David C. Makinson

Propositional relevance through letter-sharing

Article (Accepted version) (Refereed)

Original citation:
Makinson, David C. (2009) *Propositional relevance through letter-sharing*. Journal of Applied Logic, 7 (4). pp. 377-387. ISSN 1570-8683 DOI: 10.1016/j.jal.2008.12.001

Propositional Relevance through Letter-Sharing

David Makinson

Abstract

The concept of relevance between classical propositional formulae, defined in terms of letter-sharing, has been around for a 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: the relation of relevance was considered modulo the choice of a background belief set, and the belief set was put into a canonical form, called its finest splitting. In the first part of this paper, we recall the ideas of Rodrigues and Parikh, and show that they yield equivalent definitions of what may be called canonical cell/path relevance. The second part presents the main new result of the paper: while the relation of canonical relevance is syntax-independent in the usual sense of the term, it nevertheless remains language-dependent in a deeper sense, as is shown with an example. The final part of the paper turns to questions of application, where we present a new concept of parameter-sensitive relevance that relaxes the Rodrigues/Parikh definition, allowing it to take into account extra-logical sources as well as purely logical ones.

Keywords: relevance, letter-sharing, belief change, splitting, language-dependence.

From Syntactic to Canonical Cell/Path Relevance

1. Logical Relevance as a Two-Place Relation between Formulae

Attempts to give formal expression to the notion of relevance between propositional formulae go back at least to Belnap (1960), who suggested that a necessary, but not sufficient, condition for one formula to be relevant to another is that they share some elementary letter. We call this *syntactical relevance*.

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

In the same paper, Belnap went on to propose that relevance of antecedent to consequent should serve as an adequacy condition for any acceptable entailment relation in propositional logic. While classical logic fails syntactic relevance, his subclassical logic E (for 'entailment') satisfies that formal condition, as do a number of other subsystems of classical logic that came to be known as 'relevance logics'.

The present paper is not at all concerned with such relevance logics, and we have no desire to weaken the classical one. We are interested in the concept of relevance itself. Our purpose is to see how far the simple idea of letter-sharing may be developed into a well-behaved formal account of relevance in classical propositional contexts, and examine its application to the theory of belief change.

Poor behaviour of syntactic relevance. (1) The relation 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'. Moreover (2), for any a, b there are a', b' to which they are respectively classically equivalent, with a' syntactically relevant to b'.

Example 1.1. For (1), the formula $\neg p\wedge(\neg p\vee q)$ is syntactically relevant to q, but the former is classically equivalent to $\neg p$ which is not relevant to q when these letters are chosen as distinct. For (2), just put e.g. $a' = a\vee(r\wedge\neg r)$ and $b' = b\vee(r\wedge\neg r)$.

Notation. We are using $p,q,\ldots$ for elementary letters, $a,b,..,x,y,\ldots$ for arbitrary formulae, $A,B,..,K,\ldots$ for sets of formulae, and $\vdash$, $\dashv\vdash$ for the relations of classical propositional consequence and equivalence. Classical consequence as an operation is written Cn, with $Cn(A) = \{x\colon A\vdash x\}$ as usual. A formula is called contingent iff it is neither a tautology nor a contradiction.

To overcome the shortcomings of the syntactic notion, an obvious first move is to express each formula in its least letter-set, using the well-known *least letter-set theorem*:

Theorem 1.1. For every set A of formulae, there is a unique least set of elementary letters such that A may equivalently be expressed using only letters from that set.

Example 1.2. The unique least letter-set of $\neg p\wedge(\neg p\vee q)$ is $\{p\}$ since $\neg p\wedge(\neg p\vee q)$ is equivalent to $\neg p$. On the other hand, the unique least letter-set of $\neg p\wedge(\neg r\vee q)$ is $\{p,q,r\}$, as the formula is not equivalent to any other formula lacking any of those letters.

Comments on least letter-set theorem. This theorem should figure in every textbook of elementary logic, but in fact is rarely so much as mentioned. We recall:

- Strictly speaking, it holds in this simple form only when the language has a primitive zero-ary operator (propositional constant) such as the falsum $\perp$. In such a language, the least letter-set of any tautology or contradiction is $\varnothing$. Without a zero-ary connective, say with just $\neg,\wedge,\vee$, tautologies and contradictions have many minimal letter-sets (in fact, all the singleton letter sets), but no least one (since no formula is bereft of letters). For simplicity of formulation, in this paper we work with the falsum.

- Intuitively, the least letter-set theorem is just what anyone would expect, but it 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 e.g. the appendix of Makinson (2007).

- Letters in the least letter-set of A are said to be *essential* (to A) or *irredundant* (in A), those outside are called *inessential* or *redundant* (in A).

- With this terminology, the theorem may be stated in another manner: the set of all letters separately redundant in A, is jointly redundant in A, in the sense that A may equivalently be expressed without any of them.

When a formula a has no redundant letters, i.e. when all letters occurring in a are in its least letter-set, we say that it is in *least letter-set form*. It is convenient to use a choice function, writing a^* for an arbitrarily chosen formula in least letter-set form that is equivalent to a. The formula a^* is called a *least letter-set version* of a. Likewise for sets A of formulae.

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 1.3. 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 of essential relevance

- It is syntax-independent in the usual sense: when a, b are tautologically equivalent to a', b' respectively, then a is essentially relevant to b iff a' is essentially relevant to b' (immediate from definition).
- No two distinct elementary letters are relevant to each other (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 (given the presence of the falsum in our language).
- Not transitive. *Example*: p is essentially relevant to $p\wedge q$ which is so to q, but p is not to q.
- Cannot be ‘made transitive’: its transitive closure makes any two contingent formulae relevant to each other. *Verification*: Take contingent a,c. Since they are contingent, a^* and c^* contain letters p and 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.

This is all part of the folklore and well documented in the literature (see Appendix A). However, things began to take a fresh turn in the late 1990s, when a few people began thinking about relevance in the light of formal accounts of belief change. Two basic insights emerged. The first was that in that 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 choice of a background belief set*. The second was that this belief set may be given a canonical form known as its *finest splitting*.

As these developments are not very widely known, we explain and comment on them in the following two sections. To help the reader keep track of successive definitions, Appendix B contains a table of all the different kinds of relevance examined in the text.

2. Path-Relevance Modulo a Belief Set

Consider any three distinct elementary letters p, q, r. They are not essentially relevant to each other. Now consider the belief set $K = \{p \rightarrow q, q \rightarrow r\}$. Then it is natural to say that *from the point of view of* K, the letter p is relevant to q, q is relevant to r, and p is indirectly relevant to r. This suggests the following definition.

Definition 2.1. (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 $x_0,\ldots x_{n+1}$ ($n \geq 0$) of formulae with $x_0 = a^*$, $x_{n+1} = b^*$, $x_1,\ldots,x_n \in K$, and each x_i shares at least one letter with x_{i+1}.

Comments. Note that $x_1,\ldots,x_n$ 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). We do not require that K is closed under consequence (though it may be): a belief set is understood to be an arbitrary set of formulae of propositional logic.

History. Essentially this notion was introduced by Rodrigues in his thesis (1997), Appendix A, definition 8.14. It was also used by Renata Wassermann in her thesis (1999) and in subsequent papers e.g. Riana and Wassermann (2004). Actually, these authors took a,b instead of a^*,b^* in the definition, but we have refined it to ensure that it is syntax-independent in those two arguments.

Path-relevance generalizes essential relevance in a natural way: the latter amounts to the case $n = 0$ in Definition 2.1. Like essential relevance, path relevance is:

- Syntax-independent in a,b, symmetric, almost reflexive (in the same sense), not transitive (even when $n > 0$).

But with K as parameter, new features emerge. One is rather positive:

- Distinct elementary letters can be relevant to each other (mod K). *Example*: Modulo $K = \{p \rightarrow q, q \rightarrow r, \neg s\}$, p is path-relevant to r but not to s.

However, some other features are rather undesirable:

- 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. That is, when $K = Cn(K)$, any two contingent formulae a,b are path-relevant to each other modulo K. *Reason*: 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^* and note that $Cn(K)$ contains any tautology in these letters, e.g. $(p \vee \neg p) \vee (q \vee \neg q)$.

Can we get around these unpleasant features? One might try tweaking Definition 2.1 by replacing K by its least letter-set version K^*. However, this does nothing to eliminate syntax-dependence in K. *Example*: Both $K = \{p \wedge q\}$ and the equivalent $J =$

$\{p,q\}$ are already in least letter-set form, but under Definition 2.1 we have p path-relevant to q modulo K, but not so modulo J.

A better idea is needed, and one was provided by Rohit Parikh in 1999. As well as minimizing the set of elementary letters, we need to *disentangle* them. The formulae in the background belief set K need to be 'combed out' so that their letters are not mixed up with each other more than necessary. In other words, we need to render K as modular as possible. Parikh made this idea precise with his concept of *the finest splitting* of a belief set.

3. Splittings and Finest Splittings of a Belief Set

We begin with the definition of a splitting of a belief set K, and then pass to that of a finest splitting. *Notation*: We write $E(K)$ for the set of all elementary letters occurring in formulae of K, and $E_0(K)$ to be the least letter-set of K, i.e. $E_0(K) = E(K^*)$.

Definition 3.1. (Parikh) Let K be a contingent belief set, expressed in the language of classical propositional logic (with a zero-ary connective). Let $\mathbf{E} = \{E_i\}_{i\in I}$ be a partition of its least letter-set $E_0(K)$, which by contingency will be non-empty. 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. (1) 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. (2) The partitions of a set can be put in one-one correspondence with the equivalence relations over the set. (3) 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, iff every cell of the first partition is a subset of a cell of the second one. (4) Given any non-empty family of partitions of a set, the intersection of all the equivalence relations associated with partitions in the family is itself an equivalence relation over the set, and so corresponds to a partition of the set. With respect to the fineness relation, it is the infimum (alias greatest lower bound or glb) of the family of partitions.

Comments on the definition of splitting. (1) A splitting of K is thus a special kind of partition of the least letter-set of K; it is not a partition of K itself. (2) While each $E(B_i) \subseteq E_i \subseteq E_0(K)$ it is not required that the sets $B_i \subseteq K$, although their union $\cup\{B_i\}_{i\in I}$ must be classically equivalent to K. (3) Since each $E(B_i) \subseteq E_i$ and the E_i are pairwise disjoint, the B_i must be 'almost' pairwise disjoint, in the sense that they share no formulae containing elementary letters. (4) Since $E_0(K)$ is the least letter-set of K and K is assumed contingent, it follows that in a splitting each B_i is non-empty and in fact $E(B_i) = E_i$. (5) This definition (and all those that follow) may be extended to cover the limiting cases that K is inconsistent or tautologous, but at the cost of limiting-case clauses in definitions, theorems and proofs that distract from the main ideas.

History. Actually, Parikh (1999) defined splittings of K for any letter-set $E \supseteq E_0(K)$, e.g. E could be $E(K)$ or the set $E(L)$ of all letters of the language. However, it simplifies formulations to fix it at $E_0(K)$. For instance, when $E \supset E_0(K)$ then comment (4) above can fail at the edges. Example: Put $K = \{p, q\vee\neg q\}$, $E = E(K) = \{p,q\}$, $B_1 = \{p\}$ but B_2 can be $\{\neg\bot\}$ or $\varnothing$ so that $E(B_2) \subset E_2 = \{q\}$.

Example 3.1. $K = \{p\rightarrow\neg q, \neg q\rightarrow r, p\vee s, \neg s, (r\rightarrow t)\vee(t\rightarrow r)\}$.

- $E(K) = \{p,q,r,s,t\}$ but t is redundant, so $E_0(K) = \{p,q,r,s\}$.
- The coarsest splitting of K is evidently the singleton partition of $E_0(K)$ with $E_0(K) = \{p,q,r,s\}$ itself as the only cell, so that $\cup\{B_i\}_{i\in I} = B_1 = \{p\rightarrow\neg q, \neg q\rightarrow r, p\vee s, \neg s\} \dashv\vdash K$. But we can do better than that.
- A less coarse splitting of K partitions E into two cells $E_1 = \{p,q\}$ and $E_2 = \{r,s\}$, taking $B_1 = \{p,\neg q\}$, $B_2 = \{r,\neg s\}$, so that $K \dashv\vdash B_1\cup B_2$.
- The finest splitting of K partitions E into four singleton cells $\{p\}$, $\{q\}$, $\{r\}$, $\{s\}$ with $B_1 = \{p\}$, $B_2 = \{\neg q\}$, $B_3 = \{r\}$, $B_4 = \{\neg s\}$, so that $K \dashv\vdash B_1\cup\ldots\cup B_4$.

In this very simple example, the finest splitting of E has singleton cells and the associated sets B_1 to B_4 consist of literals. Of course, neither of these features need hold. For instance, take $K = \{(p\rightarrow q)\wedge(r\rightarrow s)\}$. Its finest splitting is into the two-element cells $\{p,q\}$, $\{r,s\}$, with $B_1 = \{p\rightarrow q\}$, $B_2 = \{r\rightarrow s\}$ consisting of non-literals.

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

History. Theorem 3.1 was established for the finite case by Parikh (1999). It was extended to the infinite case by Kourousias and Makinson (2007), using a new form of interpolation called ‘parallel interpolation’. Both parallel interpolation and the finest splitting theorem may be extended to first-order logic.

Comments on the theorem. Strictly speaking, it is the finest splitting $\mathbf{E} = \{E_i\}_{i\in I}$ of elementary letters that is unique. Given such a family, there will evidently be many families $\{B_i\}_{i\in I}$ with $\cup\{B_i\}_{i\in I} \dashv\vdash K$ and $E(B_i) \subseteq E_i$. However, since in fact each $E(B_i) = E_i$, the different ways of choosing a given B_i do not affect its letters. Moreover, it turns out that:

Observation 3.2. For contingent K the B_i associated with the finest splitting of K are unique up to tautological equivalence. That is: let K be a contingent belief set, and $\mathbf{E} = \{E_i\}_{i\in I}$ its finest splitting. Suppose both $K \dashv\vdash \cup\{B_i\}_{i\in I}$ and $K \dashv\vdash \cup\{B_i'\}_{i\in I}$ where $E(B_i) \subseteq E_i$ and $E(B_i') \subseteq E_i$. Then each $B_i \dashv\vdash B_i'$.

Sketch of proof. This follows from the fact that the B_i are pairwise disjoint. For the details, see Appendix C.

Again, it simplifies formulations if we use a choice function:

Definition 3.2. For contingent K, write $K^{\#}$ for $\cup\{B_i\}_{i\in I}$ for some particular such family $\{B_i\}_{i\in I}$. We abuse terminology a little by *also calling $K^{\#}$ the finest splitting* of K.

Comments on the definition of $K^{\#}$. (1) Clearly, when $K_1 \dashv\vdash K_2$ then $K_1^{\#} = K_2^{\#}$. (2) Keep in mind that the family $\{B_i\}_{i\in I}$ is not necessarily a partition of K, but is formed from a certain partition of its least letter-set $E_0(K)$. (3) Note that when the finest splitting $\mathbf{E} = \{E_i\}_{i\in I}$ of K has at least two cells, then $K^{\#}$ cannot be closed under classical consequence –the conjunction of any two formulae from different cells will be in $Cn(K^{\#})$ but cannot be in $K^{\#}$. Even when there is only one cell, $K^{\#}$ need not be closed under consequence. (4) As we have defined it, the finest splitting $K^{\#}$ is always a least letter-set version of K. However, a least letter-set version K^* of K need not be a finest splitting, as in the following simple example.

Example 3.2. Put $K = \{p\wedge q\}$. Then K is already in a least letter-set form, since there is no equivalent set of formulae in fewer letters. But it is not in a finest splitting form, since $K \dashv\vdash \{p\}\cup\{q\}$, which partitions $E(K)$ into two singleton cells.

Indeed, if we take a least letter-set version of a belief set K and tangle the letters up in any way we like, then so long as we keep it equivalent to K and do not introduce fresh letters, we are still in the least letter-set but can be far from a finest splitting.

4. Using Finest Splittings to Define Canonical Relevance (Modulo a Belief Set)

How can finest splitting help make the notion of relevance modulo a belief set fully syntax-independent in K as well as in a, x? We may see the finest splitting $K^{\#} = \cup\{B_i\}_{i\in I}$ of K 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 and at the same time (under our definition) eliminating redundant letters. We can then refine Rodrigues' notion of path-relevance by taking the path through this canonical representation $K^{\#}$ instead of through K itself. Thus, replacing $x_1,\ldots,x_n \in K$ by $x_1,\ldots,x_n \in K^{\#}$ in Definition 2.1, we have the following:

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

Features of canonical path-relevance modulo K:

- This time $x_1,\ldots,x_n$ 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).
- As desired, path-relevance becomes syntax-independent in the usual sense that it is invariant under logical equivalence in argument K as well as in a, b. This follows from the fact, noted in the comments after Definition 3.2, that equivalent belief sets have the same finest splitting $K^{\#}$.
- Like plain path-relevance, it is symmetric but not transitive (in the arguments a,b with K fixed); almost reflexive (in the same sense as before); distinct elementary letters can be relevant to each other (mod K).

There is another way of arriving at the same concept. 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. (Parikh 1999). Let a,b be formulae of classical propositional logic and K be a contingent 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 either a^* shares some letter with b^*, or 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 either $E(a^*) \cap E(b^*)$ is non-empty, or 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 of the principal case of the definition, 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 $E(a^*)$, $E(b^*)$ respectively. Thus a is canonically cell-relevant to b (mod K). However, if $E(b^*)$ consisted of just t,u then a would not be canonically cell-relevant to b (mod K).

History. Actually, Definition 4.2 is implicit rather than explicit in Parikh (1999). Moreover, both that paper and Kourousias and Makinson (2007) use $E(a)$, $E(b)$ rather than $E(a^*)$, $E(b^*)$.

Surprisingly, the following equivalence does not appear to have been noticed in the literature.

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 contingent 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).

Sketch of proof: Left to right, the essential idea is that there can be no paths across cells. Right to left, paths must span the cells. For details, see Appendix C.

Remark on the theorem. The first disjunct in the definition of canonical cell relevance corresponds to the case $n = 0$ in the definition of canonical path relevance.

Summary of the story so far. By using Parikh's notion of the finest splitting of K, we can refine Rodrigues' account of path-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 one of canonical cell-relevance. The equivalence suggests robustness of the concept, which henceforth we call simply *canonical relevance*.

Warning. Canonical relevance depends only on the logical power of K, but is not monotonically increasing in that power. When $K_1 \vdash K_2$ then it does not follow that if a is canonically relevant to b modulo K_2 then it is so modulo K_1. *Example*: Put $K_1 = \{p\wedge q\}$, $K_2 = \{p\to q\}$. Then $K_1 \vdash K_2$ and p is canonically relevant to q modulo K_2, but is not so modulo K_1.

5. Syntax-Independent – but Still Language-Dependent

As remarked in the preceding section, canonical relevance is syntax-independent in the usual sense that it is invariant under logical equivalence in all three of its arguments a,b,K. However, there is also a sense in which it is still not fully language-independent. Its definition, whether via paths or cells, gives a privileged place to elementary letters over compound formulae. As a result, it turns out that whether one state of affairs is relevant to another depends on how we deploy elementary letters in representing them. This is the first main new result of the paper.

Example 5.1. Consider the belief set $\{p\leftrightarrow q, (p\leftrightarrow q)\leftrightarrow p\}$. Clearly, p is not canonically relevant to $(p\leftrightarrow q)\leftrightarrow p$ modulo the belief set, since $p = p^*$, $((p\leftrightarrow q)\leftrightarrow p)^* = q$, and the belief set is equivalent to $\{p,q\}$. Now suppose we represent our states of affairs in a different manner. Noting that the formulae p and $p\leftrightarrow q$ are logically independent of each other (all four combinations of their truth and falsity are possible), we might let the letter p continue to represent the same state of affairs as before, but use the letter q to stand for the one previously denoted by $p\leftrightarrow q$. Then the formula $q\leftrightarrow p$ stands for the state of affairs previously represented by $(p\leftrightarrow q)\leftrightarrow p$. So under our second representation scheme, the belief set is written as $\{q, q\leftrightarrow p\}$, which is again equivalent to $\{p,q\}$. We now ask whether p (corresponding to the old p) is canonically relevant to $q\leftrightarrow p$ (corresponding to the old $(p\leftrightarrow q)\leftrightarrow p$) modulo the new belief set. This time the answer is trivially positive (modulo any belief set) since $p^* = p$ shares a letter with $(q\leftrightarrow p)^* = q\leftrightarrow p$. Changing the way in which we represent states of affairs has thus changed the answer to our question!

In case this looks like a sleight-of-hand, we review Example 5.1 more formally. First, we give the construction itself.

Observation 5.1. Let L be the propositional language generated by the letters $\{p,q\}$, and define $f: L \to L$ by putting $f(p) = p$, $f(q) = p\leftrightarrow q$, and homomorphic for compound formulae. Then there is a bijection φ between valuations on the language such that for all formulae $a \in L$, $v(a) = v'(f(a))$, where $v': L \to \{0,1\}$ is the counterpart $\varphi(v)$ of $v: L \to \{0,1\}$.

The homomorphism condition means of course that $f(\neg a) = \neg f(a)$, $f(a\wedge b) = f(a)\wedge f(b)$, $f(a\vee b) = f(a)\vee f(b)$, $f(\bot) = \bot$. The good behaviour of the formula-homomorphism f with respect to the valuation-bijection φ gives mathematical content to the intuitive idea that f does not alter semantic structure; more specifically, that the formula a

represents (under v) the same 'state of affairs' as $f(a)$ does (under v'). For a proof of Observation 5.1, see Appendix C.

Now checking for what is relevant to what, we see (Figure 5.1):

- On the one hand, trivially the formula p is canonically relevant to $q\leftrightarrow p$ modulo the belief set $K = \{q, q\leftrightarrow p\}$ (or any other) since $p^* = p$ shares a letter with $(q\leftrightarrow p)^* = q\leftrightarrow p$.

- On the other hand, $f(p) = p$ is not canonically relevant to $f(q\leftrightarrow p) = (p\leftrightarrow q)\leftrightarrow p$ modulo $f(K) = f(\{q, q\leftrightarrow p\}) = \{p\leftrightarrow q, (p\leftrightarrow q)\leftrightarrow p)\}$, since $(f(p))^* = p^* = p$ while $(f(q\leftrightarrow p))^* = ((p\leftrightarrow q)\leftrightarrow p)^* = q$ and $(f(K))^\# = \{p,q\} = K^\#$.

Figure 5.1. Example of language-dependence of canonical relevance modulo K

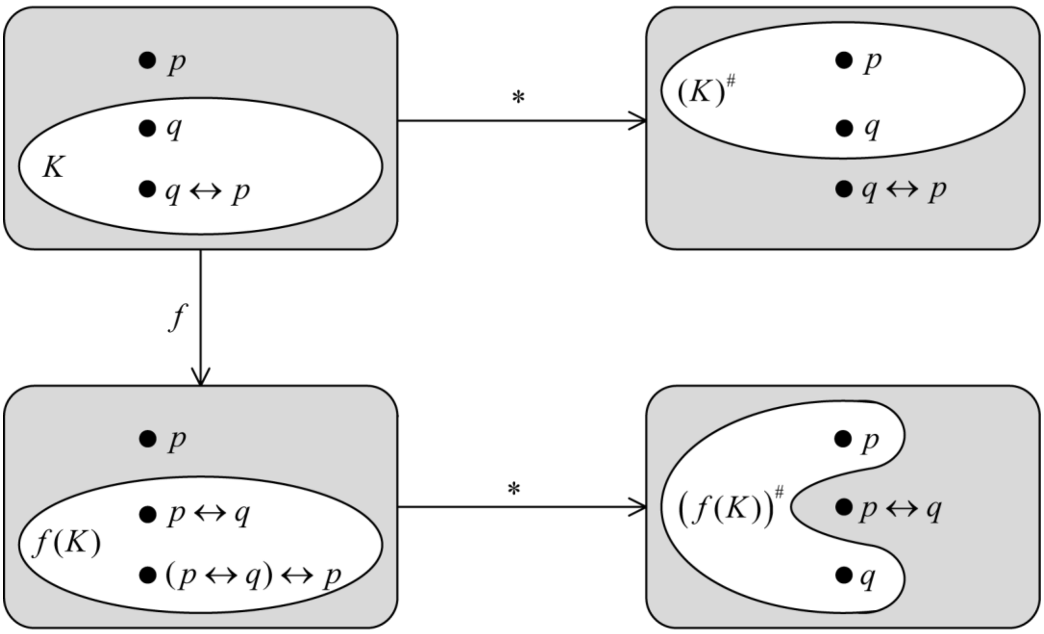

Thus, while the concept of canonical relevance is *syntax-independent* as usually understood, i.e. invariant under logical equivalence in all three of its arguments, it nevertheless remains *language-dependent*, in the deeper sense that it is not invariant under different representations of the same state of affairs – even when the representations are in the same language. To this extent it is not entirely semantic, retaining a residual syntactic element that may be difficult or impossible to eliminate.

The phenomenon does not appear to have been discussed in the literature, but may be of some importance. It goes well beyond the problem of relevance. Any concept whose definition gives a privileged role to elementary letters (or in the case of predicate logic, atomic formulae) is likely to be language-dependent in the same way. This seems to be the case, for example, with certain concepts that have been used in artificial intelligence to define particular forms of nonmonotonic reasoning, notably the closed world assumption and circumscription.

How should the language-dependence of the relation of canonical relevance be appreciated? Two contrasting attitudes suggest themselves.

- It may be felt that in view of this feature, canonical relevance is not much better behaved than its less sophisticated predecessors, which were seen to be syntax-dependent in one or more of their arguments a, x, K. For this reason it should simply be abandoned (along with all other language-dependent concepts such as circumscription).

- On the other hand, it may be felt that language-dependent notions (and perhaps even some syntax-dependent ones) do have their legitimate uses, particularly in computational contexts.

Without taking a definite stance on this delicate question, we note that canonical relevance has very interesting interactions with operations of belief change, which we now examine.

6. Respecting Relevance in Belief Change

How far do operations of belief change in the manner of Alchourrón, Gärdenfors and Makinson (1985), briefly AGM, respect relevance? We begin by reviewing the state of play, focussing on the operation of contraction (thus leaving aside revision) and omitting all proofs (which can be found in Kourousias and Makinson (2007)).

Definition 6.1. We say that an operation $-$ of contraction on a contingent belief set K *respects canonical relevance* (briefly, when no ambiguity is possible, *respects relevance*) iff whenever $K \vdash x$ but $K-a \nvdash x$ then a is canonically relevant to x (mod K). Contrapositively, when $K \vdash x$ and a is canonically irrelevant to x (mod 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 $K \vdash x$ iff $x \in K$ and likewise $K-a \vdash x$ iff $x \in K-a$. In this situation, Definition 6.1 is equivalent to one with epsilon replacing turnstile: whenever $x \in K$ and a is canonically irrelevant to x (mod K) then still $x \in K-a$.

Observation 6.1. (Parikh 1999): AGM contraction can fail to respect relevance, and this can happen independently of whether K is closed under consequence.

Example 6.1. 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, because we can split $E = \{p,q\}$ into $E_1 = \{p\}$, $E_2 = \{q\}$ with $K^{\#} = \{p\} \cup \{q\}$.

The example is robust in the sense that it goes through 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 6.2 (Kourousias and Makinson 2007). If we apply AGM contraction to the finest splitting $K^{\#}$ of a contingent belief set K, rather than to K itself, then it respects relevance.

Example 6.2. When given $K = Cn(p,q)$ or $K_0 = \{p\leftrightarrow q,q\}$ the theorem instructs us to apply the contraction operation to the canonical belief set $K^{\#} = K_0^{\#} = \{p,q\}$. Since there is just one maximal p-nonimplying subset of $K^{\#}$, namely $\{q\}$, there is just one possible output for an AGM belief contraction $K^{\#}-p$, namely $\{q\}$.

Comments. (1) Actually, the observation of Parikh (1999) was made for AGM revision, but the counterexamples for revision and contraction are essentially the same. (2) Theorem 6.2 was established by Kourousias and Makinson (2007) for the epsilon version of Definition 6.1, rather than the turnstile version. When a belief set is not closed under classical consequence (as in the case of $K^{\#}$) the two versions 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 C.

7. Should Canonical Relevance always be Respected?

Of course, we may ask whether eliminating canonically irrelevant formulae really is a shortcoming for a belief contraction operation. Assuming that canonical relevance modulo a belief set is itself a reasonable notion to work with (despite its language-dependence, already noted) we may still ask: is failure to respect it a defect, or just a feature, of AGM contraction?

It appears that the answer depends on whether we want our contractions to take into account only formal considerations, or also epistemic ones. To see this, consider again the example where we wish to contract the belief base $K_0 = \{p\leftrightarrow q,q\}$ by p.

We know that $K^{\#} = \{p,q\}$, so that while $p\leftrightarrow q$ is canonically relevant to p modulo K_0, q is not. So if the contraction is to respect relevance, it will eliminate $p\leftrightarrow q$, but not q. But it may 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 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 within the belief set.

In general, when a belief set is presented by a base, we may have differing attitudes towards the propositions in the base. They may be there by happenstance, and any other base might be deemed as just as appropriate so long as it is equivalent (and perhaps satisfies general requirements such as being computable or schematic). But other propositions may be in the base because we *want* them to be there; they may have an epistemic priority over items outside the base. Even within the base, some elements may have priority over others. From this perspective, taking epistemic matters into account, we will not need to respect canonical relevance.

8. Parameter- Sensitive Relevance

If we are interested in epistemic factors in belief change, we may well wish to develop the concept of canonical relevance to take account of them. How could we go about it? Of course, logic alone cannot specify which propositions have what epistemic status. But it can introduce into its constructions parameters that allow such specifications to play a role. In this section we introduce such a parameter. This is the second main new construction of the paper.

Definitions and theorems correspond to earlier unparametrized ones, and are numbered by their counterparts with a plus sign. We begin by observing that the finest splitting theorem may be strengthened to cover an *arbitrary family* of splittings, rather than just the family of *all* splittings of K.

Theorem 3.1^{+}. Let K be any contingent 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.

Next, we notice that the concept of canonical cell/path relevance, which was introduced in Definition 4.2 using the *finest* splitting of K, generalizes without change with respect to an *arbitrary* splitting. In terms of cells, for instance, we have:

Definition 4.2^{+}. Let a,b be formulae of classical propositional logic and K be a contingent set of formulae serving as a belief set with $\mathbf{E} = \{E_i\}_{i\in I}$ any splitting of K. We say that a is *relevant* to b (mod $\mathbf{E}$) iff a^* shares some letter with b^*, or 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.

The notion of respect for relevance, introduced Definition 6.1, similarly generalizes:

Definition 6.1^{+}. Let K be a contingent 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 $K-a \nvdash x$ then a is relevant to x modulo $\mathbf{E}$.

With these generalized definitions available, we can now introduce a parameter R to handle extra-logical (and in particular, epistemic) sources of relevance. R is a relation between elementary letters, permitting us to stipulate that certain letters are epistemically relevant to others.

Definition 8.1. Let K be a contingent set of formulae serving as a belief set, with E its least letter-set. Let R be any relation between letters in E. 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 $\mathbf{E}$.

We may now apply Theorem 3.1^+ to the family of all R-protecting splittings of K:

Corollary to Theorem 3.1^+. Let K be a contingent set of formulae serving as a belief set, with E its least letter-set. Let R be any relation between letters in E. Then K has a (unique) finest R-protecting splitting.

Proof. Since $R \subseteq E^2$, there is at least one splitting that protects R, namely the coarsest (one-cell) splitting. Hence by Theorem 3.1^+, the infimum of all R-protecting splittings of K is a splitting of K, and it is immediate that it protects R.

The relation of relevance modulo the finest R-protecting splitting of K (rather than modulo its finest splitting) may naturally be referred to as *R-sensitive* relevance. When the pairs in R represent declarations of epistemic connections between letters, it may be thought of as representing *epistemically sensitive relevance*.

The extent to which R-sensitive relevance goes beyond canonical relevance evidently depends on how much is put into the protected relation R. In the limiting case that R is empty, the two coincide; in the other limiting case that R contains all pairs of letters, we get the one-cell partition of $E = E_0(K)$ and so end up with Rodrigues' path-relevance without splitting as in section 2.

We end by noting that the same proof as for Theorem 6.2 gives us more generally:

Theorem 6.2^+. Let K be a contingent set of formulae serving as a belief set, with E its least letter-set. Let R be any relation between letters in E. If we apply AGM contraction to the finest R-protecting splitting rather than to K itself, then it respects relevance modulo that same splitting.

Appendices

Appendix A: Literature on Relevance as a Two-Place Relation

For an overview and extended bibliography of work on propositional relevance as a two-place relation between formulae, see Lang et al (2003). This paper gives particular attention to computational questions. Although the authors mention the seminal paper Parikh (1996) in passing, they do not investigate relevance modulo the finest splitting of a background belief set. In their treatment of the notion of an essential letter they follow Ryan (1991) in dividing the concept into two parts, thereby giving it a polarity. Expressed in the manner of the present paper, we may say that a formula *a sometimes depends on a positive value* for the letter p iff there is a valuation v with $v(a)$ =1 but $v_{p=0}(a) = 0$, where $v_{p=0}$ is the valuation that agrees with v on all letters other than p but gives p the value 0. Likewise, *a sometimes depends on a negative value* for the letter p iff there is a valuation v with $v(a)$ =1 but $v_{p=1}(a) = 0$. Evidently, the two kinds of dependence do not exclude each other. As Lang et al (2003) observe, it is immediate that the essential letters of a formula are just those on which it sometimes depends either positively or negatively.

Appendix B: Table of Kinds of Relevance Discussed

<table>
<tr><th>Name</th><th>Arguments</th><th>Syntax-independent?</th><th>Language-independent?</th></tr>
<tr><td>syntactic relevance</td><td rowspan="2">formulae</td><td>no</td><td rowspan="6">no</td></tr>
<tr><td>essential relevance</td><td>yes</td></tr>
<tr><td>path-relevance</td><td rowspan="3">formulae, belief set K</td><td rowspan="2">yes, except for K</td></tr>
<tr><td>cell-relevance</td></tr>
<tr><td>canonical (path/cell) relevance</td><td>yes</td></tr>
<tr><td>R-sensitive relevance</td><td>formulae, belief set K, relation R over letters</td><td>yes</td></tr>
</table>

Appendix C: Proofs

Observation 3.2. For contingent K the B_i in the finest splittings of K are unique up to tautological equivalence. That is: let K be a contingent belief set, and $\mathbf{E} = \{E_i\}_{i\in I}$ its finest splitting. Suppose both $K \dashv\vdash \cup\{B_i\}_{i\in I}$ and $K \dashv\vdash \cup\{B_i'\}_{i\in I}$ where $E(B_i) \subseteq E_i$ and $E(B_i') \subseteq E_i$. Then each $B_i \dashv\vdash B_i'$.

Notation. We write $v(X) = 1$ as shorthand for $v(x) = 1$ for all $x \in X$, while $v(X) = 0$ abbreviates $v(x) = 0$ for some $x \in X$.

Proof. Suppose otherwise. Then there is a valuation u with say $u(B_j) = 1$ and $u(B_j') = 0$. Since K is consistent, there is also a valuation v with $v(K) = 1$ so $v(B_i) = v(B_i') = 1$ for all $i \in I$. Let w be the valuation that agrees with u on all letters in E_j and agrees with v on all other letters. Then $w(K) = w(\cup\{B_i\}_{i\in I}) = 1$ while also $w(K) = w(\cup\{B_i'\}_{i\in I}) = 0$ giving a contradiction.

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 contingent 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. The theorem is immediate when a^* shares a letter with b^*. So suppose otherwise.

Left to right: Suppose that a is canonically path-relevant to b (mod K). Then there is a finite sequence $x_0,\ldots x_{n+1}$ of formulae with $x_0 = a^*$, $x_{n+1} = b^*$, all of $x_1,\ldots,x_n \in K^\#$, and each x_i sharing at least one letter with x_{i+1}. Since a^* shares no letter with b^*, we have $n \geq 1$. 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 $x_1,\ldots,x_n \in K^\#$, and each x_i shares at least one letter with x_{i+1}, it follows that all of the letters in $x_1,\ldots,x_n$ 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 (mod 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. So there are letters $p, q \in E_i$ with p occurring in a^* and q occurring in b^*. Since $p, q \in E_i$ and $K^\#$ is in least letter-set form, they must occur in formulae $y, z \in B_i \subseteq K^\#$. To complete the proof we need to show that there are $x_1,\ldots,x_n$ ($n \geq 0$) in $K^\#$ with $y = x_1$, $x_n = z$ and each x_i sharing a letter with x_{i+1}. But this must hold because otherwise we could take the closure $\{y\}^+$ of $\{y\}$ under the relation of sharing a letter, to divide E_i further into non-empty sets $E(\{y\}^+)$ and $E_i \setminus E(\{y\}^+)$ which would split K further.

Observation 5.1. Let L be the propositional language generated by the letters $\{p,q\}$, and define f: $L \to L$ by putting $f(p) = p$, $f(q) = p\leftrightarrow q$, and homomorphic for compound formulae. Then there is a bijection φ between valuations on the language such that such that for all formulae $a \in L$, $v(a) = v'(f(a))$, where v': $L \to \{0,1\}$ is the counterpart $\varphi(v)$ of v: $L \to \{0,1\}$.

Proof. We construct the bijection φ as follows: for each valuation v: $L \to \{0,1\}$ define $\varphi(v) = v'$: $L \to \{0,1\}$ by setting $v'(p) = v(p)$ and $v'(q) = v(p\leftrightarrow q)$. We need to check that (1) φ is a bijection between valuations on L, and (2) for all formulae $a \in L$, $v(a) = v'(f(a)$.

For (1), since the set of valuations is finite (4 elements), it suffices to show that φ is injective. Suppose $v \neq w$; we need to show $v' \neq w'$. *Case 1*: Suppose $v(p) \neq w(p)$. Then immediately $v'(p) = v(p) \neq w(p) = w'(p)$ and we are done. *Case 2*: Suppose $v(p) = w(p)$ but $v(q) \neq w(q)$. Then $v(p\leftrightarrow q) \neq w(p\leftrightarrow q)$ so $v'(q) = v(p\leftrightarrow q) \neq w(p\leftrightarrow q) = w'(q)$ and again we are done.

For (2), it suffices to show that $v(p) = v'(f(p))$ and $v(q) = v'(f(q))$. The former is immediate since $v'(f(p)) = v'(p) = v(p)$ by the constructions of f and v'. For the latter, $v'(f(q)) = v'(p\leftrightarrow q)$ by the construction of f. *Case 1*: Suppose $v(q) = 1$. Then $v'(q) = v(p\leftrightarrow q) = v(p) = v'(p)$ so $v'(p\leftrightarrow q) = 1$, giving us $v(q) = v'(f(q))$ as desired. *Case 2*: Suppose $v(q) = 0$. Then $v'(q) = v(p\leftrightarrow q) = v(\neg p) = v'(\neg p)$ so $v'(p\leftrightarrow q) = 0$, again giving us $v(q) = v'(f(q))$ as desired.

Theorem 6.2. If we apply AGM contraction to the finest splitting $K^\#$ of a contingent belief set K, rather than to K itself, then it respects relevance.

Proof. 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 (mod K). We need to derive the turnstile version of the theorem from the epsilon one. Assume the epsilon version, i.e. that for contingent K, whenever $x \in K^{\#}$ but $x \notin K^{\#}-a$ then a is canonically relevant to x (mod K). Suppose that K is contingent, $K^{\#} \vdash x$, $K^{\#}-a \nvdash x$; we need to show that a is canonically relevant to x (mod K).

Since $K^{\#} \vdash x$ we have $K^{\#} \vdash x^*$, so there are $a_1,\ldots,a_k \in K^{\#}$ with $a_1 \wedge \ldots \wedge a_k \vdash x^*$. Since K is consistent, we may assume without loss of generality that each a_i^* shares a letter with x^*. Since $K^{\#}-a \nvdash x$ we likewise have $K^{\#}-a \nvdash x^*$, so there is an $i \leq k$ with $K^{\#}-a \nvdash a_i$, so that $a_i \notin K^{\#}-a$. By the epsilon version of the theorem, a is canonically relevant to a_i (mod K). That is, there is a cell E_j of the finest partition **E** of K such that each of the sets $E(a^*) \cap E_j$ and $E(a_i^*) \cap E_j$ is non-empty.

To show that a is canonically relevant to x (mod K) and complete the proof it will suffice to show that $E(x^*) \cap E_j$ is non-empty. But since $a_i \in K^{\#}$ all the letters of a_i come from the same cell, so $E(a_i) \subseteq E_j$. Since $E(a_i^*) \subseteq E(a_i)$ this gives us $E(a_i^*) \subseteq E_j$. Since a_i^* shares a letter with x^*, this tells us that $E(x^*) \cap E_j$ is non-empty as desired.

It is also possible to prove Theorem 6.2 (turnstile version) directly, essentially by including the above considerations within a re-run of the proof of the epsilon version in Kourousias and Makinson (2007).

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.

Belnap, Nuel D. 1960. "Entailment and relevance". *The Journal of Symbolic Logic* 25: 144-146.

Kourousias, George and David Makinson 2007. "Parallel interpolation, splitting, and relevance in belief change. *Journal of Symbolic Logic* 72: 994-1002.

Lang, Jerôme, Paolo Libertatore and Pierre Marquis 2003. "Propositional independence: formula-variable independence and forgetting". *Journal of Artificial Intelligence Research* 18: 391-443.

Makinson, David 2007. "Friendliness and sympathy in logic" in *Logica Universalis,* 2nd edition, ed. J.-Y. Beziau: 195-224. Basel: Birkhauser Verlag. Also appeared under the title "Friendliness for logicians" in *We Will Show Them! Essays in Honour of Dov Gabbay, vol 2,* ed S. Artemov et al: 259-292. London: College Publications.

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: 266–278. Stanford, CA. CSLI Publications (CSLI Lecture Notes n° 96). An earlier version was presented in the *Preliminary Proceedings of Information Theoretic Approaches to Logic, Language and Computation*, 1996.

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.

Ryan, M.D. 1991. "Defaults and revision in structured theories", in *Proceedings of the Sixth IEEE Symposium on Logic in Computer Science* (*LICS'91*): 362-373.

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

Acknowledgements

The author wishes to thank the organizers of the Dagstuhl seminar on Formal Models of Belief Change in Rational Agents in August 2007 for the opportunity of developing these reflections. The preliminary report that was made available on the webpage of that meeting is superseded by the present paper. Many thanks to Karl Schlechta for a discussion on language independence that contributed to section 5, and an anonymous referee for helpful comments. George Kourousias prepared the diagram.

Dept. of Philosophy, Logic & Scientific Method
London School of Economics
Houghton Street, London WC2A 2AE, UK
Email: david.makinson@gmail.com