

Why Believability Cannot Explain Belief Revision

Uri Hasson (uhasson@princeton.edu)
Psychology Department, Princeton University
Princeton, NJ 08540 USA

Philip N. Johnson-Laird (phil@princeton.edu)
Psychology Department, Princeton University
Princeton, NJ 08540 USA

Abstract

A common view in epistemology is that some beliefs are more entrenched than others. This view is plausible, but we show that it fails to explain which statements individuals tend to doubt when an incontrovertible fact is inconsistent with the relevant set of statements. We report three studies that each show that the believability of statements is influenced by context. Given a conditional of the form *If P then Q* and a categorical statement *P*, individuals tend to judge the categorical as more believable than the conditional. But, when the same statements are followed by an incontrovertible fact, *not-Q*, that is inconsistent with them, individuals tend to judge the conditional as more believable than the categorical. The theory of mental models accounts in part for these and other results of the experiments, including a study of the believability of exclusive disjunctions and categoricals.

Introduction

Three fundamental questions in epistemology are: what properties do beliefs have, what counts as justification for beliefs, and how are beliefs revised in the light of incontrovertible evidence? According to the 'coherence' approach to belief revision (e.g., Harman, 1986), a belief is justified so long as it logically coheres with other beliefs. When incoherence arises, changes made to beliefs should be as conservative as possible so that the accommodation of the new information is accompanied by the rejection of a minimal number of beliefs (*minimal change*).

Beliefs can have different degrees of *epistemic entrenchment* (Gärdenfors, 1988). Hence, when you must abandon one of a set of beliefs, you should abandon the one that is the least entrenched. Several factors are likely to determine a belief's degree of entrenchment. One factor is the number of reasons supporting the belief. But even if there are no grounds to think that one belief is more justified than another, other factors may affect believability. One such factor has been termed informativeness (Bar-Hillel & Carnap, 1964). The informativeness of a statement is a function of the number of states of affairs that the statement rules out: the more possibilities ruled out, the more informative the statement. For example, a statement of the form *A or B or both* rules out fewer possibilities than a statement of the form *A*, and is therefore less informative. The less information that a statement conveys, the more probable it is. And the probability of a statement might

affect its believability. Statements that are more probable may seem more believable. Hence, a statement of the form *A or B or both* should be more believable than *A*, because the former is more probable than the latter.

Recent work has examined possible links between statement-form and believability. For example, Elio and Pelletier (1997) have suggested that conditional statements may be less believable than categorical statements, and Politzer and Carles (2001) have contended that any compound statement, i.e., one containing a sentential connective, may be less believable than a categorical statement. Such accounts are context-independent in that they suggest that belief revision can be determined by the intrinsic properties of statements, such as their syntactic form. They do not ascribe any role to the reasoning processes through which inconsistencies are recognized and resolved.

In this paper we put forward a process-account of belief revision. Its central tenet is that perceptions of believability are strongly affected by the processes of reasoning that resolve inconsistencies. Two principles underlie this account. First, the resolution of inconsistency is not a single deterministic process: multiple strategies exist to recognize inconsistencies and to resolve them. Moreover, the different ways of recognizing inconsistencies can lead to different revision outcomes. Second, the processes of resolution depend on the representation of beliefs by mental models, and so resolution is partially determined by correspondences between the models of the belief set and the model of the incoming evidence.

Our goals in the present paper are to review recent studies of the believability of different sorts of statements (part 1), to outline the model-based theory of belief revision and its principal predictions (part 2), and to describe three experiments designed to test these predictions (part 3).

1. Statement Believability and Contexts of Reasoning

In a pioneering study, Elio and Pelletier (1997) demonstrated that given an inconsistency between a set of statements and an indisputable fact, participants were more likely to abandon a conditional statement than a categorical one. The participants were presented with two statements of the following form:

1. If P then Q
2. P

These statements were followed by an incontrovertible statement that was inconsistent with the two statements taken together, i.e., they could not both be true given the fact:

3. [Fact] Not Q

In this case, the participants (in Experiment 2) tended to doubt, or to disbelieve the conditional statement rather than the categorical statement. Elio and Pelletier suggested that conditionals usually express contingencies, and so a counterexample to the conditional is likely to lead to its rejection. Politzer and Carles (2001) similarly showed that individuals tend to doubt a disjunctive statement rather than a categorical statement. They put forward an alternative account: participants are more likely to doubt a major premise, and this doubt “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.” (p. 224). The fundamental difference between the two accounts is that Politzer and Carles attribute the differences in believability to intrinsic, *context-independent* differences between types of statement. In contrast, Elio and Pelletier attribute the difference to a context in which a counterexample is salient. Their account leaves room for a different pattern of believability in a context in which there is no salient counterexample. An alternative and less interesting explanation is that the participants were more likely to reject the conditional statements because they were always presented first in each problem.

2. Towards a Processing Account of Belief Revision

Reasoning and Nondeterminacy

The first principle of our account is that belief revision is based on the ability to reason. Given a set of beliefs and a fact that is inconsistent with them, the reasoning that underlies the revision of beliefs can be carried out in more than one way, and, at present, no theoretical account exists to fix precisely in which way it occurs on any particular occasion. Furthermore, different strategies for recognizing inconsistencies can yield different patterns of belief revision. For example, given statements 1-3 above, the inconsistency can be recognized in at least three ways. First, you could infer *Q* from the first two statements (an inference of the form known as *modus ponens*), and see that its conclusion is inconsistent with the fact in the third statement. Second, you could infer *not-P* from the fact and the first statement (an inference of the form known as *modus tollens*), and see that its conclusion is inconsistent with the second statement. Third, you could conjoin the second statement with the fact to form: *P and not-Q*, and see that it is inconsistent with the conditional. These different strategies for recognizing the inconsistency are likely to yield different resolution patterns (see Harman, 1986, for a similar point). The principle of indeterminacy makes available multiple strategies for recognizing and resolving inconsistencies.

Model-based Representation

Theories of reasoning need to identify how information is mentally represented and what processes operate on these representations. According to the theory of mental models, statements are mentally represented in the form of models of the possibilities compatible with them. Each model represents a different possibility, but individuals tend to focus on certain salient possibilities and to represent within them clauses in the premises only when they are true in the relevant possibility (Johnson-Laird & Byrne, 1991). Hence, correspondences in the models of beliefs and evidence are likely to affect the revision process. We can illustrate this hypothesis by way of an example.

According to the model theory, a conditional statement of the form: *If P then Q* is represented in two mental models. One model explicitly represents the salient possibility in which *P* and *Q* co-occur, but the other model has no explicit content. It is a place-holder for the possibilities in which the antecedent of the conditional is false, and so the implicit model has a footnote to that effect. The two mental models are accordingly:

<i>P</i>	<i>Q</i>
...	

where the ellipsis denotes the implicit model. If the representation of the incontrovertible evidence matches these models of the conditional, then the conditional and evidence are consistent. In contrast, if the evidence has the form:

Not-*Q*

then it mismatches the mental model of the possibility of *P* and *Q*. Individuals will be tempted to infer that the evidence and the conditional are inconsistent, and so they may be tempted to reject the conditional. In fact, the inconsistency is apparent rather than real. If the implicit model is fleshed out into fully explicit models of the possibilities in which the antecedent of the conditional is false, then one such model represents the possibility of *not-P* and *not-Q*.

Some findings from the literature are consistent with this account. For example, Elio and Pelletier (1997) found that the tendency to reject the conditional statement varied with the way in which a counterexample to the conditional was presented. When participants were presented with configurations in the schematic form of (1-3), 35% of the participants preferred to disbelieve the conditional statement and believe the categorical statement. However, when the second and third statements were switched to the order:

4. If *P* then *Q*

5. Not *Q*

6. [Fact] *P*

only 12% of the participants chose to deny the conditional statement and believe the categorical one. In the first case, the incontrovertible fact, *not-Q*, mismatches the first mental model of the conditional: *P Q*. In the second case, the incontrovertible fact, *P*, matches the first mental model of the conditional.

Elio and Pelletier, however, suggested that the difference between the two cases originates in a parallel account based on the syntax of the sentences. When the fact is of the form: *P* (6), “neither sentence is explicitly denied by the fact.” In

contrast, when the fact is of the form: *not-Q* (3), the conditional statement contains a clause of the form: *Q*. While the syntactic account and the model account make the same predictions for the present cases, they make different predictions about cases in which the conditional contains a negative consequent, such as case 7-9 and case 10-12, in Table 1. In general, we refer to problems in which the two initial statements allow a modus ponens inference as MP configurations, and to problems in which the two initial statements allow a modus tollens inference as MT configurations. Of course, individuals may not make these inferences in the course of resolving the inconsistency, because, as we have illustrated, they have other possible strategies available to them.

Table 1: Configurations with negated conditionals

MP (negated consequent)	MT (negated consequent)
7. If P then Not Q	10. If P then Not Q
8. P	11. Q
9. [Fact]: Q	12. [Fact]: P

According to the syntactic account, the conditional in the cases in Table 1 should be more likely to be revised for the MP problems than for MT because in the MP configuration it is explicitly denied by the fact (as was the case for the affirmative conditional).

In contrast, according the mental-models account, a conditional such as *if P then not-Q* is likely to prompt the representation of two possibilities:

P	not-Q
Not-P	Q

If the second model is constructed, then the fact, *Q*, in the MP case matches the second model of the conditional, but the fact, *P*, in the MT case mismatches this second model of the conditional. This pattern is opposite to the one predicted for problems based on affirmative conditionals. Therefore, the model theory predicts that the tendency to revise a conditional would be the product of an interaction between the valence of the conditional and the form of the inference. This account posits that revision is affected by dynamic relations between a belief-set and a given fact. As a result, it is unclear whether independent judgments of believability for statements in the belief-set alone predict which statement is revised given an additional fact inconsistent with the beliefs.

3. Empirical Studies

Experiment 1: Reasoning Affects Believability

This study evaluated the model theory's predictions about mismatches, and whether patterns of believability depend on the presence or absence of an inconsistent fact.

Method. Thirty-two Princeton University undergraduates made judgments of believability in two types of tasks. In one task, they were presented with two statements that were consistent with each other, and were asked to indicate which statement they found more believable (the "two-statement"

task henceforth). The statements were made by two different speakers, e.g.:

Speaker A: If the experiment was conducted according to procedure then the helium is in liquid form.

Speaker B: The experiment was conducted according to procedure.

After they had read the statements, the participants judged which of the two they found more believable.

In the other task, the participants were presented with two statements such as those above but with a third statement, which was presented as a fact (the "three-statement" task henceforth). For example, given the two statements above, the third statement was:

You know for a fact that the helium is not in liquid form.

In this task, the participants judged whether the fact was consistent with the first two statements taken together. If they detected the inconsistency, they indicated which of the two initial statements was more believable. On all the experimental trials, the additional fact was inconsistent with the two initial statements. In both tasks, the participants made their judgments by a joint rating of the two initial statements on a six-point scale ranging from 1 (*completely believe statement A*) to 6 (*completely believe statement B*). In all experimental trials, one initial statement was a conditional and one was a categorical, and the two statements yielded either an MP or a MT inference.

The participants acted as their own controls and carried out one block of two-statement trials and one block of three-statement trials. The order of the two blocks was counterbalanced between the participants. Within each block, there were eight sorts of problem depending on whether the initial statements yielded an MP or MT inference, whether the conditional statement occurred first or second, and on whether the consequent of the conditional was affirmative or negative. The presentation of the problems in each block was in two counterbalanced orders. Sixteen thematic contents were rotated over the experimental materials. Filler problems were included to ensure that the participants critically evaluated whether a fact was consistent with the preceding two premises. The filler problems included facts that were either consistent with both statements, or clearly consistent with just one statement, but not the other.

Results and Discussion. The results showed that the task had a marked effect on the believability of statements. In the two-statement task, the participants were biased toward believing the categorical statement (19 out of the 32 participants had this bias, 4 went against it, and there were 9 ties, Wilcoxon, $z = 3.37$, $p < .001$). In contrast, in the three-statement task, they were biased towards believing the conditional statement (17 out of the 32 participants had this bias, 9 went against it, and there were 6 ties, Wilcoxon, $z = 2.10$, $p < .05$). This difference in the biases shows that the rejection of a belief in the case of an inconsistency does not depend on its intrinsic believability.

The statement revised given an inconsistency was not necessarily the one judged as less believable when there was no inconsistency. This phenomenon resolves one of the main questions that prompted this study: the reasoning

context can affect the relative believability of conditionals and categorical statements.

We converted the ratings on the 6-point scale to indicate the degree of bias towards believing the conditional assertion (which could range from +0.5 to +2.5 for bias towards believing the conditional statement and from -0.5 to -2.5 for bias towards believing the categorical statement). For the two-statement task, there were no effects of the order of the two tasks, and there were no effects of any of the main variables (a mixed ANOVA yielded F 's < 1). The analysis of the three-statement task revealed that different patterns of judgments were made depending on whether this task was done before or after the two-statement task. We first present the results for the three-statement task when it was performed in the first block of trials (Figure 1).




Figure 1: Effects of inference-form and valence

There was no main effect of the form of the inference permitted by the initial statements. But, as the figure shows, negating the consequent of a conditional increased its believability for MP trials, but decreased its believability for MT trials, and the corresponding interaction was reliable; $F(1,15) = 13.99$, $p < .01$. The results support the notion that the believability of statements in the case of an inconsistency depends on the mismatch principle.

One unexpected finding was that the bias towards believing the conditional was stronger when it was the first statement (mean bias of 0.61) than when it was second statement (mean bias of 0.12; $F(1,15) = 4.5$, $p < .05$). But, the order of statements interacted with the form of the inference; $F(1,15) = 4.9$, $p < .05$. In particular, for MP problems, the conditional was more believable than the categorical when it was presented first (mean bias of 0.72), but less believable than the categorical when the categorical was presented first (mean bias of -0.19). However, for MT problems, the effect of order was minute (mean biases of 0.50 and 0.44 respectively).

None of the effects reported above occurred for the participants who carried out the three-statement task after the two-statement task. The transfer effects are another indicator that participants reason in different ways depending on the context.

In sum, statements in a context of inference to resolve an inconsistency affected their relative believability in comparison with their believability in a context in which there was no inconsistency. These findings alone indicate that any account of believability and belief revision needs to take into account inferential processes. The mismatch-principle was corroborated as shown in Figure 1. The effect of the order of the two initial statements is problematic for theories that postulate rational principles for the maintenance of consistent beliefs. Our results are incompatible with those previous findings that showed that conditionals were more likely to be doubted in the face of inconsistency. But, in these previous studies, the first two statements were always inconsistent with the information presented as a fact. As a result, the participants did not have to employ inferential processes in order to determine whether or not the statement of fact was consistent with the initial statements. Hence, the choice of which statement to disbelieve or doubt could be made without the need to draw an inference. On this view, it is unsurprising that the results of the previous studies match our results in the two-statement task. It is possible that when participants do not have to reason, the categorical statement seems more believable than the conditional statement.

Experiment 2: Strategies in Belief Revision

The principle of indeterminacy postulates that individuals can use different strategies in reasoning, and so they may use different strategies to resolve inconsistencies. The aim of the present experiment was to test this prediction. It also aimed to determine whether different strategies lead to differences in the revision of beliefs. The experiment also allowed us to replicate the results of the previous experiment, and to investigate why reasoning in the face of inconsistency enhances the believability of conditional statements.

Method. We tested thirty-two new participants from the same population as before. They evaluated the believability of statements in two blocks of three-statement trials. In one block, they indicated which statement they found more believable given the inconsistency (henceforth, the "no justification" task). But, in another block of three-statement trials, they also had to explain the reasons for their choice (henceforth, the "justification" task). The order of the blocks was counterbalanced between participants. The materials and the procedure were the same as those in the previous experiment.

Results and Discussion. The results showed again that in the case of an inconsistency with the facts the participants found conditional statements significantly more believable than categorical statements. The bias was stronger in the justification task (mean bias of 1.06) than in the no justification task (mean bias of 0.62, $p < .05$).

The participants' written justifications revealed that they used two main sorts of strategies to resolve inconsistencies: In the first sort of strategy, they used the semantic content of the conditional and the categorical statement, and their general knowledge to assess their believability. For example, one participant argued that the conditional statement *if the turbine is operating properly then it rotates*

at over 6000 rounds per minute was not believable, since it meant that the turbine was turning too fast. In this case, the strategy assesses the plausibility of the assertions in the light of knowledge. A more general semantic consideration concerned the *form* of the assertions. For example, some participants argued that conditionals are more hypothetical and therefore less believable than categoricals, whereas others argued that conditionals expresses less commitment and are therefore more believable. All the considerations based on meaning are equally applicable to cases in which there are no inconsistencies (as in the two-statement task). Yet, the participants who explained their judgments by appealing to semantic considerations *never* explained how they had recognized the inconsistency in the first place. They did not appear to consider the method of detecting an inconsistency to be relevant to its resolution. Once they had detected an inconsistency, however, they resolved it through semantic considerations.

The second main sort of strategy was *inferential*. The participants made clear reference to the statement of incontrovertible fact in their explanations of how they detected or resolved the inconsistencies. By definition, such strategies are not available if there is no inconsistent fact. We identified three such strategies. In the *initial-statements* strategy, the participants drew their own conclusions from the initial conditional and categorical statements, and noticed that it was inconsistent with the fact. In the *fact-conditional* strategy, the participants drew a conclusion from the fact and the conditional statement and noticed that the conclusion was inconsistent with the categorical statement. In the *correspondence* strategy, the participants did not specify how they detected the inconsistency, but based their decisions on certain correspondences between the conditional statement and the fact (i.e., matches or mismatches).

Inferential strategies ought to be accessible to most people. In fact, the majority of participants used an inferential strategy at least once (24 participants out of 32; Binomial, $p < .001$). Semantic strategies were also often used; twenty-four participants used semantic strategies at least once. Overall, inferential strategies were used in 47% of the cases, and semantic strategies were used in 53% of the cases. Indeed, inferential strategies were not only available, but were routinely used so that participants did not rely solely on semantic argumentation in their decisions. This finding establishes that a context of an inconsistency leads to strategies for evaluating believability that are otherwise unavailable.

Table 2 summarizes the frequency with which the participants used the various strategies and the concomitant bias in favor of believing the conditional statement. The results show that the different strategies for *recognizing* inconsistencies did yield different ways of resolving the inconsistencies. For those participants who used both semantic and inferential strategies, the bias in favor of believing conditionals was reliably greater with an inferential strategy. Of 18 participants who used both strategies, 15 showed a greater bias for the conditional statement when they used inferential strategies (Binomial, $p < .01$). As Table 2 shows, when the participants employed

semantic strategies, the two statements were rated as about equally believable (mean bias of only 0.02). However, when participants based their decisions on inferential strategies, the bias in favor of the conditional was greatly increased (mean bias of 2.09).

Table 2: Bias in favor of conditional statements as a function of strategy, and the proportion of responses based on the use of the different strategies.

	Bias for conditional	Proportion of responses
Semantic Strategies		
Plausibility	0.10	45%
Form of assertions	-0.64	8%
<i>Means</i>	0.02	
Inference Strategies		
Fact-conditional	2.68	30%
Correspondence	0.73	11%
Initial statements	-0.33	6%
<i>Means</i>	2.09	

We conjecture that when individuals adduce knowledge (using a semantic strategy), conditional claims are vulnerable -- if only because knowledge is less likely to gainsay a specific categorical claim about a hypothetical entity, e.g., *the turbine is operating properly*. Likewise, when they consider conditionals as hypothetical assertion, they are less likely to maintain them in the face of inconsistency. But, when individuals infer a conclusion from the conditional and the incontrovertible fact, which they then recognize as incompatible with the categorical premise, they follow the mismatch principle, and reject the categorical premise. Similarly, when they make an inference from the initial statements, and detect the inconsistency of the conclusion with the fact, they again follow the mismatch principle and reject the conditional. These conjectures, however, should be the focus of further studies. Nonetheless, the study shows that participants used different strategies to resolve inconsistencies, and these strategies are linked to how believable they found the given statements.

Experiment 3: The Believability of Disjunctions and Categoricals

We have shown that the evaluation of conditional and categorical statements depends on context. The task, the order of the statements, and the correspondence between the conditional and the fact, all affected believability. The present study extended the scope of these findings by examining disjunctions.

Method. This study was analogous to our first experiment. In the two-statement task, the participants were presented with an exclusive disjunction and a categorical statement, and rated the joint believability of the statements on a scale of 1 - 6. In the three-statement task, the participants first judged whether the two statements were consistent with an indisputable fact, and made the rating only if they found the three statements to be inconsistent. Here is an example of a three-statement trial:

Speaker A: The experiment was conducted according to procedure or else the helium is in liquid form, but not both.

Speaker B: The experiment was conducted according to procedure.

Fact: You know that the helium is in liquid form

Ninety-two Princeton University undergraduates participated in this study. The two sorts of task were in separate blocks, and the order of the blocks was counterbalanced between participants. Within each block, we manipulated two factors: The form of the disjunction (p or else q vs. p or else not- q) and the form of the categorical premise (p vs. not- p). Each combination of a disjunction and a categorical statement was paired with a fact that was either consistent or inconsistent with the prior statements. The study therefore included four types of problem in which fact was inconsistent with the prior statements, and four types of problem in which the fact was consistent with the prior statements.

According to the model theory, an exclusive disjunction such as P or else Q but not both calls for two mental models:

[P] ...
... [Q]

The brackets signify that the bracketed constituent appears only in that possibility. In addition, mental models represent only what is true in each possibility. For instance, the first model represents that P is the case, but it does not represent that in this possibility Q is false. It follows that in some cases of inconsistency, the mental model of the fact matches one of the mental models of the disjunction (e.g., p or else q , p , fact: q), whereas in others cases the mental model of the fact does not match either mental model of the disjunction (e.g., p or else q , not- p , fact: not- q). Participants may also misjudge consistent cases as inconsistent (e.g., p or else q , not- p , fact: q) as inconsistent, and the predictions of the mismatch principle should hold for such cases as well.

Results and Discussion. Error rates on problems were high (33%), which is typical for reasoning with disjunctions. We therefore evaluated data for those 54 participants who correctly solved at last three inconsistent problems (from a maximum of 4). Evaluations of believability were again found to depend on context. For the two-statement trials, the participants were biased to believe the categorical statement (42 participants demonstrated this pattern, 4 went against it, and there were 8 ties; Wilcoxon $z = 4.92$, $p < .001$). However, for the three-statement trials, the participants were biased to believe the disjunctive statement (34 participants demonstrated this pattern, 11 went against it, and there were 9 ties; Wilcoxon $z = 3.74$, $p < .001$).

In addition, the disjunctions were relatively more believable when the incontrovertible fact matched one of their mental models than when it mismatched one of their mental models. This effect was marginal for inconsistent problems (mean biases were 0.66 vs. 0.32, one-tailed $t(53) = 1.37$, $p = .08$), but significant for consistent problems that were wrongly recognized as inconsistent (mean biases were 0.88 vs. 0.01; one-tailed $t(24) = 2.26$, $p = .02$).

4. General Discussion

We set out to evaluate the notion of the ‘epistemic entrenchment’ of beliefs, i.e., their intrinsic believability in the light of knowledge. Undoubtedly, some beliefs are more entrenched than others. In general, as our experiments have shown, categorical assertions were more believable than the conditionals themselves; likewise, categorical assertions were more believable than exclusive disjunctions in which they occurred. This latter result shows that individuals do not necessarily take probability into account in assessing the believability of a statement, because the probability of A cannot be greater than the probability of A or else B . Our investigation shows, however, that there is a shift in believability when statements occur in the context of an incontrovertible fact that is inconsistent with them. In this context, individuals tend to judge a compound assertion, whether it is a conditional or a disjunction, as more believable than the categorical statement. Hence, the reasoning that detects inconsistency affects believability. Individuals use a variety of strategies in their reasoning, and they in turn influence belief revision too.

One move that could be made is to separate the notion of confidence in a belief from the notion of resistance to change, and to limit entrenchment to the latter (see Hansson, in press). But, it would then be necessary to determine the relation between confidence and resistance to change. Our preference has been to explain the shift in believability as a consequence of matches, or mismatches, between the facts of the matter and mental models of the statements.

Acknowledgements

This research was supported in part by a grant from the National Science Foundation (BCS-0076287) to study strategies in reasoning. We thank Ruth Byrne, Sam Glucksberg, Geoff Goodwin, Catherine Haught, Susanna Reynolds, Louis Lee and Clare Walsh for their advice.

References

- Bar-Hillel, Y., & Carnap, R. (1964). Semantic information and its measures. In Y. Bar-Hillel (Ed.), *Language and information*, (pp. 298-312). Reading, MA: Addison-Wesley.
- Elio, R. & Pelletier, F. J. (1997). Belief change as propositional update. *Cognitive Science* 21(4), 419-460.
- Gärdenfors, P. (1988). *Knowledge in Flux*. Cambridge, MA: MIT Press.
- Hansson, S. O. (2003). Ten philosophical problems in belief revision, *Journal of Logic and Computation*, in press.
- Harman, G. (1986). *Change in view: Principles of reasoning*. London, England: Bradford/MIT Press.
- Johnson-Laird, P. N., & Byrne, R. M. J. (1991). *Deduction*. Hove (UK); Hillsdale (USA): L. Erlbaum.
- Politzer, G., & Carles, L. (2001). Belief revision and uncertain reasoning. *Thinking and Reasoning*, 7(3), 217-234.