Belief revision and uncertain reasoning

Guy Politzer, Laure Carles

Belief Revision and Uncertain Reasoning

Guy Politzer

C.N.R.S. , Saint-Denis

Laure Carles

C.R.E.A. , Ecole Polytechnique, Paris

Abstract

When a new piece of information contradicts a currently held belief, one has to modify the set of beliefs in order to restore its consistency. In the case where it is necessary to give up a belief, some of them are less likely to be abandoned than others. The concept of epistemic entrenchment is used by some AI approaches to explain this fact based on formal properties of the belief set (e. g. , transitivity). Two experiments were designed to test the hypothesis that contrary to such views, (i) belief is naturally represented by degrees rather than in an all-or-nothing manner, (ii) entrenchment is primarily a matter of content and not only a matter of form, and (iii) consequently prior degree of belief is a powerful factor of change. The two experiments used Elio and Pelletier's (1997) paradigm in which participants were presented with full simple deductive arguments whose conclusion was denied, following which they were asked to decide which premise to revise.

Belief Revision and Uncertain Reasoning

Changing belief is a pervasive mental activity. It occurs when the course of events does not meet the individual's expectations, or when indisputable facts run counter to the individual's anticipations. It also occurs during communication (e. g. , upon being convinced by one's interlocutor), in learning, in problem solving when new data transform the problem space, etc. This work addresses belief change by asking, How does an individual react when (s)he receives new information that contradicts some logical entailment of his/her current beliefs? More precisely, consider an individual who entertains a consistent set of beliefs regarding some specific topic. By deduction, that individual may infer a conclusion. Suppose a new piece of factual information on the world arrives and that it contradicts the former conclusion. Now the individual faces a conflict: On the one hand, (s)he believes the conclusion obtained by a valid deduction. But on the other hand, (s)he also believes the contradictory fact obtained from a source that is presumably reliable. The consistency of the original set of beliefs is temporarily broken.

One solution to restore consistency, called *belief revision* in Artificial Intelligence (Gärdenfors, 1988), consists of incorporating the new information and at the same time giving up some of the premises. But there may not be a unique way to do so. When there is a choice between premises to be abandoned, what are the rational bases for making a decision? One of the basic rationality principles that constrains belief change, and in particular belief revision, is that the new set which incorporates the new information should be the result of minimal changes from the original set (Harman, 1986). In addition, although this approach is non-probabilistic (that is, it considers propositional

contents to be fully believed) it nevertheless acknowledges that propositions which are candidates for abandonment may differ in their ease of disbelief. It seems clear that some beliefs are more important than others in terms of their explanatory power or informational content, which discourages the individual from abandoning them. This is what the notion *epistemic entrenchment* captures: a resistance to be given up. It is thus possible to spell out on logical grounds a few postulates which define a hierarchy among propositions of a belief set, hence a unique solution to the revision problem (Gärdenfors, 1992).

To what extent are such concepts and formalisms applicable to human knowledge representation and inferential activity? Elio & Pelletier (1997) presented a pioneering study that investigated two kinds of issue. One is to know whether the entrenchment of a belief is affected by the logical form of the associated proposition. For that purpose, they considered the special case of conditionals as opposed to non conditional sentences. Three experiments indicated that participants were more prone to abandon a conditional premise than a non conditional one. This was observed with thematic conditional statements presented in the framework of a scenario and with formal material (but was more pronounced in the former case). For example, participants in the first experiment received full arguments (here a Modus Ponens with formal material) such as: *if Lex's have a P, then they also have an R* (= P --> R); *Max is a Lex that has a P* (= P); *therefore Max has an R* (= R) followed by the contradictory information that in fact *Max does not have an R* (= ¬R). Participants were then asked to choose among several belief sets the one they thought "the best way to reconcile all the knowledge"; in the present case, they had to choose among: (1) P --> R, ¬P, ¬R; (2) ¬(P --> R), ¬R, P uncertain; (3) ¬(P --> R), P, ¬R and it was observed that option (1) was chosen more

frequently than (2) or (3). Overall, it was also observed that giving up a conditional premise was more frequent with Modus Ponens than with Modus Tollens arguments.

The other issue is related to the principle of minimal change. The problem is to operationalise this concept. From a syntactic view this refers to the change computed as a function of the number of propositions added to, or subtracted from, the initial set. From a model-theoretic approach there are various possibilities for computing the change. The authors considered several of these possibilities and tried to identify whether subjects complied with any one. The results of their final two experiments showed that no single algorithm could be identified. Rather, participants chose options that expressed uncertainty, or as the authors put it, "... revisions that were not minimal with respect to what was changed, but were instead minimal with respect to what they believed to hold true without doubt..." (page 454).   To illustrate, given the initial set *P, Q, R, S or ¬P, ¬Q, ¬R, ¬S* and the new information *P, ¬R, ¬S* participants tended to opt for the revision alternative *P, ¬R, ¬S, ?Q* (where ?Q means that Q is uncertain) rather than for *P, ¬Q, ¬R, ¬S.* Given that there was a strong tendency in the third experiment towards the selection of a revision option that expressed uncertainty (and no opposite trend in the first two experiments), what seems to be the most remarkable feature in the data overall is that when offered this kind of options to revise their set of beliefs, participants preferred to select them rather than options that expressed certainty. Now this tendency might be even more clear-cut than appeared in the data because participants were not always offered an uncertain option (as opposed to a certain option), especially in the first two experiments where no option expressed doubt in the conditional premise alone.

One reason to consider the importance of options that give participants an opportunity to doubt propositions stems from the observation that an important part of deductive reasoning occurs with uncertain premises. On some occasions they are treated as if they were certain (George, 1995), that is, people perform sound arguments to conclude with certainty. But on other occasions, such as after the reception of a new piece of information or when the premise is doubtful by itself, people attribute a degree of belief that is carried over to the conclusion. This has been demonstrated in a growing number of experiments, the most significant of which will be reviewed below.

In some studies, a piece of information was added to the premises of a deductive argument in order to manipulate the degree of belief in one of the premises.

Politzer & George (1992) presented the argument form Modus Tollens: *If somebody touches an object on display then the camera starts up; the camera did not start up;* conclusion: *nobody touched an object on display* to be evaluated on a five-point scale (certainly false; probably false; maybe true, maybe false; probably true; certainly true). One group of participants received this standard argument while a second group received the two premises plus the additional information *the apparatus was correctly plugged in* (the moderate reassuring condition) and a third group received the two premises plus the information *there was nothing faulty in the material or in the electrical supply* (the strong reassuring condition). While roughly the same proportion of participants (80%) rated the conclusion as true (either *probably true* or *certainly true)* in the three groups, the ratio of *probably true* to *certainly true* responses increased from 1:1 for the standard group to 1:2 for the second group to 1:3 for the third group, showing that the conclusion had a degree of belief that varied

monotonically as a function of the strength of the major (defined by the strength of the reassuring conditions).

These results were generalised by Chan & Chua (1994) who used various non causal conditional rules, two different populations and both Modus Ponens and Modus Tollens arguments. For example, with a Modus Ponens whose major was *If Steven is invited then he will attend the party* three levels of necessary conditions for the conclusion to hold (and therefore three degrees of belief in the major) were defined and introduced by an additional premise *(If Steven knows the host well/knows at least some people well/completes the report to night, then he will attend the party).* Essentially the same phenomenon was observed: the rate of endorsement of the conclusion was a decreasing function of the degree of necessity, whereas the rate of expression of doubt (maybe or maybe not) was an increasing function of the degree of necessity.

Stevenson and Over (1995) manipulated the degree of belief in the major premise by introducing various levels of frequency in the additional premise. For example, given the major premise *If John goes fishing, he will have a fish supper,* there were five levels of frequency*(John always/ almost always/ sometimes/ almost never/ never catches a fish).* The likelihood of the conclusion *John will have a fish supper* evaluated on a 10-point scale increased monotonically with the degree of belief.

Some other experiments demonstrate even more directly the effect of belief in the premises on belief in the conclusion. No additional information was used to manipulate belief in the premise but either the likelihood of the premise was independently measured or the premise was formulated in probabilistic terms. George (1995) asked participants to evaluate on a 7-point scale the truth status they attributed to controversial cause-effect conditional statements (e. g. , *if*

exports decrease, then unemployment will rise) before using these statements in Modus Ponens arguments with conclusions to evaluate on the same scale. There was a high correlation between belief in the conditional and belief in the conclusion, suggesting that the degree of belief in the conditional (the major of the argument) was conveyed to the conclusion *(unemployment will rise).*

Cummins (1995; Cummins, Lubart, Alksnis, and Rist, 1991) demonstrated that the acceptance rate of the conclusion of Modus Ponens and Modus Tollens arguments with a causal conditional premise depended on the availability of competing factors that could prevent the effect from occurring (which the author called disabling conditions). For example, of the following two arguments, *If the match was struck, then it lit; the match was struck; therefore it lit,* and *If Joe cut his finger, then it bled; Joe cut his finger; therefore it bled,* people accept more readily the conclusion of the first than that of the second. Thompson (1994) obtained similar results with causal, obligation, permission and definition rules by using conditionals with either low level of sufficiency (many necessary conditions were missing) or high level of sufficiency (few necessary conditions missing).

Liu, Lo, and Wu (1996) defined three levels of "perceived sufficiency" for conditional statements: high level (uncontroversial because definitionally true), medium (expressing common regulations or habits) and low (expressing obviously debatable stereotypes) so defining three levels of conditional probability of the consequent given the antecedent. They observed that the rate of endorsement of the conclusion of the standard conditional arguments was an increasing function of the level of perceived sufficiency.

In one of George's (1997) experiments, the degree of belief was manipulated by formulating the premises themselves in probabilistic terms. Two kinds of conditional were contrasted and each of them used as the major premise

in a Modus Ponens, e. g. : *If Peter is in the kitchen, then it is very likely that Mary is in the garden* (vs *not very likely).* The conclusion was evaluated on a 9-point scale. The modal response was identical to the probability term used in the conditional, suggesting inheritance of the degree of belief from premise to conclusion.

In brief, there is extensive evidence that shows that belief in premises can be operationalised and manipulated by various means. In particular when the premise is a conditional the manipulation always involves people's consideration of tacit conditions that are necessary in order to render the antecedent of the conditional actually sufficient. When this is achieved, the truth status of the conclusion is spontaneously treated by degree rather than in an all-or-nothing manner and this degree is highly dependent on the degree of belief in the premises.

From this point of view, belief change looks like a reassignment of strength of belief or weights to various propositions in the set (and giving up a conclusion, the result of a weight that has become null). In other words, belief change could be cast within the more general problem of reasoning under uncertainty.

The two experiments reported below were inspired by the foregoing considerations. They represent a short and preliminary step in the direction just outlined, namely an attempt to test the hypothesis that belief change can be measured by degrees, that is, in a qualitative rather than an all-or-nothing manner and is affected by the degree of belief in the premises.

In the first experiment, participants received each revision option in two versions, one certain and the other uncertain, so allowing a fair comparison of their frequency of choice. On the basis of research on uncertain reasoning and

on some of Elio & Pelletier's data, it was hypothesised that choosing an uncertain revision option is at least as natural a way to change belief in the premises as choosing a certain option.

The question of the relative entrenchment of the two premises of the Modus Ponens and Modus Tollens arguments was also addressed. In Elio and Pelletier's data with thematic content, the conditional premise was more frequently abandoned. We wanted to make sure that this result was not linked to the pecularities of the sample of response options that were offered to the participants. Finally, another objective was to generalise the study of premise entrenchment by introducing two other basic arguments which are similar to the first two but contain no conditional premise.

<div align="center">Experiment 1</div>

<u>Method</u>.

Materials. Each participant received a booklet that contained a page of instructions followed by four revision problems, one of each of the following four kinds of deductive argument:

- Modus Ponendo Ponens (henceforth MPP, referred to above, following the psychological tradition, as Modus Ponens): *If P then Q; P; therefore Q.*

- Modus Tollendo Tollens (MTT, referred to above, following again the psychological tradition, as Modus Tollens): *If P then Q; not Q; therefore not P.*

- Modus Tollendo Ponens (MTP, also called disjunctive syllogism): *P or Q; not P; therefore Q.*

- Modus Ponendo Tollens (MPT): *not both (P and Q); P; therefore not Q.*
From an information point of view all four are equivalent in that they amount to the detachment of the conclusion from the conjunction of the minor and the

conclusion. For example, Modus Ponendo Ponens amounts to the detachment of Q from (P and Q), etc.

Each problem contained (i) two premises supposed to have been uttered by two different persons (in order to allow independent questioning of the premises); (ii) the statement of the conclusion (not to be evaluated) which followed validly from the premises in accordance with one of the four arguments above; (iii) a piece of factual information presented as indisputable that contradicted the conclusion; (iv) five revision options presented in the same order that constituted various ways to reconcile the data; (v) a request to justify the choice of revision.

The first four revision options expressed:

- the negation of the minor premise

- a doubt on the minor premise

- the negation of the major premise

- a doubt on the major premise

Participants were instructed to choose either one and only one of these four options or to write up (as a fifth open option) any other solution (a combination from the first four or a choice not offered). This fifth option was proposed because it was felt that the first four might be too restrictive and fail to capture the whole range of revision choices that participants were likely to express. A sample problem is presented in the Appendix.

There was a common story (a space flight) covering all four problems. In order to limit belief bias, the premises were factual statements related to the planet being explored.

In order to avoid a confounding between argument form and content, each argument form was framed in four different scenarios (four events in the

exploration trip). In order to keep some coherence in the story, the four events appeared in the same order. Participants were required not to go back to a previous page. Four orders of presentation of the arguments were chosen out of the 24 theoretically possible.

Participants and procedure. Eighty students of Commerce and Business Administration at a higher learning institution in Paris, untutored in logic, participated on a voluntary basis. Questionnaires were administered in small groups and completed in 10 to 15 minutes without time limit.

Design. There were four argument types; subjects acted as their own controls.

Results.

Two kinds of comparison are of interest. The first one concerns the relative frequency of uncertain responses. It was predicted that these responses would be preferred to categorical denials. Table 1 indicates the frequency of choice (certain vs uncertain) for each of the four arguments.

—————————————

Insert Table 1 about here

—————————————

Collapsing across arguments, participants opted for an uncertain choice 77.8 per cent of the time and for a certain choice 16.6 per cent of the time. The differences across argument form are small and non significant ($F_{(3,\ 228)}$ ratios < 1). The frequencies of certain and uncertain responses were submitted to a bayesian analysis for each argument form separately. The ratio f1/f2 = frequency of uncertain responses /frequency of certain responses was chosen as the relevant indicator. The credibility intervals at p = .95 for these ratios were: [3.4; 13.2] for MPP; [2.3; 7.4] for MTT; [3.3; 12.6] for MTP; and [2.4; 7.5] for MPT. Medians of the associated distributions were equal to 6.37, 4.05, 6.12, and

4.12, respectively (note 1). In brief, considering all four argument forms, when people were offered to revise a premise, they produced an uncertain response four to six times more often than a categorical negative response.

These results are confirmed by individual results. For each of the participants (excluding five of them who produced more than one ambiguous response) if we subtract the number of certain choices from the number of uncertain choices (both varying from zero to 4), there were 63 positive values, 3 negative (and 9 ties), showing an overwhelming individual preference for uncertain choices ($z = 7,26$; $p << 10^{-7}$, sign test).

The second comparison concerns the choice of the premise to revise. Considering first the two conditional arguments, the prediction based on Elio & Pelletier's results that the conditional premise would be more often revised than the categorical one was confirmed: as shown in Table 2, the major (conditional) premise was selected about two to three times more often than the minor (categorical) premise.

————————————

Insert Table 2 about here

————————————

The reliability of this result was again submitted to a bayesian analysis, taking the ratio f1/f2 = frequency of revision of the major / frequency of revision of the minor as the indicator of interest. The 95% credibility intervals were [1.0; 3.1] for MPP and [1.9; 6.1] for MTT with medians equal to 1.81 and 3.34, respectively. Again, this group analysis was confirmed by an individual analysis. Sign tests calculated by using the four basic response categories of the questionnaire indicated that the preference for the major (conditional) premise over the minor (categorical) premise was reliable at the level of significance of

.01 for MPP (frequency of conditional = 37, non conditional = 21, z = 1.97) and for MTT (conditional = 48, non conditional = 15, z = 4.03).

Considering now the two arguments that have no conditional premise, table 2 indicates that there is again a clear difference in the frequency of revision choices: the major premise was selected about twice as often as the minor. The same type of bayesian analysis as above yielded the following results: the 95% credibility intervals were [0.9; 2.6] for MTP and [1.5; 4.5] for MPT, with medians equal to 1.54 and 2.57, respectively. The individual analysis using sign tests indicated a difference that fell short of significance at the .05 level for MTP (frequency of major = 37, minor = 24) and a significant difference for MPT (major = 46, minor = 18, z =3.37).   For the four arguments collapsed, the data indicated a ratio of 3:1 in favour of the major (major = 39, minor = 13, a highly reliable result (z = 4.18). Notice also that about one participant out of six spontaneously decided to revise both the major and the minor.

Finally, an ANOVA performed on Table 2 failed to reveal an effect of logical argument on the frequency of choice of the minor or of the major ((F $_{(1, 76)}$ ratios < 1).

Discussion.

In the present experiment, participants were offered the same opportunity to select a revision option in terms of the negation of a premise (a certain option) and in terms of the expression of a doubt in that premise (an uncertain option), which was not always the case in Elio and Pelletier's experiments. Whereas in their data there was only a trend towards the choice of the uncertain option, in the present experiment there is an overwhelming preference for the uncertain option, confirming the hypothesis that the natural way to express revision of a set of beliefs might be to doubt one premise (or more) rather than categorically

denying it. Of course, this claim is limited in scope to the specific kind of data base that has been considered, namely factual knowledge in the frame of an imaginary world.

Comparing the relative entrenchment of the premises of simple conditional arguments, Elio and Pelletier's observation was confirmed: the conditional premise was more often questioned than the minor. How to interpret this phenomenon? One might attribute it to some property of factual conditional sentences. Rather than looking for a *post hoc* explanation, we prefer to interpret it in the light of the data obtained on the other two arguments. On both MPT and MTP the major premise (which is non conditional) was also less entrenched than the minor. This suggests that it is not the conditional nature of the premise that is important, but its status *qua* major premise, that is, a compound proposition obtained by the connection of two atomic propositions, one of which is affirmed or negated as a minor. From this point of view, questioning the major more than the minor stems from the rather trivial fact that it is more complex in that it contains a connection and has more chances to be the source of error.

Even if future research confirms the weaker entrenchment of major premises, a formal property of the kind advocated by some AI approaches, the role of premise content is in no way ruled out. By using premises that had arbitrary factual content, this factor was eliminated from the first experiment. However, the observation that participants were essentially sensitive to the uncertainty of the premises, together with the hypothesis that belief revision consists of the reassignment of degrees of belief leads to the prediction that, *ceteris paribus,* the revision process will result in a premise being all the more often questioned as people have a lower degree of belief in that premise. This hypothesis, which was tested in the second experiment, has already got some

support from Elio's (1997, 1998) results. She reports several studies with causal conditionals in which the number of disabling conditions was manipulated, a factor which, as noted above, seems to act as a mediator for the plausibility of such conditionals. Participants were more prone to retain the conditional proposition than the categorical one when it had few disabling conditions, and the reverse obtained when it had many. However, in one study (indirectly reported in Elio, 1997), the differences were small and in another study (1997, first experiment) the results were not so clean-cut with one type of arguments (MTT) as they were with the other (MPP). This might be attributed to the rather crude and indirect estimation of prior belief based on number of disabling conditions. Although the author took the trouble to have the prior plausibility of the conditionals estimated independently, she did not relate those estimates to the posterior belief ratings. But examination of the data suggests that higher prior ratings of the conditionals were associated with greater entrenchment as indicated by their higher posterior ratings and the lower ratings of the associated categorical proposition. Unluckily, this observation is *post hoc* and limited to causals. In the next experiment, strong and weak prior degrees of belief were defined and the effect of this factor on the entrenchment of the major premise was studied with conditional propositions expressing various social rules (causal, decision and means-end rules).

<div align="center">Experiment 2</div>

<u>Method</u>.

    <u>Materials</u>.   Each participant received a booklet that contained a page of instructions followed by eight revision problems. The structure of a problem was nearly identical to that used in the first experiment, that is, two premises uttered by two different people, a conclusion, the negation of the conclusion presented

as an observed fact, then a request to choose a revision; the only difference is that the choice was from four options (minor declared false, minor declared uncertain, major declared false, major declared uncertain) instead of five. Only Modus Ponendo Ponens arguments were used.

On each booklet there were two blocks of four problems separated by two filler items. Problems were paired between the two blocks so that two members of a pair differed by the plausibility (high or low) of the major (conditional) premise while the thematic content was kept common to both problems. This was achieved by giving different antecedents but the same consequent to each pair of conditional premises. (These are given in the Appendix). In brief, the two blocks were made of four matched problems, with a high plausibility major in one block and the paired low plausibility major in the other block. Each block contained two high and two low plausibility problems. There were two kinds of booklets that differed by the order of presentation of the problems, so that one half of the participants received one order (randomly attributed), and the other half received the reverse order. The problems in the direct sequence were in the following order: H1, L2, H3, L4, L1, H2, L3, H4 (where H stands for high plausibility and L for low plausibility, and paired problems have the same number).

The sentences were chosen as follows. Ten pairs of conditional sentences (presented in counterbalanced orders) were submitted to nineteen participants who served as judges. They were asked to estimate the plausibility of the sentences on a five-point scale. Four pairs turned out to be highly discriminant, the mean difference in ratings between the two members of each pair being nearly equal to, or slightly above, three points on the scale (whose range was four). These sentences were used to constitute the four paired problems.

Participants and procedure. Thirty-six psychology students at the University of Paris VIII solved the problems before a lecture. They answered in 10 to 20 minutes. All of them were untutored in logic and belonged to the same population as the judges.

Design. All participants received two kinds of problems (high and low plausibility) so that they acted as their own controls. In addition, there were two groups defined by the order of presentation of the problems.

Results.

No difference between the two orders of presentation was observed; consequently the data were pooled across all participants whenever applicable.

It was predicted that premises with low plausibility would be more often questioned than premises with high plausibility. Overall, the low plausibility major premise was revised 91% of the time and the high plausibility major premise only 43.7% of the time. Considering the 36 (participants) x 4 (pairs of problems) = 144 paired observations, the number of times a low plausibility conditional was revised while the paired plausible conditional was not revised was equal to 73, whereas the reverse occurred only 4 times.

A within-subject comparison indicated that thirty-two participants out of thirty-six questioned the low plausibility premise more often than the high plausibility premise, whereas one participant questioned it less often (and there were three ties); this result is very highly significant ($z = 5.12$ $p < 2 \times 10^{-7}$, sign test).

A between- subject comparison was also conducted as follows. The two sub-groups defined by the two orders of presentation of the problems were compared against each other. Given that the direct sequence was D = {H1, L2, H3, L4, L1, H2, L3, H4} and the reverse sequence R = {H4, L3, H2, L1, L4, H3,

L2, H1}, performance on problems H1+H3 from D was compared with performance on problems L1+L3 from R, and similarly performance on problems H2+H4 from R was compared with performance on problems L2+L4 from D. In that way, performance was always based on problems borrowed from the first block, that is, problems solved before their paired counterpart had been seen. To analyse performance, the dependent variable that was chosen was the number of participants who answered by revising the major premise of both problems such as L1+L3, etc. This choice allowed the use of chi-square statistics in spite of the fact that measures were repeated across subjects. Results are clear-cut as shown in Table 3.

———————————

Insert Table 3 about here

———————————

For problems 1 and 3 as well as for problems 2 and 4, most participants revised both low plausibility major premises, whereas only a minority did so for high plausibility major premises (χ^2 = 18.7, p < .001 and χ^2 = 5.35, p < .05, respectively).

There are two other questions of interest in relation to the first experiment. One regards the choice of the premise to revise: after pooling both levels of plausibility, again a majority of choices (67,4%) concerned the major. Table 4 presents the distribution of the frequencies of revision (in percent) as a function of the level of plausibility, of the type of premise revised (major or minor), and of the response option (a doubt or a denial).

———————————

Insert Table 4 about here

———————————

Comparing across levels of plausibility, for high plausibility problems there were roughly equal percentages of revision of the major (43.8) and of the minor (56.2)

but for low plausibility problems there was a sharp superiority in the percentage of revision of the major (91 percent) over the minor (9 percent).

The other question regards the kind of revision expressed by participants (uncertain or categorical). Overall, there were roughly as many expressions of doubt (51.7 percent) as of denial (48.3 percent). This was the case for the high plausibility problems as well as for the low plausibility problems separately. However, this was not the case for the two kinds of premise separately, since the major was significantly more often doubted than denied (Wilcoxon test applied within subjects, $z = 1.98$, $p < .05$) while there was an opposite tendency (albeit non significant) for the minor premise.

Discussion.

As expected, there was a sharp difference in the rate of revision of the major premises as a function of their credibility. They were more frequently revised when they were less plausible. Since in this experiment judgments of plausibility crucially depended on world knowledge, and the logical form of the premises as well as the form of the arguments were kept constant, the hyposis that epistemic entrenchment is based solely on formal properties of the propositions seems hard to maintain.

The problem of which premise people are more prone to revise, whether the major or the minor, is highly dependent on the question of the major premise credibility, since a variation in the rate of revision of the major automatically implies a variation in the opposite direction in the rate of revision of the minor, by complementarity. This experiment shows that one cannot even discuss this problem without taking into account the credibility of the major premise. It sheds some light on the question raised by the first experiment: the major was much more often revised than the minor only when it had a low plausibility. However,

the fact that even with a high plausibility the major was revised over 40 percent of the time attests to its frail entrenchment.

Finally, the question of the choice of the revision mode, by doubting the premise or by categorically denying it is somehow clarified by the results. The major premise expressed a conditional relation referring to daily life. Participants doubted it more often than they denied it. In contrast, this did not occur for the minor premise, which expressed a fact. So, more complex propositions that convey relations rather than facts, that refer to a well-known domain and that are plausible are more susceptible to being doubted than denied, a view that agrees with Elio and Pelletier's. But notice that the tendency to doubt rather than to deny remains powerful since overall, roughly one half of all the responses in the second experiment expressed uncertainty (while the denial-to-doubt ratio was about one to five in the first experiment).

General discussion

Assuming that belief is a matter of degree, it is reasonable to hypothesise that the aim of belief change is not necessarily to give up belief but more comprehensively to alter the degree of belief. If that is the case, all-or-nothing *formats of response* are too crude to evaluate belief change. This is a methodological point which should not be confused with the issue of the *format of representation* of belief strength, a theoretical problem.

Empirically, the first experiment addressed the question of how to better capture people's expression of belief change. It was inspired by the notion that premises and conclusion of deductive reasoning often have degrees of belief that people can express, and by the observation that participants of Elio and Pelletier's (1997) studies showed a tendency to express degrees of disbelief rather than sheer denials, even though they were not always given a fair chance

to do so. The main aim of the first experiment was to test the hypothesis that to question the truth of a premise, qualifying it as doubtful is at least as natural a way as qualifying it as false. This turned out to be the case first with a set of propositions containing arbitrary rules of categorisation in a fictitious environment, for which over three quarters of the revisions were observed to be uncertain. This was confirmed in the second experiment where the environment referred to everyday life and the propositions expressed conditions for a decision to be made, for which there were about one half of uncertain revisions, which suggests that the claim that is being made is susceptible to generalisation.

Theoretically, there are opposite views regarding the representation of belief. Harman (1986) endorses the notion that beliefs have degrees. However, he claims that degrees of belief are implicit (that is, obtainable by inference) and have the status of an epiphenomenon resulting from the revision and the reasoning processes which act on all-or-nothing representations. This is not of course the only possibility. The other main hypothesis is that principles of reasoning and revision act on degrees of belief directly. Numeric formalisms using probability are based on this view.

In the all-or-nothing case, inconsistency results from a clash between belief in the new information and disbelief in the old one and coincides with contradiction. The degree-of-belief approach faces the problem of defining a criterion of inconsistency for two different degrees of belief where at least one of them is not full belief or full disbelief. However, inconsistency could be a graded concept too, that is, an increasing function of the difference in belief between old and new information with a threshold value that triggers revision in the belief set. The problem may also receive a solution if one adopts a connectionist approach of the type investigated by Thagard (1989). In that case, the degree of belief in a

proposition could be represented by the level of excitation of a unit and a change in the degree of belief would result in an adjustment of the level of excitation of the other units. This leads us to specify which propositions are assumed to be affected. In principle, all propositions in the set may incur a change in belief. Although the first four options in the multiple choice of the first experiment distinguished single premises, the fifth option offered an opportunity to express wider changes; indeed, about 15 to 20 percent of the participants decided to alter both premises; that only a minority did so seems to reflect the psychological plausibility of the principle of minimal change.

The other main point addressed in this paper revolves around the notion of epistemic entrenchment, the property of some premises to better resist abandonment than others. As recalled in the introduction, one approach assumes that it depends on formal relations between propositions. It is therefore a reasonable investigation to compare the relative entrenchment of the major conditional premise against the minor premise of elementary arguments. Elio and Pelletier found that the former were more readily abandoned than the latter. This result was confirmed in both experiments, but it seems more economical at this stage to link it to the syntactic characteristic of the major of the arguments under consideration, namely the presence of a connective, rather than to a specific feature of the conditional. This conclusion stems from the result of the first experiment in which arguments with a non conditional major premise led to the same choices of premise to revise as arguments with a conditional major premise and from that of the second experiment in which all arguments had conditional premises as a major but led to different choices of premise to revise as a function of the credibility of the major.

However, even though it seems confirmed that some formal, possibly syntactic, characteristics of propositions (such as connected propositions as opposed to atomic ones) affect their entrenchment, the notion that belief has varying degrees deeply alters the determination of the entrenchment of premises, as well as the concept itself. As shown in the second experiment, from the moment there are premises which are more or less implausible, these premises are more or less susceptible of being abandoned. In brief, in addition to formal factors, content factors affect people's change of belief. Of course, not all AI formalisms adopt an all-or-nothing representation of belief. The psychological literature reviewed above and the present results tend to suggest that there is more psychological plausibility in the degree-of- belief approach. Consequently, entrenchment as defined solely by formal properties seems to have limited psychological import. On the other hand, viewed as people's unwillingness to give up a premise for reasons linked with content, it is psychologically highly relevant.

References

Bernard, J.-M. (1998). Bayesian inference for categorised data. In Rouanet, H. , Bernard, J.-M. , Bert, M.-C. , Lecoutre, B. , Lecoutre, M.-P. , & Le Roux, B. (Eds.), *New ways in statistical methodology* (pp. 159-226). Bern: Peter Lang.

Chan, D. , & Chua, F. (1994). Suppression of valid inferences: Syntactic views, mental models, and relative salience. *Cognition, 53,* 217-238.

Cummins, D. D. (1995). Naive theories and causal deduction. *Memory and Cognition, 23,* 646-658.

Cummins, D. D. (1996). Evidence of deontic reasoning in 3- and 4-year-old children. *Memory and Cognition, 24,* 823-829.

Cummins, D. D. , Lubart, T. , Alkinis, O. , & Rist, R. (1991). Conditional reasoning and causation. *Memory and Cognition, 19,* 274-282.

Elio, R. (1997). What to believe when inferences are contradicted: The impact of knowledge type and inference rule. *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 211-216). Mahwah, N. J. : Lawrence Erlbaum.

Elio, R. (1998). How to disbelieve P-->Q: Resolving contradictions. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 315-320). Mahwah, N. J. : Lawrence Erlbaum.

Elio, R. , & Pelletier, F. J. (1997). Belief change as propositional update. *Cognitive Science, 21,* 419-460.

Galliers, J. R. (1992). Autonomous belief revision and communication. In P. Gärdenfors (Ed.), *Belief revision* (pp. 220-246). Cambridge: Cambridge University Press.

Gärdenfors, P. (1988). *Knowledge in flux.* Cambridge: Cambridge University Press.

Gärdenfors, P. (1992). Belief revision: An introduction. In P. Gärdenfors (Ed.), *Belief revision* (pp. 1-28). Cambridge: Cambridge University Press.

George, C. (1995). The endorsement of the premises: Assumption-based or belief-based reasoning. *British Journal of Psychology, 86,* 93-111.

George, C. (1997). Reasoning from uncertain premises. *Thinking and Reasoning, 3,* 161-189.

Harman, G. (1986). *Change in view.* Cambridge, Ma: MIT Press.

Liu, I.-M. , Lo, K.-C. , & Wu, J.-T. (1996). A probabilistic interpretation of "If-Then". *Quarterly Journal of Experimental Psychology, 49A,* 828-844.

Politzer, G. , & George, C. (1992, July). *Reasoning under uncertainty and non monotonicity: A pragmatic account.* Paper presented at the 25th International Congress of Psychology, Brussels.

Stevenson, R. J. , & Over, D. E. (1995). Deduction from uncertain premises. *Quarterly Journal of Experimental Psychology,* 48A, 613-643.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences, 12,* 435-467.

Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory and Cognition, 22,* 742-758.

Appendix

Experiment 1. Example of a revision problem (case of a Modus Ponendo Tollens
   argument).

You are going to the closest city. On your way you meet a first informant who

tells you:

- The city is not inhabited both by Klingons and Bagors.

Soon after you meet a second informant who tells you:

- The city is inhabited by Klingons.

Of course, you conclude that the city is not inhabited by Bagors.

However, on arriving the next day you notice that the city is inhabited by Bagors.

How to reconcile all these data? In your opinion:

1) It is false that the city is inhabited by Klingons.

2) It is not sure that the city is inhabited by Klingons.

3) It is false that the city is not inhabited both by Klingons and Bagors.

4) It is not sure the city is not inhabited both by Klingons and Bagors.

5) Other:

What are the reasons of your choice?

Experiment 2. The four pairs of conditional sentences (translated from French) in

their high plausibility version (a) and low plausibility version (b).

1) a. If his car is beyond repair, Alex will change his car.

    b. If his neighbour buys a new car, Bastien will change his car.

2) a. If the teacher is absent, the lecture will be canceled.

    b. If a chair is lacking, the lecture will be canceled.

3) a. If she is the first in her class, Aline will be admitted to the higher   form.

   b. If she is polite with the teacher, Amélie will be admitted to the   higher

form.

4) a. If Isabelle feels very sick, she will leave before the end of the show.

   b. If Valérie has no more pop corn, she will leave before the end of the show.

Footnotes

1. The computation is based on BAYCAT (Bernard, 1988). Credibility intervals
   are similar to confidence intervals. The associated medians indicate that the
   ratio f1/f2 has a probability of .50 to be inferior to the median and a probability
   of .50 to be superior.

Table 1.

Experiment 1. Frequency distribution (in percent) of the revision choice type for each of the four argument forms. N = 80.

argument form	revision choice type			
	certain	uncertain	unclas-sified(*)	total
P --> Q; P; /Q	12.5	80.0	07.5	100
P --> Q; ¬Q / ¬P	17.5	77.5	05	100
P or Q; ¬P /Q	15	76.2	08.7	100
not both P&Q; P /¬Q	21.2	77.5	01.2	100
mean	16.6	77.8	05.6	100

* A few choices of option 5 did not allow classification as certain or uncertain

Table 2.

Experiment 1. Frequency distribution (in percent) of the revision choice for each

of the four argument forms. N = 80.

	revision choice				
argument form	minor	major	minor and major	other (*)	total
P --> Q; P; /Q	27.5	46.2	18.7	07.5	100
P --> Q; ¬Q / ¬P	20	61.2	16.2	02.5	100
P or Q; ¬P /Q	30	46.2	15	08.7	100
¬ (P&Q) ; P/ ¬Q	21.2	58.7	18.7	01.2	100
mean	24.7	53.1	17.2	05.0	100

* A few choices of option 5 did not allow classification

Table 3.

<u>Experiment 2. Frequency of revision of the major premise as a function of the</u>

<u>plausibility of the premise.</u>

problems 1 + 3

	<u>plausibility of major premise</u>	
number of problems where major is revised	High	Low
0 or 1	17	3
2	2	14

problems 2 + 4

	<u>plausibility of major premise</u>	
number of problems where major is revised	High	Low
0 or 1	10	4
2	7	15

Table 4.

Experiment 2. Frequency of revision (in percent) as a function of level of

plausibility, type of premise revised, and response option.

	premise revised				
	major		minor		
	response	option	response	option	
Plausibility	*false*	*uncertain*	*false*	*uncertain*	
High	13.2	30.6	34.0	22.2	100
Low	43.1	47.9	06.3	02.7	100