

Language splitting, Relevance and the Logic of Campaigning

Rohit Parikh

City University of New York

Brooklyn College and CUNY Graduate Center

IMSc, Chennai
January 7, 2009

Abstract: When a theory is updated with new information, few problems arise if the new information is consistent with the theory. If, however, the new information is inconsistent with the theory, then some adjustment is needed. The AGM axioms provide guidelines for the *form* of the adjustment though not for its exact nature.

Abstract: When a theory is updated with new information, few problems arise if the new information is consistent with the theory. If, however, the new information is inconsistent with the theory, then some adjustment is needed. The AGM axioms provide guidelines for the *form* of the adjustment though not for its exact nature.

What we would like to do is to maximize the information retained and also to adjust only in those areas of the old theory to which the new information is pertinent. We use the notion of language splitting to this end. This notion allows us to carve up a theory into disjoint pieces about different subject matters. We show that such a splitting is unique. Tools that we use include a notion of mixing truth assignments and Craig's lemma.

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.

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.

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.

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. 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.

Notation: In the following, L is a finite propositional language. 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 $\neg A \in T$, then $T * A = T \dot{+} A$.
- ▶ 6. $T * A$ is consistent if A is.

The operator $*$ is non deterministic. Also, we want to retain as much information from T as possible. So one solution is to choose $T' = T \dot{-} \neg A$ be a *maximal subtheory* of T which does not contain $\neg A$ and then let $T * A = (T' \dot{+} A)$.

The operator $*$ is non deterministic. Also, we want to retain as much information from T as possible. So one solution is to choose $T' = T \dot{-} \neg A$ be a *maximal subtheory* of T which does not contain $\neg A$ and then let $T * A = (T' \dot{+} A)$.

The trouble is that for every B , both $A \rightarrow B$ and $A \rightarrow \neg B$ were in T and one of them will be in T' as well. Thus we are going to have $T' \dot{+} A$ be a complete theory. It will decide questions about which T itself had no opinion!

Remark: Sometimes two more axioms having to do with revision by conjunctions are also included. They are:

▶ $T * (A \wedge B) \subseteq (T * A) \dot{+} B$

▶ If $\neg B \notin T * A$ then $(T * A) \dot{+} B \subseteq T * (A \wedge B)$

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 $T * A = \text{Con}(A)$.*

Thus in case A is inconsistent with T , under this update, all information in T is simply discarded.

Tennant's triviality result (J. Symbolic Logic, 2006)

Let the operator $*'$ be defined by taking the closure of $\Theta \cup \{A\}$ if the two are consistent, and to take it to be $Con(A)$ otherwise.

Now for theory T and A inconsistent with T , define $T * A$ to be $A *' \Theta$. Then this operator $*$ satisfies all eight AGM axioms.

This means that when $T * A$ is non-trivial, it *completely ignores* T !

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.

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.

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 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.

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 $T = \text{Con}(A, B)$. Similarly we say that (mutually disjoint) languages L_1, L_2, \dots, L_n split T if there exist formulae $A_i \in L_i$ such that $T = \text{Con}(A_1, \dots, A_n)$. We may also say that $\{L_1, \dots, L_n\}$ is a T -splitting.

The notion of splitting languages does this for us.

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 $T = \text{Con}(A, B)$. Similarly we say that (mutually disjoint) languages L_1, L_2, \dots, L_n split T if there exist formulae $A_i \in L_i$ such that $T = \text{Con}(A_1, \dots, A_n)$. We may also say that $\{L_1, \dots, L_n\}$ is a T -splitting.

2) If $L_1 \subset L$ then we say that T is *confined* to L_1 if $T = \text{Con}(T \cap L_1)$. (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 refines P' then P' will also be a T splitting.¹ 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 $T = \text{Con}(A_1, A_2, A_3)$ where $A_i \in L_i$ for all i . Then $T = \text{Con}(A_1 \wedge A_2, A_3)$ and $A_1 \wedge A_2 \in L_1 \cup L_2$ so that P' is also a T -splitting.

¹ 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.

Example: Let $L = \{P, Q, R, S\}$, and $T = \text{Con}(P \wedge (Q \vee R))$. Then $T = \text{Con}(P, Q \vee R)$, 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.

Kourousias and Makinson have proved a generalization of Craig's interpolation theorem and used it to give an alternative proof of lemma 1, including for infinite languages. (*J. Symbolic Logic*, 2007).

Theorem: Let T in language L be split into theories $T_i : i \in I$ in disjoint languages $L_i : i \in I$. Let $T \vdash A$ then there are a finite number of theories $T_i : i \leq n$ and formulas $B_i : i \leq n$ such that $L(B_i) \subseteq L_i \cap L(A)$ and $B_1, \dots, B_n \vdash A$.

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.

In fact, the notion of language splitting seems intrinsic to any attempt to form a theory of anything at all.

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.

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 $T = \text{Con}(A, B)$ where A, B are in L_1, L_2 respectively and C is in L_1 , then $T * C = \text{Con}(A) *' C \dot{+} B$, 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.

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

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 .

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

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 .

It is easy to see that P implies $P1$.

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*.

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 $T = \text{Con}(C, D)$ 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 $T = \text{Con}(P, Q)$ and let $A = P$ and $B = \neg Q$. Then the trivial update yields $T * A * B = T * B = \text{Con}(\neg Q)$ and $T * B * A = \text{Con}(B) * A = \text{Con}(\neg Q, P)$. This violates P2. Also, $T * B = \text{Con}(\neg Q)$ 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.

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 $T * \text{true} = T$ for consistent T and *disallow* it for an inconsistent 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_A^T be the smallest language containing L'_A such that $\{L_A^T, L - L_A^T\}$ is a T -splitting. Thus L_A^T 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 $T = \text{Con}(P \wedge (Q \vee R) \wedge S)$. Then $T = \text{Con}(P, Q \vee R, S)$ 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_A^T , 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 $T = \text{Con}(B, C)$ where B, C are in $L_A^T, L - L_A^T$ respectively. Then let $T * A = \text{Con}(A, C)$. B, C are unique up to logical equivalence, hence this procedure yields a unique theory $T * A$. 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. \square

Example:

Let, as before, $T = \text{Con}(P, Q \vee R, S)$. Then the partition $\{\{P\}, \{Q, R\}, \{S\}\}$ is the finest T -splitting. Let A be the formula $\neg P \wedge \neg Q$, then L_A^T is the language $\{P, Q, R\}$.

Example:

Let, as before, $T = \text{Con}(P, Q \vee R, S)$. Then the partition $\{\{P\}, \{Q, R\}, \{S\}\}$ is the finest T -splitting., Let A be the formula $\neg P \wedge \neg Q$, then L_A^T 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 $\text{Con}(\neg P \wedge \neg Q, S)$. 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_A^T , 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 $L' = L_A^T$. What this does is to update B by A on L_A^T according to the old update procedure, but preserves all the information C in $L - L_A^T$.

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 $T * A * B = T * B * A = T * (A \wedge B)$.

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 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 $\{\{P, Q, R\}, \{S\}, \{T\}\}$.

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 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.

Computational Considerations:

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.

Other bases for L :

We recall that the set of truth assignments over some language $L = \{P_1, \dots, P_n\}$ 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 V stands for *Vikram is home*, U stands for *Uma is home*, and R is $V \leftrightarrow U$.

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 V stands for *Vikram is home*, U stands for *Uma is home*, and 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 our results continue to hold for such adjustments of the “atomic” symbols.

The Proofs:

The proof of Lemma 1 depends on lemma A below.

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

The Proofs:

The proof of Lemma 1 depends on lemma A below.

Definition: Let $\{L_1, \dots, L_n\}$ be a partition of L , and t_1, \dots, t_n be truth assignments. Then by $Mix(t_1, \dots, t_n; L_1, \dots, L_n)$ 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:

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

Lemma A:

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

Proof of lemma A: (\Rightarrow) Suppose $T = Con(A_1, \dots, A_n)$ where $A_i \in L_i$. If t_1, \dots, t_n satisfy T , let $t = Mix(t_1, \dots, t_n; L_1, \dots, L_n)$. Then, for each i , since t agrees with t_i on L_i , t satisfies A_i . Hence it satisfies T .

Lemma A:

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

Proof of lemma A: (\Rightarrow) Suppose $T = Con(A_1, \dots, A_n)$ where $A_i \in L_i$. If t_1, \dots, t_n satisfy T , let $t = Mix(t_1, \dots, t_n; L_1, \dots, L_n)$. 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 $Mod(T) = \{t \mid t \models T\}$. Now let X_i be the projection of $Mod(T)$ to L_i . I.e. $X_i = \{t' \mid (\exists t)(t \models T \wedge t' = t \upharpoonright L_i)\}$, where $t \upharpoonright L_i$ is the restriction of t to L_i . If $t \in Mod(T)$ then for all i , $t \upharpoonright L_i \in X_i$, i.e. $t \in X_1 \times X_2 \times \dots \times X_n$. Hence $Mod(T) \subseteq X_1 \times X_2 \times \dots \times X_n$.

But the reverse inclusion is also true. For if $t \in X_1 \times X_2 \times \dots \times X_n$, 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 $t = \text{Mix}(t_1, \dots, t_n; L_1, \dots, L_n)$ so $t \models T$ also. Thus $\text{Mod}(T) = X_1 \times X_2 \times \dots \times X_n$.

But the reverse inclusion is also true. For if $t \in X_1 \times X_2 \times \dots \times X_n$, 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 $t = \text{Mix}(t_1, \dots, t_n; L_1, \dots, L_n)$ so $t \models T$ also. Thus $\text{Mod}(T) = X_1 \times X_2 \times \dots \times X_n$.

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

Proof of lemma 1: Suppose that $P = \{L_1, \dots, L_n\}$ 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.

Proof of lemma 1: Suppose that $P = \{L_1, \dots, L_n\}$ 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 $P' = \{L'_1, \dots, L'_m\}$ 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 $\{L'_1, (L'_2 \cup \dots \cup L'_m)\}$ 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 $\{L_1 \cap L'_1, \dots, L_n \cap L'_1, L_1 \cap L'_2, \dots, L_n \cap L'_2\}$ 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 $\{L_1 \cap L'_1, \dots, L_n \cap L'_1, L_1 \cap L'_2, \dots, L_n \cap L'_2\}$ is also a T -splitting, we use lemma A. Let $t_1, \dots, t_n, t'_1, \dots, t'_n$ be any truth assignments which satisfy T . Let

$$t''' = \text{Mix}(t_1, \dots, t_n, t'_1, \dots, t'_n; L_1 \cap L'_1, \dots, L_n \cap L'_2)$$

We have to show that t''' satisfies T . Let

$t' = \text{Mix}(t_1, \dots, t_n; L_1, \dots, L_n)$ and let

$t'' = \text{Mix}(t'_1, \dots, t'_n; L_1, \dots, L_n)$. Since $\{L_1, \dots, L_n\}$ is a T -splitting, both t', t'' satisfy T . Also, $\{L'_1, L'_2\}$ is a T -splitting and hence $t^{iv} = \text{Mix}(t', t''; L'_1, L'_2)$ 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 $t''' \in \text{Mod}(T)$. Now use lemma A.

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 $t''' \in \text{Mod}(T)$. Now use lemma A.

Thus $\{L_1 \cap L'_1, \dots, L_n \cap L'_1, L_1 \cap L'_2, \dots, L_n \cap L'_2\}$ 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. \square .

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. \square

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 $\{L_1, \dots, L_n\}$ be the finest T -splitting and let $T = \text{Con}(A_1, \dots, A_n)$ where $A_i \in L_i$. Given a formula C , $L_C^T = \bigcup L_i : L_i \text{ overlaps } C$
 $L'_C = \bigcup L_i : L_i \subseteq L_C^T$.

Moreover, if $*$ is the update procedure of theorem 1, then

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

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 ι

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 it 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 $T = \text{Con}(A, B)$ where A, B are in L_1, L_2 respectively and C is in L_1 . Let $\{L'_1, \dots, L'_n\}$ be the finest T -splitting of L (and therefore refines $\{L_1, L_2\}$). Let $T = \text{Con}(A_1, \dots, A_n)$ with $A_i \in L'_i$. Note that both L_A^T and L_C^T 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_A^T .

Under the update of theorem 1, those A_i which lie in L_C^T will be replaced by C . The others will remain. Hence,

$$\text{Con}(A) * C = \text{Con}(A_i : L_i \subseteq (L_A^T - L_C^T), C).$$

Also $B = \bigwedge A_i : L_i$ does not overlap L_1 . So $\text{Con}(A) * C \dot{+} B$ equals $\text{Con}(A_i : L_i \subseteq (L_A^T - L_C^T), C, \bigwedge A_i : A_i \text{ does not overlap } L_1) = T * C$. \square

Remark: The notion of splitting languages and the lemmas can easily be extended to first order logic without 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' .² Given a partition L_1, \dots, L_n of a language L and standard L -structures $\mathcal{M}_1, \dots, \mathcal{M}_n$ we can define $Mix(\mathcal{M}_1, \dots, \mathcal{M}_n, L_1, \dots, L_n)$ to be that standard structure \mathcal{M}' which agrees 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.

²For 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}) .

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 .

The Logic of Campaigning

Joint work with Walter Dean

NEW YORK After Sen. Barack Obama's comments last week about what he typically eats for dinner were criticized by Sen. Hillary Clinton as being offensive to both herself and the American voters, the number of acceptable phrases presidential candidates can now say is officially down to four. "At the beginning of 2007 there were 38 things candidates could mention in public that wouldn't be considered damaging to their campaigns, but now they are mostly limited to 'Thank you all for coming,' and 'God bless America,'" ABC News chief Washington correspondent George Stephanopoulos said on Sunday's episode of *This Week*.

The Onion,³ May 8, 2008

³ *The Onion* is a tongue-in-cheek weekly newsmagazine.

Abstract: *We (Walter Dean and I) consider the issue of the sort of statements which a candidate running for office might make during the course of the campaign. We assume that there are various issues on which different groups of voters have preferences and that the candidate is aware of these different groups and their preferences. Her problem is to choose statements which will have the net effect of improving her perception among these groups of voters. We represent the situation using the propositional calculus with each propositional variable representing one issue, formal theories (representing the perception of the candidate in the minds of the voters), and for each voter, a distance function from that voter's preferred world to the possible worlds which might come to exist should the candidate be elected. We assume that candidates talk in such a way as to decrease the distance which voters perceive between their own preferences and the candidate's position, keeping in mind that different groups of voters have different preferences.*

When a candidate utters a sentence A , she is evaluating its effect on several groups of voters, G_1, \dots, G_n with one group, say G_1 being her primary target *at the moment*.

Thus when Clinton speaks in Indiana, the Indiana voters are her primary target but she is well aware that other voters, perhaps in North Carolina, are eavesdropping. Her goal is to increase the likelihood that a particular group of voters will vote for her, but without undermining the support she enjoys or hopes to enjoy from other groups. If she can increase *their support* at the same time as wooing group G_1 , so much the better, but at the very least, she does not want to undermine her support in G_2 while appealing to G_1 . Nor does she want to be caught in a blatant contradiction. She may not always succeed, as we all know, but remaining consistent, or even truthful, is surely part of her strategy. Lies are expensive.

We will represent a particular group of like minded voters as *one* formal voter, but since the groups are of different sizes, these formal voters will not all have the same influence. A formal voter who represents a larger group of actual voters will have a larger *size*. We will assume that each voter has a preferred ideal world – how that voter would like the world to be as a result of the candidate's policies, should she happen to be elected.

Thus suppose the main issues are represented by $\{p, q, r\}$, representing perhaps, policies on the Iraq war, energy, and taxes. If the agent's ideal world is $\{p, q, \neg r\}$, then that means that the voter wants p, q to be true, and r to be false. But it may be that p is more important to the voter than q . Then the world $\{\neg p, q, \neg r\}$ which differs from the ideal world in just p will be worse for the voter than the one, $\{p, \neg q, \neg r\}$, which differs in just q .

We represent this situation by assigning a utility of 1 to the ideal world, and assigning weights to the various issues, adding up to at most 1. If the weights of p , q , r are .4, .2, and .4 respectively and the ideal world is $p, q, \neg r$, then a world in which p, q, r are all true will differ from the ideal world in just r . It will thus have a utility of $(1 - .4)$, or .6.

Each voter also has a theory T_c of the candidate, and in the first pass we will assume that the theory is simply generated by things which the candidate has said in the past. If the candidate has uttered (presumably consistent) assertions A_1, \dots, A_5 , then T_c will be just the logical closure of A_1, \dots, A_5 . If the candidate is truthful, then T_c will be a subtheory of T_a which is the candidate's own theory of the world.

The voter will assume that if the candidate is elected, then one of the worlds which model T_c will come to pass. The voter's utility for the candidate will be obtained from the utilities of these worlds, perhaps by calculating the expected utility over the (finitely many) models of T_c .

(Note that we are implicitly assuming that all the worlds are equally likely, something which is not always true, but even such a simple setting turns out to be rich enough for some insights.)

Suppose now that the candidate (who knows all this) is wondering what to say next to some group of voters. She may utter some formula A , and the perceived theory T_c will change to $T'_c = T_c + A$ (the logical closure of T_c and A) if A is consistent with T_c , and $T_c * A$ if not. Here the $*$ represents an AGM like revision operator. (Note: The AGM operator $*$ accommodates the revision of a theory T by a formula A which is inconsistent with T . For the moment we will assume that A is in fact something which the candidate believes and is consistent with T_c which is a subtheory of T_a , and thus $T_c * A$ really amounts to $T_c \dot{+} A$, i.e., the closure of $T_c \cup \{A\}$.)

Thus the candidate's utterance of A will change her perceived utility in the minds of the voters and her goal is to choose *that* A which will maximize her utility summed over all groups of voters. We can now calculate the utility to her of the utterance of a particular formula A .

Each group of voters will revise their theory of the candidate by including the formula A , and revising their utility evaluation of the candidate.

Let the old utility to group G_i calculated on the basis of T_c be U_i and the new utility calculated on the basis of $T_c * A$ be U'_i . Let w_i be the weight of the group G_i calculated on the basis of size, likelihood of listening to A which is greater for the current target group, and the propensity to actually vote. Then the change in utility on the basis of uttering A , or the *value* of A , will be

$$val(A) = val(A, T_c) = \sum w_i (U'_i - U_i)$$

The rational candidate should utter *that* A which will have the largest value for $val(A)$.

Example 1: Sometime during the primaries, Hillary Clinton announced that she had shot a duck as a child. Now ducks do not vote, so we know she was not appealing to them. Who was she appealing to? Clearly those voters who oppose gun control. Other things being equal, a world in which there is gun control is worse for them than a world in which there isn't, and Hillary's remark will clearly decrease the set of worlds (in the voters' perception) in which Hillary is president and there *is* gun control. Presumably this will increase her utility in the minds of these voters.

But what about other voters who *do* prefer gun control? Now first of all, note that the fact that she shot a duck as a child does not *eliminate* worlds in which she is president and there is gun control – it merely decreases their number. Moreover, when she is campaigning in Pennsylvania or Indiana, these voters are not her primary voters. The likelihood that Massachusetts voters will be affected by the duck story will be (hopefully) less than the likelihood of a Pennsylvania voter being so affected. There is even the likelihood that voters who disfavor gun control – perhaps because they own a gun, will be more passionate about gun control (against it, and give it more weight), than voters who favor gun control for more abstract reasons who may be for it, but with less passion.

Definitions and Notation

C will denote the candidate under consideration. v will denote the group of voter (in the single block case).

$T_c =$ voters' current theory of candidate C

$T_a =$ candidate C 's *actual theory*

$A_t = \{P_1, \dots, P_n\}$ are atomic propositions corresponding to issues (which we may identify with the integers $\{1, \dots, n\}$).

\mathcal{W} is a finite set of worlds. Worlds will be seen as truth assignments, i.e., as functions $w : At \rightarrow \{1, -1\}$ such⁴ that $w(i) = 1$ if $w \models P_i$ and $w(i) = -1$ if $w \not\models P_i$ and we write $w(i)$ to denote the i th component of w . It may well happen that there is a non-trivial theory T_0 which is shared by both voters and candidates, and then of course the worlds to consider (even initially) will be those which model T_0 .
 \mathcal{L} = the propositional language over At , which we may occasionally identify with the corresponding *propositions*, or subsets of \mathcal{W} .

⁴We are using truth values 1,-1 for truth and falsehood so as to make some formulas simpler.

$p_v : At \rightarrow \{1, 0, -1\} = V$'s preferred world, represented as follows

$$p_v(i) = \begin{cases} 1 & \text{if } V \text{ would prefer } p_i \text{ to be true} \\ 0 & \text{if } V \text{ is neutral about } p_i \\ -1 & \text{if } V \text{ would prefer that } p_i \text{ be false} \end{cases}$$

$x_v : At \rightarrow [0, 1]$ assigns weight $x_v(i)$ to proposition i . To simplify thoughts, we assume $\sum_{1 \leq i \leq n} x_v(i) \leq 1$. Thus all worlds have non-negative values. Nothing will turn on this since utilities are only characterized up to a linear transformation which takes u into $au + b$ with $a > 0$.

$u_v(w)$ = the utility of world w for V

$$u_v(w) = \sum_{1 \leq i \leq n} p_v(i) \cdot x_v(i) \cdot w(i)$$

Voter types: [o]ptimistic, [p]essimistic, [e]xpected value.
Given a possible set of worlds, according to the candidate's position T_c so far, the optimistic voters will assume that the candidate will implement the best one which is compatible with T_c . The pessimistic voters will assume the worst, and the expected value voters will average over the possible worlds which satisfy T_c .

$ut_v^t(T)$ = the utility of the theory T for V of type t (we leave out the subscript v below).

▶ $ut^o(T) = \max\{u(w) : w \models T\}$

▶ $ut^p(T) = \min\{u(w) : w \models T\}$

▶ $ut^e(T) = \frac{\sum_{w \models T} u(w)}{|\{w : w \models T\}|}$

Note that these three values for T are only defined when T is consistent. We could think, with slight abuse of language, of the ut functions as applying to sets of worlds rather than to theories, and if X, Y are sets of worlds, we will have,

- ▶ $ut^o(X \cup Y) = \max(ut^o(X), ut^o(Y))$
- ▶ $ut^p(X \cup Y) = \min(ut^p(X), ut^p(Y))$
- ▶ $ut^e(X \cup Y) \in$ the closed interval of $ut^e(X), ut^e(Y)$

The last claim about ut^e requires that X, Y be disjoint. For suppose $X = \{w_1, w_2, w_3, w_4, w_5\}$ and $Y = \{w_3, w_4, w_5, w_6, w_7\}$ where w_1, w_2, w_6, w_7 have low values, and w_3, w_4, w_5 have high values, then in each of X, Y , high will dominate, but in $X \cup Y$, low will dominate.

Note that a voter of type p always prefers to be better informed. Although such a voter always “assumes the worst” he is also such that hearing additional messages can never *decrease* the utility he assigns to his current theory of the candidate. As such, such a voter will always *prefer* to hear more information on which to base his vote. This seems like a rational strategy. Of course, a pessimistic voter can also be regarded as a ‘play it safe’, or ‘worst outcome’ voter.

Let $val(A, T) =$ the value of announcement $A \in \mathcal{L}$ be what a particular announcement A is worth to the candidate.

$$val(A, T) = ut(T \dot{+} A) - ut(T)$$

What are the sets of formulas from which the candidate might choose? Possible sets $\mathcal{X} \subseteq \mathcal{L}$ of statements from which C might select to choose the message A she will utter:

- ▶ $\mathcal{X}_\ell = \mathcal{L}$ (this would allow for contradicting a statement already in T_c or lying about statements in T_a)
- ▶ $\mathcal{X}_a = T_a$ (only select a message from her actual theory)
- ▶ $\mathcal{X}_m = \mathcal{L} - \{\neg A : A \in T_c\}$ (allow for any message which is consistent with T_c)
- ▶ $\mathcal{X}_t = \mathcal{L} - \{\neg A : A \in T_a\}$ (allow for any message which is consistent with T_a)

An honest candidate will only choose a message (from \mathcal{X}_a) which she actually believes, but a Machiavellian candidate may well choose a message (from \mathcal{X}_m) which she does not believe, perhaps even disbelieves, but which is compatible with what she has said so far. (The subscript t in \mathcal{X}_t stands for 'tactical') But as we see, even an honest candidate has options.

$best(T, \mathcal{X})$ = the most advantageous message for C which is an element of \mathcal{X} .

$$best(T, \mathcal{X}) = \operatorname{argmax}_{A \in \mathcal{X}} val(A, T)$$

Clearly the goal of the candidate is to find the formula $best(T, \mathcal{X})$.

Example 2: Given that there will be a class of maximally beneficial messages given $T_\emptyset = \emptyset$, we next need to illustrate how certain of these are ruled out in cases where T_c is non-empty, corresponding to situations in which candidate is on the record with respect to certain issues, and still wants to be either honest or, at worst, Machiavellian. Here's an example:

Candidate = McCain

A_1 = "We must achieve victory in Iraq"

A_2 = "We should pull out from Iraq"

A_3 = "We should stay in Iraq but change strategy"

A_4 = "We should maintain the Bush administration's strategy in Iraq"

Lets assume that i) McCain has chosen $\mathcal{X} = \mathcal{X}_m$, ii) $best(\mathcal{X}_\ell, T_\emptyset) = A_2$ but iii) $A_1 \in T_c$. Since we can then assume that $T_c \vdash \neg A_2$, A_2 is ruled out for McCain. Thus despite the greater benefit which could be had from A_2 were he not already on the record about Iraq, McCain must choose between A_3 and A_4 . We can assume that he chooses A_3 because $best(T_c, \mathcal{X}) = A_3$.

Results

In the two propositions below we assume that all theories whose values are being considered are consistent. For instance for **3.1.1** we assume that both A and $\neg A$ are consistent with T .

Proposition

Assume e-voters. For all $A \in \mathcal{X}$, there exist positive a, \dots, f such that

1. $a \cdot \text{val}(A, T) + b \cdot \text{val}(\neg A, T) = 0$
2. $\text{val}(A \wedge B)$
 $= \text{val}(A, T) + \text{val}(B, T \dot{+} A)$
 $= \text{val}(B, T) + \text{val}(A, T \dot{+} B)$
3. $c \cdot \text{val}(A \vee B) + d \cdot \text{val}(A \wedge B, T) = e \cdot \text{val}(A, T) + f \cdot \text{val}(B, T)$

Here the numbers a, \dots, f represent the number of worlds satisfying particular (relevant) theories.

Proof: The proofs involve simple algebra. For instance the models of T are divided into (nonempty) sets of models satisfying A and $\neg A$ respectively. The average of any quantity over the models of T is therefore in the open interval whose endpoints are the average over the models of $T + A$ and over the models of $T + \neg A$. \square

Comment: Note that **1.1.2** appears to say that the order in which statements are made does not matter. This is clearly not true in general, but in our model we do not address the reasons why non-commutativity might hold.

We immediately get,

Proposition

Assume e-voters. For all $A \in \mathcal{X}$, either exactly one of $val(A, T)$, $val(\neg A, T)$ is positive and the other negative, or they are both zero.

This follows immediately from **3.1.1**.

Proposition

Assume e-voters. Then there is at least one complete extension of T whose value is at least as great as that of T .

This result follows from the fact that a linear function defined on a convex set always takes its maximum value on the boundary. An alternative argument would follow Lindenbaum's procedure for creating a complete extension of a theory, and each time there is a choice, choosing the formula which yields a benefit rather than a loss.

It follows that if the voters are all e-voters, then the candidate should be as explicit as possible about her positions – she cannot lose and she might gain.

With optimistic voters, it is likely that one of $A, \neg A$ yields a loss and the other one yields no change. With pessimistic voters, one will in general yield a benefit and the other no loss or gain. It follows that if a candidate leaves her position ambiguous (this corresponds to the theory T_c being incomplete) then this is because she is counting on the optimistic voters. Thus if T_c is satisfied by exactly two truth assignments s, t then if all voters are e-voters, she would always be better off with the voters perceiving her as pure s or as pure t , and she has nothing to gain from ambiguity. This is even more so with pessimistic voters. It is with optimistic voters that ambiguity between s, t helps her, and if she *is* ambiguous, it means she is counting on the optimists to give her the benefit of the doubt.

Example: Assuming e-voters, while it is clear that statements A and $\neg A$ cannot both benefit (or both hurt) a candidate, we *could* have a situation where $A \wedge B$ and $(\neg A) \wedge (\neg B)$ are both beneficial or both harmful. Worlds which satisfy T fall into four groups:

1. **X:** worlds which satisfy both A, B .
2. **Y:** worlds which satisfy $A, \neg B$
3. **Z:** worlds which satisfy $\neg A, B$
4. **U:** worlds which satisfy $\neg A, \neg B$

Each of the sets **X,Y,Z,U** could have on average, better or worse worlds than the full set of worlds which satisfy T_C . It can be that any of them are good and any of them are bad, provided that if one is good, at least one is bad, and if one is bad then at least one is good.

For instance, voters may feel that it is good to have a military buildup *and* go to war with Iran, or *not* have a buildup and not go to war, while thinking it foolish to have *just* one and *not* the other.

	A	$\neg A$
B	X	Z
$\neg B$	Y	U


Figure 1

The utilities over the models of $T + A \wedge B$, $T + A \wedge \neg B$, $T + \neg A \wedge B$, $T + \neg A \wedge \neg B$, must average out to utilities over models of T . But any of them can be higher or lower, as long as they average out.

For optimistic voters, the best world could be in any of sets **X,Y,Z,U**, and the next best in any of them also.

If the voters are optimistic, the candidate is best off saying nothing. In other words, the best thing to say is something like, "These pancakes are nice."

If $T_c \subseteq T_d$, where c, d , are two candidates, then for all types of voters, the best possible extension of T_c is at least as good as the best possible extension of T_d . But for optimistic voters and e-voters, the best extension of T_c is likely to be strictly better.



The good world \circ in **Figure 2** might stay when more information is received, and we move from the larger square to the smaller one, or it can be deleted. So for optimistic voters, more information can only make things worse. For pessimistic voters, it is the other way.