

Propositional Relevance through Letter-Sharing: Review and Contribution

David Makinson

Abstract

The concept of relevance between classical propositional formulae, defined in terms of letter-sharing, has been around for a very long time. But it began to take on a fresh life in the late 1990s when it was reconsidered in the context of the logic of belief change. Two new ideas appeared in independent work of Odinaldo Rodrigues and Rohit Parikh. First, the relation of relevance was considered modulo the belief set under consideration, Second, the belief set was put in a canonical form, known as its finest splitting. In this paper we explain these ideas; relate the approaches of Rodrigues and Parikh to each other; briefly report some recent results of Kourousias and Makinson on the extent to which AGM belief change operations respect relevance; and finally show how the introduction of a further parameter allows one to take into account epistemic and other components of relevance as well as purely logical ones.

1. Logical Relevance as a Two-Place Relation

The idea of defining a notion of relevance between propositional formulae goes back a long way. In the context of classical logic, which will be our focus, the simplest definition to suggest itself is the following:

Definition 1.1 Let a, b be formulae of classical propositional logic. They are *syntactically relevant* to each other iff they share some elementary letter.

Bibliographical: It is not clear when this definition was first formulated. Perhaps traces of it may even be found in Boole in the mid-nineteenth century. In the 1950s it was used as an adequacy condition by the founders of so-called relevance logic – a subsystem of classical logic satisfying the condition that syntactically irrelevant formulae never imply one another. In this paper we are not concerned with those logics: our concern is with refinements and deployments of the concept in *classical contexts*.

Shortcoming: As defined, the notion of syntactic relevance is syntax-dependent. In other words, formulae a, b may be classically equivalent to a', b' respectively, and a relevant to b but a' not relevant to b' .

Example: $\neg p \wedge (\neg p \vee q)$ is syntactically relevant to q , but the former is classically equivalent to $\neg p$ which is not relevant to q . Here and always in the paper, p, q, \dots are understood to be elementary letters while a, b, \dots are arbitrary formulae.

To overcome this, the obvious move is to express each formulae in its least letter-set, using the well-known:

Fact: for every formula a , there is a unique least set of elementary letters such that a may equivalently be expressed using only letters from that set.

Examples: The unique least letter-set of $\neg p \wedge (\neg p \vee q)$ is $\{p\}$. On the other hand, the unique least letter-set of $\neg p \wedge (\neg r \vee q)$ is $\{p, q, r\}$.

Remarks: (1) Strictly speaking, the result holds in this simple form only when the language has a primitive zero-ary operator (propositional constant) such as the falsum. In such a language, the least letter-set of any non-contingent formula is \emptyset . Without a zero-ary connective, say with just \neg, \wedge, \vee , any tautology or contradiction has many minimal letter-sets (in fact, all the singleton letter sets), but has no least one (since no formula is bereft of letters). For simplicity of formulation, in this paper we work with a language that does have a primitive zero-ary connective, e.g. the falsum. (2) The result stated is intuitively obvious, but needs proof. Getting *minimal* letter sets is trivial since every formula contains only finitely many letters; but getting a *least* one (which, by the antisymmetry of set-inclusion, will be unique) requires a bit more work. See the appendix of Makinson (2007), where the result is shown to hold, more generally, for arbitrary sets A of formulae.

We write a^* for an (arbitrarily chosen) formula equivalent to a that is built in the least letter-set for a .

Definition 1.2 Let a, b be formulae of classical propositional logic. They are said to be *essentially relevant* to each other iff a^*, b^* share some elementary letter. Equivalently: iff every formula equivalent to a shares a letter with every formula equivalent to b .

Example: Although $\neg p \wedge (\neg p \vee q)$ is syntactically relevant to q , it is not essentially so, since $(\neg p \wedge (\neg p \vee q))^* = \neg p$ shares no letter with $q^* = q$.

Features: Essential relevance has the following properties:

- It is syntax-independent (immediate from definition).
- It is symmetric (immediate from definition).
- Reflexive? Nearly: every *contingent* formula is relevant to itself. Non-contingent formulae are not relevant to anything.
- Not transitive. *Example:* p is essentially relevant to $p \wedge q$ which is so to q , but p is not so to q .
- Cannot be ‘made transitive’: its transitive closure makes any two contingent formula relevant to each other. *Verification:* Take contingent a, c . Since a is contingent, a^* contains a letter p , likewise c^* contains a letter q . Put $b = p \wedge q = b^*$. Then a is essentially relevant to b , also b to c , so transitive closure would make a relevant to c .
- No two distinct elementary letters are relevant to each other (immediate from definition).

This is all part of the folklore, and dates back a long way. However, things began to take a fresh turn in the late 1990s, when a few people began thinking about relevance in the context of formal accounts of belief change. Two basic insights emerged. The

first was that in this context, the relevance or irrelevance of one formula to another may be taken to depend not only on the formulae themselves but *also on the belief set under consideration*. The second was that this belief set may be given a canonical form known as its *finest splitting*. The following two sections explain and comment on them.

2. Path-Relevance Modulo a Belief Set

Suppose that we have the belief set $K = \{p \rightarrow q, q \rightarrow r\}$, where p, q, r are distinct elementary letters. As such, they are irrelevant to each other. But it is natural to say that *from the point of view of the belief set K* , p is relevant to q , q is relevant to r , and p is thus indirectly relevant to r . This suggests the following definition.

Definition 2.1. (Odinaldo Rodrigues). Let a, b be formulae of classical propositional logic, and let K be a set of formulae serving as a belief set. We say that a is *path-relevant* to b (mod K) iff there is a finite sequence x_0, \dots, x_{n+1} ($n \geq 0$) of formulae with $x_0 = a^*$, $x_{n+1} = b^*$, $x_1, \dots, x_n \in K$, and each x_i shares at least one letter with x_{i+1} .

Comments: Note that x_1, \dots, x_n are required to be elements of K . Thus we are looking at finite paths through K . On the other hand, it is not required that either of $x_0 = a^*$, $x_{n+1} = b^*$ is in K (although of course they may be).

Bibliographical: This notion was introduced by Rodrigues in his thesis (1997), Appendix A, definition 8.14. It was used by Renata Wassermann in her thesis (1999) and in subsequent papers e.g. Riana and Wassermann (2004). Actually, all these authors used a, b instead of a^*, b^* in the definition, but we make it syntax-independent in those two arguments.

Path-relevance generalizes essential relevance in a natural way: the latter amounts to requiring that $n = 0$ in Definition 2.1. Like essential relevance, has the following properties:

- Syntax-independent in a, b .
- Symmetric.
- Almost reflexive: Every contingent formula is relevant to itself, but no tautology or contradiction is relevant to anything.,
- Not transitive.

With K as parameter, new features emerge:

- Distinct elementary letters can be relevant to each other (mod K). *Example:* With $K = \{p \rightarrow q, q \rightarrow r, \neg s\}$, p is path-relevant to q but not to s .
- However, the relation is syntax-dependent in K . *Example:* Add to the above K the formula $(r \rightarrow s) \vee (s \rightarrow r)$. As this is a tautology, it does not change the strength of K . But p is now path-relevant to s .
- The relation trivializes when the belief set is closed under classical consequence. When $K = Cn(K)$, any two contingent formulae a, b are path-relevant to each other modulo K . *Verification:* Since a, b are contingent, each

of a^* , b^* has at least one letter. Take any letter p in a^* , any letter q in b^* , note that $(p \vee \neg p) \rightarrow (q \vee \neg q) \in Cn(K)$.

Can we get around the syntax-dependence in the argument K ? It is natural to try tweaking Definition 2.1, replacing x_1, \dots, x_n by their least letter-set versions x_1^*, \dots, x_n^* . However, this does not eliminate syntax-dependence. *Example:* Compare $K = \{p \wedge q\}$ with p relevant to q (mod K) versus the equivalent $K' = \{p, q\}$ with p irrelevant to q (mod K'). A better idea is needed, and one was provided by Rohit Parikh in 1999, with his concept of a *splitting* of a belief set.

3. Splittings of a Belief Set

Definition 3.1. Let K be a (non-empty) belief set, expressed in the language of classical propositional logic (with a zero-ary connective). Let E be the set of all elementary letters of the language (or occurring in formulae in K). Let $\mathbf{E} = \{E_i\}_{i \in I}$ be a partition of E (be careful: we partition the letter-set E , not the belief set K). We say that \mathbf{E} is a *splitting* of K iff there is a family $\{B_i\}_{i \in I}$ of sets of formulae such that each $E(B_i) \subseteq E_i$ and $K \dashv\vdash \cup\{B_i\}_{i \in I}$. In other words, iff K can be represented as the union of belief sets each of which uses only letters from one of the cells of the partition.

Background on partitions: Recall that a partition of a (non-empty) set is a family of disjoint non-empty subsets of that set, whose union exhausts the set. The partitions of a set can be put in one-one correspondence with the equivalence relations over the set. One partition is said to be finer than another iff the equivalence relation associated with the former is included (set-theoretically) in the equivalence relation associated with the latter. Equivalently, if every cell of the first partition is a subset of a cell of the second one. Recall that the infimum under fineness of any non-empty family of partitions of a set (i.e. the partition corresponding to the intersection of all the equivalence relations associated with partitions in the family) is also a partition of that set.

Example: Let $K = \{p \rightarrow \neg q, \neg q \rightarrow r, p \vee s, \neg s, (r \rightarrow t) \vee (t \rightarrow r)\}$. Here $E = \{p, q, r, s, t\}$.

- The coarsest splitting of K is the singleton partition with E itself as its only cell, putting $B = K$. But we can do better than that.
- Slightly less coarse is the partition into two cells $E_1 = \{p, q\}$ and $E_2 = \{r, s, t\}$ with $B_1 = \{p, \neg q\}$, $B_2 = \{r, \neg s\}$. The letter t does not appear in either of the B_i , but that is not a problem – the definition requires only the inclusion $E(B_i) \subseteq E_i$ for each $i \in I$.
- The finest splitting of K partitions E into five singleton cells $\{p\}, \{q\}, \{r\}, \{s\}, \{t\}$, with $B_1 = \{p\}$, $B_2 = \{\neg q\}$, $B_3 = \{r\}$, $B_4 = \{\neg s\}$, $B_5 = \emptyset$. Although each E_i must be non-empty (since it is a cell of a partition) the corresponding B_i may be empty.
- In this example, for simplicity, the finest partition has singleton cells and the associated sets B_1 to B_5 consist only of literals. Of course, neither need always be the case. For instance take $K = \{(p \rightarrow q) \wedge (r \rightarrow s)\}$. Its finest partition is into

the two-element cells $\{p,q\}$, $\{r,s\}$ with $B_1 = \{p \rightarrow q\}$, $B_2 = \{r \rightarrow s\}$ containing non-literals.

Theorem 3.1. (Rohit Parikh 1999). Every set K of formulae of classical propositional logic has a unique finest splitting.

Comments:

- It is the splitting $\mathbf{E} = \{E_i\}_{i \in I}$ of elementary letters that is unique. Given such a family, there will evidently be many families $\{B_i\}_{i \in I}$ of sets of formulae with $\cup\{B_i\}_{i \in I} \dashv\vdash K$ and $E(B_i) \subseteq E_i$.
- However, it simplifies formulations if we take a choice function associating with each K , having finest splitting $\mathbf{E} = \{E_i\}_{i \in I}$ of elementary letters, some particular such family $\{B_i\}_{i \in I}$, and write $\cup\{B_i\}_{i \in I}$ as $K^\#$. We abuse terminology a little by also calling $K^\#$ the finest splitting of K .
- Note that when two belief sets are classically equivalent, they will have exactly the same finest splitting $\mathbf{E} = \{E_i\}_{i \in I}$ and thus exactly the same finest splitting $K^\# = \cup\{B_i\}_{i \in I}$.

Bibliographical: This theorem was proven by Parikh (1999) for the finite case. The infinite case proven by Kourousias and Makinson (2007), using a new form of interpolation called “parallel interpolation”. Both parallel interpolation and the finest splitting theorem can be extended to first-order logic.

4. Canonical Relevance (Modulo a Belief Set)

How can splitting help us make the notion of relevance modulo a belief set fully syntax-independent? The finest decomposition $K^\# = \cup\{B_i\}_{i \in I}$ of K may be seen as a canonical form for the belief set K , disentangling the roles of the different elementary letters as far as is possible without altering the power of K . We can refine Rodrigues’ definition of path relevance by taking the path through this canonical representation instead of through K itself. Thus in Definition 2.1, replacing $x_1, \dots, x_n \in K$ by $x_1, \dots, x_n \in K^\#$, we have the following.

Definition 4.1. Let a, b be formulae of classical propositional logic, K be a set of formulae serving as a belief set, and $K^\#$ be the finest splitting of K . We say that a is *canonically path-relevant* to b (mod K) iff there is a finite sequence x_0, \dots, x_{n+1} ($n \geq 0$) of formulae with $x_0 = a^*$, $x_{n+1} = b^*$, $x_1, \dots, x_n \in K^\#$, and each x_i sharing at least one letter with x_{i+1} .

Comments: This time x_1, \dots, x_n are required to be elements of the canonical form $K^\#$, so we are looking at finite paths through $K^\#$ (rather than through K itself). As before, it is not required that either of $x_0 = a^*$, $x_{n+1} = b^*$ is in $K^\#$ (although of course they may be).

To help the reader keep track of successive definitions, Appendix 1 contains a table of the different kinds of relevance considered in the paper.

Features of canonical path relevance modulo K :

- As desired: syntax-independent in each of its arguments a, b, K . For the argument K , this follows from the fact, noted above, that equivalent belief sets have the same finest partition.
- Like plain path-relevance, it is symmetric but not transitive (in the arguments a, b); almost reflexive; distinct elementary letters can be relevant to each other (mod K).

There is another way of doing the same thing. It also uses Parikh's notion of the finest splitting $K^\#$ of K , but does not consider *paths*. Instead, it looks at *cells*.

Definition 4.2. (Rohit Parikh 1999). Let a, b be formulae of classical propositional logic, let K be a set of formulae serving as a belief set, with $\mathbf{E} = \{E_i\}_{i \in I}$ the finest splitting of K . We say that a is *canonically cell-relevant* to b (mod K) iff there is a cell E_i of \mathbf{E} such that each of a^* and b^* shares some letter (not necessarily the same letter) with E_i . More formally: iff for some $i \in I$, each of the sets $E(a^*) \cap E_i$ and $E(b^*) \cap E_i$ is non-empty.

Table 4.1: Illustration of canonical cell-relevance

E_1		E_2		E_3	
p	q	r	s	t	u
$E(a^*)$			$E(b^*)$		

In this illustration, the finest partition \mathbf{E} of K has three cells, each containing two elementary letters. The letters in a^* and b^* are disjoint, but there is a cell (the middle one) that contains letters r, s from both $E(a^*), E(b^*)$ respectively.

Bibliographical: Actually, Parikh (1999) worked with a, b rather than with a^*, b^* , and so did Kourousias and Makinson (2007). This makes no difference to the particular applications to AGM belief change operations made in those papers. But when considering the notion of relevance from a general perspective, it is evidently better to work with the least letter-set forms a^*, b^* .

Theorem 4.1. Canonical path-relevance is equivalent to canonical cell-relevance. In detail: let a, b be formulae of classical propositional logic, and let K be a set of formulae serving as a belief set. Then a is canonically path-relevant to b (mod K) iff it is canonically cell-relevant to b (mod K).

Proof: See Appendix 2.

Summary of the story so far: By using Parikh's notion of the finest splitting of K , we can refine Rodrigues's account of relevance to make it syntax-independent in all three arguments a, b, K . This notion of canonical path-relevance is equivalent to the more semantic-looking definition of canonical cell-relevance. This equivalence confirms the robustness of the concept, which henceforth we call simply *canonical relevance*.

Warning: When $K_1 \vdash K_2$ then since each $K_i^\# \dashv\vdash K_i$ we have $K_1^\# \vdash K_2^\#$. But it does not follow that when $K_1 \vdash K_2$ and a is canonically relevant to $b \pmod{K_2}$ then a is canonically relevant to $b \pmod{K_1}$. Example: Put $K_1 = \{p \wedge q\}$, $K_2 = \{p \rightarrow q\}$. Then $K_1 \vdash K_2$ and p is canonically relevant to $q \pmod{K_2}$ but not so $\pmod{K_1}$. Canonical relevance depends on the logical power of K , but is not monotonically increasing in that power.

5. Respecting Canonical Relevance in Belief Change

We recall briefly some applications of canonical relevance to the study of AGM belief change in the manner of Alchourrón, Gärdenfors and Makinson 1985. We focus on the operation of contraction (omitting revision) and omit all proofs, which can be found in Kourousias and Makinson (2007).

Definition 5.1. We say that an operation $-$ of contraction on a belief set K respects canonical relevance iff whenever $K \vdash x$ but $K-a \not\vdash x$ then a is canonically relevant to $x \pmod{K}$. Contrapositively, whenever $K \vdash x$ and a is canonically irrelevant to $x \pmod{K}$ then still $K-a \vdash x$.

Comment: When K is closed under classical consequence, i.e. when $K = Cn(K)$ then for AGM contraction $K-a$ is also closed under consequence, so we have $K \vdash x$, $K-a \vdash x$ iff respectively $x \in K$, $x \in K-a$, and thus Definition 5.1 becomes equivalent to one with epsilon replacing turnstile: whenever $x \in K$ and a is canonically irrelevant to $a \pmod{K}$ then still $x \in K-a$.

Observation (Parikh 1999): AGM contraction can fail to respect canonical relevance, and this can happen independently of whether K is closed under consequence.

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 canonically irrelevant to p modulo K . This is because we can split $E = \{p, q\}$ into $E_1 = \{p\}$, $E_2 = \{q\}$ with $K^\# = \{p, q\}$, and neither of these two cells contains *both* of the letters p and q .

The example is robust in the sense that it goes through even when 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 eliminated letter q is canonically irrelevant to p modulo K_0 for the same reason as before.

Theorem 5.1 (Kourousias and Makinson 2007). If we apply AGM contraction to the finest splitting $K^\#$ of a consistent belief set K , rather than to K itself, then it respects canonical relevance.

Example: In the above example, we would be applying the contraction operation to the canonical belief set $K^\# = \{p, q\}$ rather than to $K = Cn(p, q)$ or to $K_0 = \{p \leftrightarrow q, q\}$. Since there is just one maximal p -nonimplying subset of $K^\#$, namely $\{q\}$, it follows

that in this example there is just one possible output of an AGM belief contraction operation $K^\#-p$, namely $\{q\}$.

Comments: (1) Actually, the observation of Parikh (1999) was made for AGM revision, but the counterexamples for the two operations are essentially the same. (2) Theorem 5.1 was established by Kourousias and Makinson (2007) only for the epsilon version of respecting equivalence, rather than the turnstile version. The two are not the same, as remarked by Pavlos Peppas (personal communication). However, it is not difficult to obtain the turnstile version of the theorem from the epsilon one, as is done in Appendix 3.

6. Should Canonical Relevance Always be Respected?

Of course, we may ask whether it is really a shortcoming in a belief contraction operation to eliminate canonically irrelevant formulae. Is this failure to respect canonical relevance a defect, or just a feature of AGM contraction?

To answer this question, we need to distinguish between logical and epistemic components of relevance. The notion of canonical relevance appears to capture well the logical or formal component, but leaves aside entirely the epistemic one.

Consider again the example where we wish to contract the belief base $K_0 = \{p \leftrightarrow q, q\}$ to discard p . If we are interested only in logical matters, then we note as above that $K_0^\# = \{p, q\}$ so that q is not canonically relevant to p modulo K_0 . In that context, discarding p should not lead us to eliminate q , but rather $p \leftrightarrow q$, which *is* canonically relevant to p modulo K_0 .

But it may also happen that the formula $p \leftrightarrow q$ has a special place among our beliefs. It may be more deeply entrenched, less vulnerable, or in some other way epistemically more basic than either of the letters p, q or their conjunction $p \wedge q$, all of which are elements of $Cn(K_0)$. In that context, when discarding p we should keep the biconditional $p \leftrightarrow q$ and jettison the letter q . The eliminated formula q is not *logically* relevant to the formula p that we are discarding, but it is *epistemically* so, since it occurs in a formula $p \leftrightarrow q$ to which we are attributing special epistemic status.

In general, when a belief set is presented by a base, we may have differing attitudes towards the propositions in the base. Some may be there for convenience or by happenstance, and another equivalent base lacking them may be deemed as just as appropriate. But others may be in the base because we *want* them to be there; they may have an epistemic priority over consequences outside the base and even some others within it.

Essentially this perspective was suggested in the brief discussion in Makinson and Kourousias (2007). Now we take the analysis further by showing how we can take formal account of such extra-logical considerations.

7. Generalizations: Parametrized and Epistemic Relevance

Of course, logic alone cannot specify which propositions are deemed to have a particular epistemic status. But it can introduce into its constructions parameters that allow such considerations to play a role. We now express these intuitive ideas more formally. This requires generalizing some definitions and results of previous sections.

First, we observe that the notion of canonical cell/path relevance, which was introduced (Definition 4.2) using the finest splitting of K , may be generalized to a notion of cell/path relevance with respect to an *arbitrary* splitting. In terms of cells, this may naturally be done as follows:

Definition 7.1. Let a, b be formulae of classical propositional logic, let K be a set of formulae serving as a belief set, and $\mathbf{E} = \{E_i\}_{i \in I}$ any splitting of K . We say that a is *relevant to b modulo \mathbf{E}* iff there is a cell E_i of \mathbf{E} such that each of a^* and b^* shares some letter (not necessarily the same letter) with E_i . More formally: iff for some $i \in I$, each of the sets $E(a^*) \cap E_i$ and $E(b^*) \cap E_i$ is non-empty.

In turn, the notion of respect for canonical relevance, which was formulated with respect to the finest splitting, may be reformulated with respect to an arbitrary splitting.

Definition 7.2. Let K be a set of formulae serving as a belief set, and $\mathbf{E} = \{E_i\}_{i \in I}$ any splitting of K . We say that an operation $-$ of contraction on K *respects relevance modulo \mathbf{E}* iff whenever $K \vdash x$ but $K-a \not\vdash x$ then a is relevant to x modulo \mathbf{E} . Contrapositively, whenever $K \vdash x$ and a is irrelevant to x modulo \mathbf{E} then still $K-a \vdash x$.

Next, we note that two key theorems may also be strengthened to cover the more general context of arbitrary splittings.

Theorem 7.1. Let K be any set of formulae of classical propositional logic. The infimum of any non-empty family of splittings of K is also a splitting of K .

Proof: The proof of Theorem 3.1 (the finest splitting theorem) that is given in Kourousias and Makinson (2007) may be applied without change to yield this generalization.

Theorem 7.2 Let K be a consistent set of formulae serving as a belief set. If we apply AGM contraction to an arbitrarily chosen splitting of K , then it respects relevance modulo that splitting.

Proof: Simply re-run the one given for Theorem 5.1.

With these generalizations in hand, we can introduce a parameter to handle extra-logical (and in particular, epistemic) sources of relevance.

Definition 7.3. Let K be a set of formulae serving as a belief set, and let R be any relation between elementary letters. We say that a splitting $\mathbf{E} = \{E_i\}_{i \in I}$ of K *protects R* iff whenever $(p, q) \in R$ then p, q are in the same cell E_i of \mathbf{E} .

Theorem 7.3. Let K be a set of formulae serving as a belief set, and let R be any relation between elementary letters. Then K has a (unique) finest R -protecting splitting.

Proof. By Theorem 7.1, the infimum of all R -protecting splittings of K is a splitting of K , and it is straightforward to check that it also protects R .

Theorem 7.4. Let K be a consistent set of formulae serving as a belief set, and let R be any relation between elementary letters. If we apply AGM contraction to that finest R -protecting splitting rather than to K itself, then it respects relevance modulo that splitting.

Proof. Immediate from Theorem 7.2.

There are many ways, semantic and syntactic, of specifying a relation R between letters that we may want to be protected. One simple method is to distinguish a special subset of the belief set consisting of formulae whose syntactic expression we regard as carrying epistemic information about the formulae in it and the connections between the elementary letters occurring in those formulae. This suggests the following:

Definition 7.4. Let K be a belief state and let $K_1 \subseteq K$. We say that a splitting $\mathbf{E} = \{E_i\}_{i \in I}$ of K protects K_1 iff it protects the relation R defined by putting $(p, q) \in R$ iff there is a formula $x \in K_1$ containing both p and q .

The relation of protection evidently depends on the syntax of the formulae in K_1 . For example, a conjunction in K_1 will not have the same effect as the two conjuncts in it, since we are looking at two formulae rather than a single one. But this is deliberate: a formula is put in K_1 only when we see its syntactic form as carrying epistemic information. From Theorems 7.3 and 7.4 we have immediately:

Corollary to Theorems 7.3 and 7.4. Let K be a set of formulae serving as a belief set, and let $K_1 \subseteq K$. Then K has a (unique) finest K_1 -protecting splitting. Moreover, for consistent K , if we apply AGM contraction to the splitting rather than to K itself, then it respects relevance modulo that splitting.

The kinds of relevance defined in earlier sections come out as limiting cases. In particular:

- In the limiting case that $K_1 = \emptyset$, we are taking no notice of the syntactic features of any of the formulae in K . Nothing is protected, the finest K_1 -protecting splitting of K is just the finest splitting of K , and thus epistemic relevance coincides with Parikh's canonical relevance (Definition 4.2).
- At the other end of the spectrum, when $K_1 = K$, we are taking notice of the syntactic formulation of every formula in K , and epistemic relevance coincides with Rodrigues' path-relevance modulo (Definition 2.1).

Summary of this section: Thus the introduction of a parameter into the central definitions provides sufficient flexibility to represent extra-logical sources of

relevance. In particular, it allows the believing agent to specify whether there are any special formulae in the belief set that create connections between letters, beyond the purely logical ones. The extent to which epistemic relevance goes beyond canonical relevance depends on how much of its belief set the agent puts into the set K_1 .

Appendices

Appendix 1: Table of Kinds of Relevance Discussed in this Paper

Name	Arguments	Syntax-independent?
syntactic relevance	formulae	no
essential relevance		yes
path-relevance	formulae plus belief set K	except in K
cell-relevance		
canonical (path/cell) relevance		yes
parametrized relevance	formulae plus belief set K plus relation R over letters	yes
epistemic relevance	formulae plus belief set K plus subset $K_1 \subseteq K$	except in K_1

Appendix 2. Proof of Theorem 4.1

Theorem 4.1. Canonical path-relevance is equivalent to canonical cell-relevance. In detail: let a, b be formulae of classical propositional logic, and let K be a set of formulae serving as a belief set. Then a is canonically path-relevant to b (mod K) iff it is canonically cell-relevant to b (mod K).

Proof. Left to right: Suppose that a is canonically path-relevant to b (mod K). Then there is a finite sequence x_0, \dots, x_{n+1} ($n \geq 0$) of formulae with $x_0 = a^*$, $x_{n+1} = b^*$, all of $x_1, \dots, x_n \in K^\#$, and each x_i sharing at least one letter with x_{i+1} . Let p be a letter shared by $x_0 = a^*$ and x_1 , and let q be a letter shared by x_n and $x_{n+1} = b^*$. Since all of $x_1, \dots, x_n \in K^\#$, and each x_i shares at least one letter with x_{i+1} , it follows that all of the letters in x_1, \dots, x_n come from the same cell E_i of the finest splitting of K . Thus in particular p and q come from the same cell E_i , so each of the sets $E(a^*) \cap E_i$ and $E(b^*) \cap E_i$ is non-empty as required for canonical cell-relevance.

Right to left: Suppose that a is canonically cell-relevant to $b \pmod{K}$. Then there is a cell E_i of the finest splitting \mathbf{E} of K such that each of the sets $E(a^*) \cap E_i$ and $E(b^*) \cap E_i$ is non-empty. By the former, there is a letter $p \in E_i$ occurring in a^* . Since $p \in E_i$ it occurs in some formula $x \in B_i \subseteq K^\#$ (otherwise the splitting would not be finest, as we could further split E_i into $\{p\}$ and $E_i \setminus \{p\}$). Likewise by the second there is a letter q occurring in a^* and in some formula $y \in B_i \subseteq K^\#$. We need to show that there are x_1, \dots, x_n in $K^\#$ with $x = x_1$, $y = x_n$ and each x_i sharing a letter with x_{i+1} . But this must hold because otherwise we could take the closure $\{x\}^+$ of $\{x\}$ under the relation of sharing a letter, to split E_i further into $E(\{x\}^+)$ and $E_i \setminus E(\{x\}^+)$.

Appendix 3. Derivation of Theorem 5.1 (turnstile version) from its epsilon counterpart

Theorem 5.1 states that if we apply AGM contraction to the finest splitting $K^\#$ of a consistent belief set K , rather than to K itself, then it always respects relevance. In other words, whenever $K^\# \vdash x$ but $K^\# - a \not\vdash x$ then a is canonically relevant to $x \pmod{K}$. In Kourousias and Makinson (2007) this was proven in an ‘epsilon version: whenever $x \in K^\#$ but $x \notin K^\# - a$ then a is canonically relevant to $x \pmod{K}$.

To derive the turnstile version of the theorem from the epsilon one, it suffices to show that for AGM contraction on a consistent belief set K , respect for relevance (epsilon version) implies respect for relevance (turnstile version).

Assume the epsilon version. Suppose that K is consistent, $K^\# \vdash x$, $K^\# - a \not\vdash x$; we need to show that a is canonically relevant to $x \pmod{K}$. Since $K^\# \vdash x$ we have $K^\# \vdash x^*$, so there are $a_1, \dots, a_k \in K^\#$ with $a_1 \wedge \dots \wedge a_k \vdash x^*$. Since K is consistent, we may assume without loss of generality that each a_i^* shares a letter with x^* . Since $K^\# - a \not\vdash x$ we likewise have $K^\# - a \not\vdash x^*$, so there is an $i \leq k$ with $K^\# - a \not\vdash a_i$, so that $a_i \notin K^\# - a$. By the epsilon version of the theorem, a is canonically relevant to $a_i \pmod{K}$. That is, there is a cell E_j of the finest partition \mathbf{E} of K such that each of the sets $E(a^*) \cap E_j$ and $E(a_i^*) \cap E_j$ is non-empty. Since $a_i \in K^\#$ all the letters of a_i come from the same cell, so $E(a_i) \subseteq E_j$. But $E(a_i^*) \subseteq E(a_i)$, and so also $E(a_i^*) \subseteq E_j$. Since a_i^* shares a letter with x^* , this tells us that $E(x^*) \cap E_j$ is non-empty. Putting together the non-emptiness of $E(a^*) \cap E_j$ and of $E(x^*) \cap E_j$ we may conclude that that a is canonically relevant to $x \pmod{K}$ as desired, and the proof is complete.

It is also possible to prove Theorem 5.1 (turnstile version) directly, essentially by including the above considerations within a re-run of the proof of the epsilon version.

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.

Kourousias, George and David Makinson 2007. “Parallel interpolation, splitting, and relevance in belief change. To appear in *The Journal of Symbolic Logic* 2007: 994-1002.

Makinson, David “Friendliness and sympathy in logic” in *Logica Universalis*, 2nd edition, ed. J.-Y. Beziau (Basel: Birkhauser Verlag, 2007).

Makinson, David and George Kourousias, “Respecting relevance in belief change”. To appear in *Análisis Filosófico* 2007 (special issue commemorating the tenth anniversary of the death of Carlos Alchourrón).

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

Rodrigues, Odinaldo Teixeira 1997. “A methodology for iterated information change” PhD thesis, Imperial College London,
http://www.dcs.kcl.ac.uk/staff/odinaldo/research/phd_thesis.pdf.

Riana, Joselyto and Renata Wassermann 2004. “Using relevance to speed up inference: some empirical results”, in Bazzan and Labidi eds *SBIA 2004*, LNAI 3171: 21-30. Berlin: Springer-Verlag.

Wassermann, Renata 1999. “Resource-bounded belief revision”, PhD thesis, University of Amsterdam, <http://www.ime.usp.br/~renata/Thesis.html>.

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