.

It can be shown that when an AGM (partial meet) contraction is performed on the finest form K' of K then it always respects relevance: it never eliminates from K' any formula a that is irrelevant to the discarded formula x in the sense of Parikh. That is, whenever a is irrelevant to x (modulo K'), then if $a \in K'$ to start with, still $a \in K' - x$. For the proof, we refer the reader to Kourousias and Makinson (to appear).

The corresponding result for revision follows immediately: whenever a is irrelevant to x (modulo K'), then if $a \in K'$ to start with, still $a \in K' * x$. For by the definition of revision from contraction using the Levi identity, $K' * x$ equals $(K' - \neg x) \cup \{x\}$ or its closure $Cn((K' - \neg x) \cup \{x\})$. Since the elementary letters occurring in $\neg x$ are just the same as those occurring in x, a is irrelevant to x iff it is irrelevant to $\neg x$ (modulo K' each time), so the result for contraction tells us that $a \in K' - \neg x \subseteq (K' - \neg x) \cup \{x\} \subseteq Cn((K' - \neg x) \cup \{x\})$ and we are done.

To avoid any misunderstanding, it should however be emphasized that these results do not give us any guidance on *which* contraction or revision operation we should apply to the finest form K' of K, even if we already have one for K. All they do is tell us that *whatever* AGM belief change operation we apply to K', it will respect Parikh's relevance criterion.

7. Parallel Interpolation in Classical Logic

In the preceding section we remarked that while Parikh proved the finest splitting theorem for the finite case, it can be extended to cover the infinite case as well. The proof is rather intricate, and we do not wish to describe it here. But we do want to draw attention to an interesting aspect: it can be carried out elegantly using a hitherto unnoticed strengthening of the well-known interpolation theorem for classical propositional logic.

The standard interpolation theorem (also known as Craig's Lemma) tells us that whenever $K \vdash x$ there is a formula b all of whose elementary letters are common to K and to x, such that $K \vdash b \vdash x$.

Now consider the case where $K \dashv\vdash \cup\{B_i\}_{i\in I}$ where the sets $E_i = E(B_i)$ are pairwise disjoint. Suppose $\bigcup\{B_i\}_{i\in I} \vdash x$.

We know from the standard interpolation theorem that there is a formula b all of whose elementary letters are common to $\bigcup\{B_i\}_{i\in I}$ and to x, such that $K \vdash b \vdash x$. But since the sets B_i do not *separately* imply x, interpolation does not tell us immediately whether we may treat the B_i in parallel, i.e. whether we can find

formulae b_i such that all of the elementary letters in each b_i are common to B_i and x, such that $B_i \vdash b_i$ and $\{b_i\}_{i \in I} \vdash x$, as in Figure 1 below.

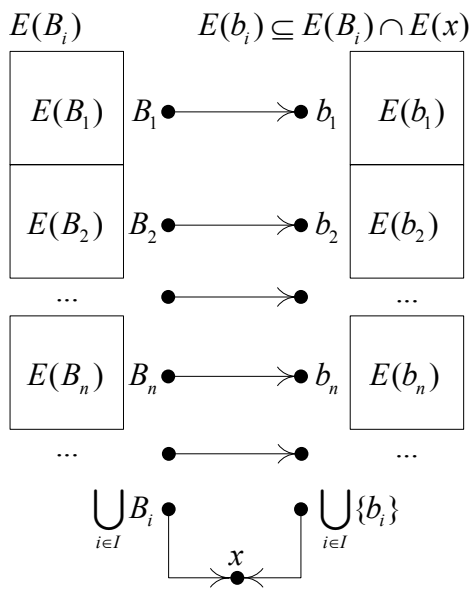

FIGURE 1. Parallel Interpolation

It turns out that parallel interpolation does hold for classical propositional and first-order logic. This can be proven by a direct argument but, as shown by Georg Gottlob (personal communication), it is simpler to obtain it by iterated applications of standard interpolation. The details are given in Kourousias and Makinson (to appear).

In this way, the investigation of a rather specialized problem in the logic of theory change that would have been of close interest to Carlos, namely the extent to which AGM belief revision operations respect relevance, leads us back into fresh views of classical logic. It led to an important theorem of Parikh (the finest splitting theorem) and its extension to the infinite case, and also takes us to an interesting strengthening of one of the fundamental theorems of classical logic (interpolation).

8. Final Remark

Despite all the accomplishments of the last century and a half, we see that there are still new things to be discovered in such an elementary area as classical propositional logic - not just matters of detail but also general concepts. Parallel interpolation and finest splitting are not the only examples; another is the concept of 'logical friendliness'. But that is another story, for which we refer the reader to Makinson 2005.

References

Alchourrón, Carlos, Peter Gärdenfors, and David Makinson 1985. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50:510–530.

Chopra, Samir and Rohit Parikh 2000. Relevance sensitive belief structures. *Ann. Math. Artif. Intell.*, 28(1-4):259–285.

Kourousias, George and David Makinson (to appear). Parallel interpolation, splitting, and relevance in belief change. The Journal of Symbolic Logic.

Levi, Issac 1980. *The Enterprise of Knowledge.* Cambridge, MA. MIT Press.

Makinson, David 2005. Friendliness for logicians. In Sergei N. Artemov, Howard Barringer, Artur S. d'Avila Garcez, Luis C. Lamb, and John Woods, editors, *We Will Show Them! Essays in Honour of Dov Gabbay, Volume Two*, pages 259–292. London. College Publications.

Parikh, Rohit 1999. Beliefs, belief revision, and splitting languages. In Lawrence Moss, Jonathan Ginzburg, and Maarten de Rijke, editors, *Logic, Language, and Computation, Vol. 2*, pages 266–278. Stanford, CA. CSLI Publications.

Peppas, Pavlos, Samir Chopra, and Norman Foo 2004. Distance semantics for relevance-sensitive belief revision. In Didier Dubois, Christopher A. Welty, and Mary-Anne Williams, editors, *KR2004: Principles of Knowledge Representation and Reasoning*, pages 319–328. Menlo Park, CA. AAAI Press.

Dept. of Philosophy, Logic & Scientific Method, London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom
E-mail address: `david.makinson@googlemail.com`

Dept. of Computer Science, Group of Logic & Computation, King's College London, Strand, London WC2R 2LS, United Kingdom
E-mail address: `george.kourousias@googlemail.com`