Beliefs, Belief Revision, and Splitting Languages

Rohit Parikh[1]

Computer Science Department

Brooklyn College and City University Graduate Center[2]

Much current work in the study of belief revision goes back to a now classic paper due to Alchourron, Gärdenfors and Makinson [AGM]. The central issue is how to revise an existing set of beliefs T to a new set of beliefs $T * A$ when a new piece of information A is received. If A is consistent with T, then it is easy: we just add A to T and close under logical inference to get the new set of beliefs. The harder problem is how to revise the theory T when a piece of information A *inconsistent* with T is received. Clearly, as Levi has suggested, T must first be contracted to a smaller theory $T' = T \dot{-} \neg A$ which *is* consistent with A and then A added to T'. However, it is not clear how $T \dot{-} \neg A$ should be obtained. The mere deletion of $\neg A$ from T will clearly not leave us with a theory and there is in general no unique way to get a theory T' which is contained in T and does not contain $\neg A$.

Suppose, for example, that I believe that country Saturnia is hot and country Urania is cold. Now I discover that the two countries have very similar climates. Do I drop my belief that Saturnia is hot or that Urania is cold? Clearly I cannot retain them both.

[1] Research supported in part by NSF grant CCR-9208437 and a grant from the Research Foundation of CUNY. This paper appeared in *Proc. Logic, Language and Computation*, Ed. Moss, Ginzburg and de Rijke, CSLI 1999, pp. 266-278 (earlier version appeared in 1996 in the prliminary proceedings).

[2] I thank Sam Buss, Samir Chopra, Horacio Arlo Costa, Konstantinos Georgatos and Bernhard Heinemann for comments. Research supported in part by NSF grant CCR-9208437 and a grant from the Research Foundation of CUNY.

The AGM approach does not actually tell us what to think about the two lands in question. What it does tell us is *if* we do have some procedure for updating, what logical properties such a procedure should satisfy. These properties (the AGM axioms) have been widely studied and model theoretic results proved for them (see [G], [KM]. [S]). Yet some issues remain.

Notation: In the following, L is a finite propositional language.[3] We assume that the constants *true, false* are in L. We shall use the letter L both for a set of propositional symbols and for the formulae generated by that set. It will be clear from the context which is meant. $A \Leftrightarrow B$ means that A and B are logically equivalent, i.e. that $A \leftrightarrow B$ is a tautology, i.e. true under all truth assignments. Similarly, $A \Rightarrow B$ means that $A \rightarrow B$ is a tautology. If X is a set of formulae then $Con(X)$ is the logical closure of X. In particular, X is a theory iff $X = Con(X)$. We shall use letters T, T' etc. for theories. $T * A$ is the revision of T by A, and finally, $T \dot{+} A$ is $Con(T \cup \{A\})$, i.e. the result of a brute addition of A to T (followed by logical closure) without considering the need for consistency.

AGM have proposed the following widely accepted axioms for the revision operator $*$:

1. $T * A$ is a theory.
2. $A \in T * A$
3. If $A \Leftrightarrow B$, then $T * A = T * B$.
4. $T * A \subseteq T \dot{+} A$
5. If A is consistent with T, i.e. it is not the case that $\neg A \in T$,

[3]This restriction is only made for convenience. The results continue to hold for a countably infinite first order language without equality.

then $T * A = T \dot{+} A$.

6. $T * A$ is consistent if A is.

Remark: Sometimes two more axioms having to do with revision by conjunctions are also included. Since we do not find a strong intuitive reason behind them we have omitted them. Please see the end of the paper for a discussion.

Unfortunately, the AGM axioms are consistent with the trivial update, which is defined by:
If A is consistent with T, then $T * A = T \dot{+} A$, otherwise $T * A = Con(A)$.
Thus in case A is inconsistent with T, under this update, all information in T is simply discarded. Clearly this is unsatisfactory because we would like to keep as much of the old information as is feasible. Hence the AGM axioms need to be supplemented to rule out the trivial update. However, various actual proposals have run into trouble, either by being too flexible and allowing implausible update operators, or worse, by allowing *only* the trivial update (see [DP]; [L] and [AP] give examples of some difficulties with the approach of [DP]).

We propose axioms for update operators which are consistent with the AGM axioms and which block the trivial update. The axioms are based on the notion of splitting languages. We shall explain first the intuitive idea behind this. The existing set of beliefs T may contain information about various matters. E.g. my current state of beliefs contains beliefs about the location of my children, the state of health of my teeth, and beliefs about the forthcoming election in India. In case one of my beliefs about the location of my children turns out to be false, it surely ought not to affect my beliefs about the election, since the subject matters of the two beliefs do not interact in any way. In order to model this intuition mathematically, we need to define in a rigorous way what it means to say that some given set of beliefs can be split among various unrelated matters. The notion of splitting languages does this for us[4]. Intuitively a theory T in language L *splits* if L is a union of two or more disjoint sub-languages, and the beliefs in T are generated by separate beliefs in the various sub-languages.

Definition 1: 1) Suppose T is a theory in the language L and let $\{L_1, L_2\}$ be a partition of L. We shall say that L_1, L_2 *split* the theory T if there are formulae A, B such that A is in L_1, B is in L_2 and $T = Con(A, B)$. Similarly we say that (mutually disjoint) languages $L_1, L_2, .., L_n$ split T if there exist formulae $A_i \in L_i$ such that $T = Con(A_1, ..., A_n)$. We may also say that $\{L_1, ..., L_n\}$ is a T-splitting.

2) If $L_1 \subset L$ then we say that T is *confined* to L_1 if $T = Con(T \cap L_1)$. (Note that in that case T also splits between L_1 and $L - L_1$, with the $L - L_1$ part being trivial, i.e. any formula of $L - L_1$ which is a theorem of T will be a tautology.)

In part 1 of the definition, we can think of T as being generated by the various T_i in languages L_i. Then the condition implies that T contains no "cross-talk" between L_i and L_j for distinct i, j. Part 2 of the definition says that T knows nothing about the part $L - L_1$ of L.

Remark: If P and P' are partitions of L, P is a T-splitting and P

[4] This is a very natural notion and indeed we suspect that the notion of splitting has a wider application than just in belief revision.

refines P' then P' will also be a T splitting.[5] For example suppose that $P = \{L_1, L_2, L_3\}$ is a T-splitting and let $P' = \{L_1 \cup L_2, L_3\}$. Then P' is a 2-element partition, P is a 3-element partition which refines P' and P' is also a T-splitting. For let $T = Con(A_1, A_2, A_3)$ where $A_i \in L_i$ for all i. Then $T = Con(A_1 \wedge A_2, A_3)$ and $A_1 \wedge A_2 \in L_1 \cup L_2$ so that P' is also a T-splitting.

Example: Let $L = \{P, Q, R, S\}$, and $T = Con(P \wedge (Q \vee R))$. Then $T = Con(P, Q \vee R)$, and the partition $\{\{P\}, \{Q, R\}, \{S\}\}$ will be (the finest) T-splitting. $\{\{P, Q, R\}, \{S\}\}$ is also a T-splitting, but not the finest. Also, T is confined to the language $\{P, Q, R\}$ and knows nothing about S.

Lemma 1: Given a theory T in the language L, there is a unique *finest* T-splitting of L, i.e. one which refines every other T-splitting.

Lemma 1 says that there is a unique way to think of T as being composed of disjoint information about certain subject matters.

Lemma 2: Given a formula A, there is a *smallest* language L' in which A can be expressed, i.e., there is $L' \subseteq L$ and a formula $B \in L'$ with $A \Leftrightarrow B$, and for all L'' and B'' such that $B'' \in L''$ and $A \Leftrightarrow B''$, $L' \subseteq L''$.

Although A is equivalent to many different formulas in different languages, lemma 2 tells us that nonetheless, the question, "What is A actually *about*?" can be uniquely answered by providing a smallest language in which (a formula equivalent to) A can be stated.

All proofs are at the end of this paper.

[5] P *refines* P' if every element of P is a subset of some element of P'. Equivalently, the equivalence relation corresponding to P extends the equivalence relation corresponding to P'. P will have smaller members than P' does and more of them.

The axioms:

The general rationale for the axioms is as follows. If we have information about two subject matters which, as far as we know, are unrelated (are split) then when we receive information about *one* of the two, we should only update our information in that subject and leave the rest of our beliefs unchanged. E.g. suppose I believe that Barbara is rich and Susan is beautiful and only that. Later on I meet Susan and realize that she is not beautiful. My beliefs about Barbara should remain unchanged since I do not connect Susan and Barbara in any way. If on the other hand I had initially believed that Barbara had made her money *as the agent* for Susan who was a beautiful model, then the two beliefs would be *connected* and finding that Susan was not beautiful could require an adjustment also of my beliefs about Barbara.

In fact, the notion of language splitting seems intrinsic to any attempt to form a theory of anything at all. Any observation or experiment gives us an enormous amount of information. E.g. when we are dealt a hand of cards, we are dealt them in a certain order, either by the right hand or the left hand of the dealer, who may have grey or brown or blue eyes. We usually ignore all this extra information and concentrate on the *set* of cards received. There is a tacit assumption, for instance, that the color of the dealer's eyes will not affect the probability that the hand contains two aces. This assumption that we can ignore some aspects while we are considering others is inherent in almost all intellectual activity.[6]

Now we give our new axioms P1–P3, giving intuitive justification for each.

[6] The third chapter of Cherniak's book [C] gives arguments why beliefs must thus be divided into subsets, and cites supporting statements from Quine and Ullian [QU] as well as from Herbert Simon [Si].

We also give a single axiom P which implies all of P1–P3.

Axiom P1: If T is split between L_1 and L_2, and A is an L_1 formula, then $T * A$ is also split between L_1 and L_2.

Justification: The two subject areas L_1 and L_2 were unconnected. We have not received any information which *connects* these two areas, so they remain separate.

Axiom P2: If T is split between L_1 and L_2, A, B are in L_1 and L_2 respectively, then $T * A * B = T * B * A$.

Justification: Since A and B are unrelated, they do not affect each other and so it should not matter in which order they are received.

Axiom P3: If T is confined to L_1 and A is in L_1 then $T * A$ is just the consequences in L of $T *' A$ where $*'$ is the update of T by A in the sub-language L_1.

Justification: Since we had no information about $L - L_1$ and have received none in this round, we should update as if we were in L_1 only. $L - L_1$, about which we have no prior opinions and no new information, should simply not have any impact.

All these axioms follow from axiom P, below.

Axiom P: If $T = Con(A, B)$ where A, B are in L_1, L_2 respectively and C is in L_1, then $T * C = Con(A) *' C \dot{+} B$, where $*'$ is the update operator for the sub-language L_1.

Justification: We have received information only about L_1 which does not pertain to L_2 so we should revise only the L_1 part of T and leave the rest alone.

It is easy to see that P implies P1. To see that P3 is implied, we use the special case of P where the formula B is the trivial formula *true*. To see that P implies axiom P2, suppose T is split between L_1 and L_2, and A, B are in L_1 and L_2 respectively. Let $T = Con(C, D)$ where $C \in L_1$ and $D \in L_2$. Then we get $T * A * B = (T * A) * B =_P [(Con(C) *' A) \dot{+} D] * B =_P (Con(C) *' A) \dot{+} (Con(D) *'' B)$. The two occurrences of $=_P$ indicate where we used the axiom P. Now the last expression $(Con(D) * B) \dot{+} (Con(C) * A)$ is symmetric between the pairs (C, A) and (D, B) and calculating $T * B * A$ yields the same result.

Remark: The trivial update procedure cannot satisfy P2 (or P), though it does satisfy P1 and P3. It follows that any procedure that does satisfy P cannot be the trivial procedure.

Justification: Let $T = Con(P, Q)$ and let $A = P$ and $B = \neg Q$. Then the trivial update yields $T * A * B = T * B = Con(\neg Q)$ and $T * B * A = Con(B) * A = Con(\neg Q, P)$. This violates P2. Also, $T * B = Con(\neg Q)$ which violates P. Thus P, or P2 alone, rules out the trivial update.

In theorem 1 we shall restrict AGM 1–6 to the case of those updates where both T and A are individually consistent and only their union might not be. This is because we can suppose that our current state of belief T about the subject matter of L originates in a state T_0 where we only believe tautologies (or if not, T_0 is at least consistent) and T is obtained from T_0 through zero or more revisions. Suppose that at some stage we are told an inconsistent formula A. Then axiom 3 tells us that this is equivalent to being told a blatant contradiction like $P \wedge \neg P$ and we would simply not believe A in that case. Hence if A is inconsistent, then $T * A$ should be just T. The AGM axioms 1–6 in their original form *force* that $T * true = T$ for consistent T and *disallow* it for an inconsistent T. Our restriction has the fortunate consequence that $T * true$ always equals T.

Definition 2: Given a theory T, language L and formula A, let L'_A be the smallest language in which A can be expressed and L^T_A be the smallest language containing L'_A such that $\{L^T_A, L - L^T_A\}$ is a T-splitting. Thus L^T_A is a union of certain members of the finest T-splitting of L, and in fact the smallest in which A can be expressed.

Example: Let $L = \{P, Q, R, S\}$, and $T = Con(P \wedge (Q \vee R) \wedge S)$. Then $T = Con(P, Q \vee R, S)$ and $\{\{P\}, \{Q, R\}, \{S\}\}$ will be the finest T-splitting. If A is the formula $P \vee \neg Q$, then L'_A is the language $\{P, Q\}$. But L^T_A, the smallest language *compatible* with the T-splitting, will be the larger language $\{P, Q, R\}$ which is the union of the sets $\{P\}$ and $\{Q, R\}$ of the finest T-splitting.

Theorem 1: There is an update procedure which satisfies the six AGM axioms and axiom P.

Proof: We define $T * A$ as follows. Given T and A, if A is consistent with T then let $T * A = T \dot{+} A$.

Otherwise, if A is not consistent with T, then write $T = Con(B, C)$ where B, C are in L^T_A, $L - L^T_A$ respectively.[7] Then let $T * A = Con(A, C)$. B, C are unique up to logical equivalence, hence this procedure yields a unique theory

[7] We allow that $L^T_A = L$, in which case $L - L^T_A$ will be the empty language with no

$T * A$. To see that it satisfies axioms 1–6 of AGM is routine. For the proof that it satisfies axiom P, see the proofs section of this paper. □

Example: Let, as before, $T = Con(P, Q \vee R, S)$. Then the partition $\{\{P\}, \{Q, R\}, \{S\}\}$ is the finest T-splitting., Let A be the formula $\neg P \wedge \neg Q$, then L^T_A is the language $\{P, Q, R\}$. Thus B will be the formula $P \wedge (Q \vee R)$ and $C = S$. B represents the part of T incompatible with the new information A. Thus $T * A$ will be $Con((\neg P \wedge \neg Q), S)$. The update procedure of theorem 1 notices that A has no quarrel with S and keeps it. As we will see, axiom P requires us to keep S.

Remark: In this update procedure we used the trivial update on the sub-language L^T_A, but we did not need to. Thus suppose we are given certain updates $*_{L'}$ for sub-languages L' of L. We can then build a new update procedure $*$ for all of L by letting $T * A = (B *_{L'} A) \dot{+} C$ in the proof above, where $L' = L^T_A$. What this does is to update B by A on L^T_A according to the old update procedure, but preserves all the information C in $L - L^T_A$.

Georgatos' axiom: K. Georgatos has suggested that axiom P2 be strengthened to require that in fact we should have $T * A * B = T * B * A = T * (A \wedge B)$. Thus axiom P2 would be revised as follows:
Axiom P2g: If T is split between L_1 and L_2, A, B are in L_1 and L_2 respectively, then $T * A * B = T * B * A = T * (A \wedge B)$.
But this is easily achieved. Suppose we are given a current theory T with its partition P_1 and a new piece of information C, and the theory $Con(C)$ has its own partition P_2. Now let P be the (unique) finest partition such that both

non-logical symbols. L^T_A will still contain the constants $true$ and $false$, and we take C to be the trivial formula $true$.

P_1 and P_2 are refinements of P. E.g. if P_1 is $\{\{P,Q\},\{R\},\{S\},\{T\}\}$ and P_2 is $\{\{P\},\{Q,R\},\{S\},\{T\}\}$, then P will be $\{\{P,Q,R\},\{S\},\{T\}\}$. Now P is also *a* partition for both T and C though not necessarily the finest partition for either. Say $P = (L'_1, L'_2, L'_3)$ where T is axiomatized by (A_1, A_2, A_3) in (L'_1, L'_2, L'_3) respectively, and C by (C_1, C_2, C_3), also in (L'_1, L'_2, L'_3). Then let $T * C$ be $Con(D_1, D_2, D_3)$ where $D_i = A_i * C_i$. Now we get the AGM axioms, axiom P, and also the Georgatos axiom. We skip the proofs which are similar to the other proofs earlier.

Computational Considerations: The following are rather trite observations, but may be of value. If we have a theory which has a large language, but which is split up into a number of small sub-languages, then the revision procedure outlined above is going to be computationally feasible. This is because when we get a piece of information which lies in one of the sub-languages (or straddles only two or three of them) then we can leave most of the theory unchanged and revise only the affected part. All of this is just common sense and very likely how we humans actually think. The results above give us a formal framework which shows that this is actually doable in a precise way.

Other bases for L: We recall that the set of truth assignments over some language $L = \{P_1, ..., P_n\}$ is just an n-dimensional vector space over the prime field of characteristic 2. Hence there is more than one basis of atomic predicates for this language. We regard the P_i as basic, but this is not necessary. For example, the language $L = \{P, Q\}$ is also generated by R, Q where R is $P \leftrightarrow Q$. To see this we merely need to see that P can be expressed in terms of Q, R. But this is easy, for $P \Leftrightarrow (Q \leftrightarrow R)$. Such a way of formalizing a language may be natural at times. For example, if my children Vikram and Uma have gone to a movie together, then it is very likely that both have come home or neither has. So it may be natural for me to formalize my belief in the language $\{R, U\}$ where R is $V \leftrightarrow U$. I may originally think that they are together but not home, so my theory will be $Con(R, \neg U)$. Later, if I find that Uma is home, I will revise to $Con(R, U)$, retaining my belief that they are together. It is obvious that the results of this paper will continue to hold for such adjustments of the "atomic" symbols.

The Proofs:

The proof of Lemma 1 depends on lemma A below.

Definition: Let $\{L_1, ..., L_n\}$ be a partition of L, and $t_1, ..., t_n$ be truth assignments. Then by $Mix(t_1, ..., t_n; L_1, ..., L_n)$ we mean the (unique) truth assignment t on L which agrees with t_i on L_i.

Example: Suppose that $L = \{P, Q, R\}$, $L_1 = \{P, Q\}$ and $L_2 = \{R\}$. Let the truth assignment t_1 be (1,1,1) on P, Q, R respectively, where 1 stands for *true*. Similarly, t_2 is (0,0,0). Then $Mix(t_1, t_2, L_1, L_2)$ will be (1,1,0) and $Mix(t_2, t_1, L_1, L_2)$ equals (0,0,1). Now the formula $A = (Q \rightarrow R)$ does not respect the splitting L_1, L_2. Hence both t_1 and t_2 satisfy A, but $Mix(t_1, t_2, L_1, L_2)$ does not (although $Mix(t_2, t_1, L_1, L_2)$ does).

Lemma A: $\{L_1, ..., L_n\}$ is a T-splitting iff for every $t_1, ..., t_n$ which satisfy T, $Mix(t_1, ..., t_n; L_1, ..., L_n)$ also satisfies T.

Proof of lemma A: ($\Rightarrow$) Suppose $T = Con(A_1, ..., A_n)$ where $A_i \in L_i$.

If $t_1, ..., t_n$ satisfy T, let $t = Mix(t_1, ..., t_n; L_1, ..., L_n)$. Then, for each i, since t agrees with t_i on L_i, t satisfies A_i. Hence it satisfies T.

($\Leftarrow$) Write $t \models T$ to mean that T is true under t and let $Mod(T) = \{t | t \models T\}$. Now let X_i be the projection of $Mod(T)$ to L_i. I.e. $X_i = \{t' | (\exists t)(t \models T \wedge t' = t \uparrow L_i\}$, where $t \uparrow L_i$ is the restriction of t to L_i. If $t \in Mod(T)$ then for all i, $t \uparrow L_i \in X_i$, i.e. $t \in X_1 \times X_2 \times ... \times X_n$. Hence $Mod(T) \subseteq X_1 \times X_2 \times ... \times X_n$.

But the reverse inclusion is also true. For if $t \in X_1 \times X_2 \times ... \times X_n$, then for each i, there must exist t_i which agree with t on L_i and such that $t_i \models T$. But then $t = Mix(t_1, ..., t_n; L_1, ..., L_n)$ so $t \models T$ also. Thus $Mod(T) = X_1 \times X_2 \times ... \times X_n$.

Let $B_i \in L_i$ be such that $X_i = Mod(B_i)$. Then it is immediate that $Mod(T) = X_1 \times X_2 \times ... \times X_n = Mod(B_1) \times Mod(B_2) \times ... \times Mod(B_n)$. Hence $T = Con(B_1, ..., B_n)$ and so $\{L_1, ...L_n\}$ is a T-splitting. $\square$

Proof of lemma 1: Suppose that $P = \{L_1, ...L_n\}$ is a *maximally fine* T-splitting. Such a P must exist but there is no a priori reason why it should also be *finest*; there might be more than one maximally fine partition. However, we will show that this does not happen and that a maximally fine partition is actually *finest*, i.e. refines every other T-splitting. If it does not, then there must exist a T-splitting $P' = \{L'_1, ..., L'_m\}$ such that P does *not* refine P'. Then there must be i, j such that L_i overlaps L'_j but is not contained in it. By renumbering we can take $i = j = 1$, so that L_1 overlaps L'_1 but is not contained in it. Consider $\{L'_1, (L'_2 \cup ... \cup L'_m)\}$ which is also a (2-element) T-splitting which P does not refine.

So without loss of generality we can take $m = 2$ and P' to be a two-element partition $\{L'_1, L'_2\}$. We now show that $\{L_1 \cap L'_1, ..., L_n \cap L'_1, L_1 \cap L'_2, ..., L_n \cap L'_2\}$ must also be a T-splitting. But this is a contradiction, for even if we throw out all empty intersections from these $2 \times n$ intersections, we still have at least $n + 1$ non-empty ones. For L_1 gives rise to two non-empty pieces, and all the other L_i must give rise to at least one non-empty piece. We thus get a *proper* refinement of P which was supposedly *maximally* refined.

To show that $\{L_1 \cap L'_1, ..., L_n \cap L'_1, L_1 \cap L'_2, ..., L_n \cap L'_2\}$ is also a T-splitting, we use lemma A. Let $t_1, ..., t_n, t'_1, ..., t'_n$ be any truth assignments which satisfy T. Let

$$t''' = Mix(t_1, ..., t_n, t'_1, ..., t'_n; L_1 \cap L'_1, ..., L_n \cap L'_2)$$

We have to show that t''' satisfies T. Let $t' = Mix(t_1, ..., t_n; L_1, ..., L_n)$ and let
$t'' = Mix(t'_1, ..., t'_n; L_1, ..., L_n)$. Since $\{L_1, ..., L_n\}$ is a T-splitting, both t', t'' satisfy T. Also, $\{L'_1, L'_2\}$ is a T-splitting and hence $t^{iv} = Mix(t', t''; L'_1, L'_2)$ does satisfy T.

To see that t''' satisfies T, it suffices to show that $t''' = t^{iv}$. We note now that for example, t^{iv} agrees with t' on L'_1 and the latter agrees with t_1 on L_1. Hence t^{iv} agrees with t_1 (and hence with t''') on $L_1 \cap L'_1$. Arguing this way we see that t^{iv} agrees with t''' everywhere so that $t^{iv} = t'''$. Thus $t''' \in Mod(T)$. Now use lemma A.

Thus $\{L_1 \cap L'_1, ..., L_n \cap L'_1, L_1 \cap L'_2, ..., L_n \cap L'_2\}$ is a T-splitting which properly refines P, which was maximally fine. This is a contradiction. Since assuming that P was not finest led to a contradiction, P is indeed the finest

T-splitting. $\Box$.

Proof of lemma 2: Let us say that A is *expressible* in L' if there is a formula B in L' with $A \Leftrightarrow B$. We want to show that there is a smallest such L'. So let L_1 (with B_1) and L_2 (with B_2) be *minimal* such languages. Then $A \Leftrightarrow B_1$ and $A \Leftrightarrow B_2$ and so we have that $B_1 \Rightarrow B_2$. By Craig's lemma there is a B_3 in $L_1 \cap L_2$ such that $B_1 \Rightarrow B_3 \Rightarrow B_2$. But then we must have $A \Rightarrow B_1 \Rightarrow B_3 \Rightarrow B_2 \Rightarrow A$ and all are equivalent. Hence A is expressible in $L_1 \cap L_2$. By the minimality of L_1 and L_2 we must have $L_1 = L_1 \cap L_2 = L_2$ and L_1, L_2 were in fact not just minimal but actually smallest. $\Box$

For the proof that the procedure of theorem 1 satisfies axiom P, we need lemma B. The language L_A^T is as in definition 2.

Lemma B: Suppose that C is inconsistent with T. Let $\{L_1, ..., L_n\}$ be the finest T-splitting and let $T = Con(A_1, ..., A_n)$ where $A_i \in L_i$. Given a formula C, $L_C^T = \bigcup L_i : L_i$ overlaps $L'_C = \bigcup L_i : L_i \subseteq L_C^T$.
Moreover, if $*$ is the update procedure of theorem 1, then

$$\begin{aligned} T * C &= Con(C\ ,\ A_j : L_j \text{ does not overlap } L_C^T) \\ &= Con(A_j : L_j \text{ does not overlap } L_C^T) \dot{+} C. \end{aligned}$$

The proof of the lemma B is quite straightforward and relates updates to the finest splitting of T. During the update, those A_j such that L_j does not overlap L_C^T are exactly the A_j that remain untouched by the update. C has no quarrel with *them.* The other A_j are dropped and replaced by C.

Proof of theorem 1 (continued): It is sufficient to show that axiom P holds.

To see that P holds, suppose that $T = Con(A, B)$ where A, B are in L_1, L_2 respectively and C is in L_1. Let $\{L'_1, ..., L'_n\}$ be the finest T-splitting of L (and therefore refines $\{L_1, L_2\}$). Let $T = Con(A_1, ..., A_n)$ with $A_i \in L'_i$. Note that both L^T_A and L^T_C will be subsets of L_1 and indeed they will each be a union of some members, contained in L_1, of some of the L'_i. A will be a consequence of some of the A_i, each from some member of this finest T-splitting and contained in L^T_A.

Under the update of theorem 1, those A_i which lie in L^T_C will be replaced by C. The others will remain. Hence, $Con(A) * C = Con(A_i : L_i \subseteq (L^T_A - L^T_C), C)$.
Also $B = \bigwedge A_i : L_i$ does not overlap L_1. So $Con(A) * C \dotplus B$ equals
$Con(A_i : L_i \subseteq (L^T_A - L^T_C), C, \bigwedge A_i : A_i$ does not overlap $L_1) = T * C$. $\Box$

Remark: The notion of splitting languages and the lemmas can easily be extended to first order logic without[8] equality. We use the fact that if a first order theory (without equality) is consistent then it has a countably infinite model, in particular, a model whose domain is the natural numbers. Call such models standard. Now given a standard first order structure $\mathcal{M}$ which interprets a language L and a sub-language L' of L, we can define a reduct $\mathcal{M}'$ of $\mathcal{M}$ which is just $\mathcal{M}$ restricted to L'.[9] Given a partition $L_1, ..., L_n$ of a language L and standard L-structures $\mathcal{M}_1$,...,$\mathcal{M}_n$ we can define $Mix(\mathcal{M}_1, ..., \mathcal{M}_n, L_1, ..., L_n)$ to be that standard structure $\mathcal{M}'$ which agrees

[8]Sam Buss has pointed out that problems arise if equality is present. Once we have equality present, one can express the formula for instance which says that the model has cardinality 2. Another formula in a disjoint language can say that the cardinality is at least 3. These two formulae will conflict though they have no non-logical symbols in common.

[9]For instance suppose $L = \{P, R\}$ and $L' = \{P\}$. If $\mathcal{M}$ is $(N, \mathcal{P}, \mathcal{R})$ then $\mathcal{M}'$ would be $(N, \mathcal{P})$.

with $\mathcal{M}_i$ on L_i. Then lemma 1 can be generalized in the obvious way and all our arguments go through without any trouble.

Comment on AGM axioms 7 and 8: Axiom 7 of AGM says that $T * (A \wedge B) \subseteq (T * A) \dot{+} B$ and axiom 8 says further that if B is consistent with $T * A$ then the two are equal. We do not feel that these axioms are consistent with the spirit of our work for the following reason. Suppose that $A = (\neg P \vee Q)$ and $B = (P \vee Q)$, then $A \wedge B$ is equivalent to Q and says nothing about P. Now revising a theory T first by A could cause us to drop some P-related beliefs we had, and revising after that with B we might not recover them. But revising with $A \wedge B$ should leave our P beliefs unchanged, provided that our beliefs about P and Q were not connected. Thus contrary to 7, revising with the conjunction may at times preserve more beliefs than revising first with A and then with B. This is why it does not seem to us that axioms 7 and 8 *should* hold in general.

To give a somewhat different, concrete example, suppose you believe that to *reach* a certain place, the *path* should be clear. Write this as $R \rightarrow P$. You also believe (strongly) that your grandmother is afraid of flying and therefore is *not taking flying lessons*. Call this (latter) $\neg F$. You now receive the news $A = (\neg P \vee F)$, that either the path is not clear or your grandmother is taking flying lessons. You conclude that the path is not clear and that you will not reach in time. Later you are told $B = (P \vee F)$, that either the path is clear or that your grandmother is taking flying lessons. You will conclude that you will reach your destination after all. But you will never acquire the belief that your grandmother is taking flying lessons. But you would have acquired that belief if you had been told $A \wedge B$, i.e. F.

Postscript: Subsequent to our submission of this paper to ITALLC, a paper by del Cerro and Herzig [CH] has appeared which also mentions dependence in the context of belief revision. However, the sort of dependence they discuss is from probability theory and not the one that we have discussed above.

References

The technical material depends only on the reference [AGM] and then only on the definitions. The other references are given for background. [QU], [C] and [Si] discuss motivation. The others consist mostly of technical developments arising from [AGM].

[AGM] Carlos Alchourron, Peter Gärdenfors and David Makinson, "On the logic of theory change: partial meet contraction and revision functions", *J. Symbolic Logic* 50 (1985) 510–530

[AP] Horacio Arlo Costa and Rohit Parikh, "On the inadequacy of C2 (research note)", typescript, May 1995.

[C] Christopher Cherniak, *Minimal Rationality*, MIT Press 1986.

[CH] Luis Farinas del Cerro and Andreas Herzig, "Belief change and dependence", proceedings of *Theoretical Aspects of Rationality and Knowledge*, Ed. Y. Shoham, Morgan Kaufmann 1996, 147-161.

[DP] Adnan Darwiche and Judea Pearl, "On the logic of iterated theory revision", *Theoretical Aspects of Reasoning about Knowledge*, Edited by R. Fagin, Morgan Kaufmann 1994, 5–22

[G] Adam Grove, “Two modellings for theory change”, *J. Phil. Logic* 17 (1988) 157–170.

[KM] Hirofumi Katsuno and Alberto Mendelzon, “Propositional knowledge base revision and minimal change”, *Artificial Intelligence*, 52 (1991) 263–294.

[L] Daniel Lehmann, "Belief Revision, revised”, *Proceedings of 14th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann 1995, 1534–1541.

[QU] Willard van Orman Quine and Joseph Ullian, *The Web of Belief*, 2nd edition, Random house 1978.

[S] Krister Segerberg, “Belief revision from the point of view of doxastic logic”, *Bull. of the IGPL* 4 (1995) 535–553.

[Si] Herbert Simon, *Administrative Behaviour*, Macmillan 1947.