

## **UC Merced**

# **Proceedings of the Annual Meeting of the Cognitive Science Society**

### **Title**

The Source of Beliefs in Conflicting and Non Conflicting Situations

### **Permalink**

<https://escholarship.org/uc/item/7rr9p9m0>

### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

### **ISSN**

1069-7977

### **Authors**

Meunier, Fanny

Seigneuric, Alix

Spinelli, Elsa

### **Publication Date**

2005

Peer reviewed

# The Source of Beliefs in Conflicting and Non Conflicting Situations

**Hugo Mercier (hmercier@isc.cnrs.fr)**

Institut Jean Nicod, 1bis, Avenue Lowendal  
Paris, 75007 France and  
Institut des Sciences Cognitives, 67, boulevard Pinel  
Bron, 69675 France

**Jean-Baptiste Van der Henst (vanderhenst@isc.cnrs.fr)**

Institut des Sciences Cognitives, 67, boulevard Pinel  
Bron, 69675 France

## Abstract

What stance does one take towards information received as part of communication? On the one hand, a great deal of knowledge one acquires in one's life comes from communication. One is therefore likely to be trustful towards information communicated by others. On the other hand, communication can be used to manipulate others. If an audience were to constantly trust information provided by a communicator, she would often be misled. Hence, being cautious with communicated information can reduce the risk of acquiring false beliefs. From an evolutionary standpoint, people should process a new piece of information distinctively as a function of its source. A particularly interesting case arises when a communicated piece of information conflicts with beliefs that have been acquired on one's own. Are we inclined to revise our beliefs or are we inclined to conserve our beliefs and reject the speaker's information? A first experiment showed that participants are biased towards the beliefs they acquired on their own: People are less likely to revise an initial set of beliefs obtained on their own when it is contradicted by communicated information, than to revise an initial set of beliefs obtained by communication when it is contradicted by information obtained on their own. A second experiment showed that the more likely the communicator is in position of manipulating her audience the less likely are participants to revise their beliefs on the basis of what he has communicated. Finally, a third experiment showed that even when beliefs do not have to be revised, the source has some influence on the degree to which these beliefs are endorsed.

**Keywords:** Belief revision; Communication; Machiavellian intelligence.

## Introduction

How did we learn that the number 'pi' equals 3.14...? How do we know the price of the barrel of oil? How do we know that Britney Spears married a childhood friend? To obtain this knowledge, there is no need to calculate the ratio between the perimeter of a circle and its diameter, to buy a thousand barrels of oil from the *Kuwait Petroleum Corporation* or to have been present at the *Little White Wedding Chapel* of Las Vegas on the third of January 2004. All those pieces of information have been transmitted by others: by schoolteachers, financial journalists or our neighbour,

respectively. Clearly, a good deal of our knowledge has a social origin.

What does this imply for the way our cognitive processes operate? On the one hand, they must be able to assimilate this socially transmitted knowledge and to give it some credit. The ability to communicate is one of the cognitive tools playing this role. On the other hand, those processes may have to face the danger of manipulation. If we can reasonably consider that individually acquired knowledge does not result from processes susceptible to deceive us (at least in our natural environment, see (Sperber, 2001)), things become different for knowledge obtained by communication. Our interlocutor may want us to believe something so that we will behave in a certain way. As long as what is the most important for her is not the piece of information she communicates to us but the way our behaviour will be altered if we believe in this piece of information, she might choose to communicate something false on purpose. Thus, communication brings not only the risk that our interlocutors may be mistaken, but also the danger that they may want to manipulate us (Dawkins & Krebs, 1978; Sperber, 2001).

The dangers of manipulation have been underlined for some years by evolutionary minded researchers. The "Machiavellian Intelligence Hypothesis" proposes that "high level" cognitive abilities have evolved by natural selection in order to deal with the complexity of the social environment (Byrne & Whiten, 1988; Whiten & Byrne, 1997). Among these abilities, some may be dedicated to manipulation: either to try it on others or to ward off manipulation attempts (Cosmides & Tooby, 1992; Sperber, 2001). More specifically, since language is a major tool of potential manipulation, we might be endowed with mechanisms designed to deal with the special kind of manipulation allowed by language. One way to defend ourselves against liars or would-be manipulators may be to pay attention to some of their external features: Are they sweating? Are they shaking? Do they seem nervous? Has their pitch changed? Do they look at you in the eyes? However, we seem to perform badly when asked to detect cheaters using this kind of clue (Ekman, 2001). Another way to detect liars can be to check the internal and external consistency of what other people say (Sperber, 2000). Every liar knows how hard it is to remain consistent as the lie expends. The possibility that this

consistency checking may be the source of some of our logical abilities remains to be empirically tested.

Still another solution is to “decouple” the content of what we are told from our other beliefs (Cosmides & Tooby, 2000; Sperber, 1997). This way we can keep some beliefs “quarantined” in order to test them further before accepting them. It could allow us to check if the source is reliable or if there is some fact that contradicts the quarantined belief. For example, you notice that the box of chocolates is empty. Seeing your disappointed figure, Julie tells you that Paul finished them all last evening. At least for some time, the belief that Paul has finished the chocolates will be embedded that way:

- (1) Julie has said that
- (2) Paul finished the chocolates

Before incorporating (2) with your other fixed beliefs, you may want to confront it with some other information: perhaps Julie dislikes Paul and is quick to gossip, perhaps there is another witness, perhaps you know that Paul does not like chocolates since he refuses all the chocolates you offer him, etc. Such a mechanism allows us to keep a safe distance towards transmitted beliefs. However, at one point we will have to make a choice: either we incorporate the communicated belief or not. If it has passed all the tests, we can go on and add it to our other fixed beliefs, but what of the cases in which it contradicts one of our previously held beliefs? In these cases where beliefs conflict, how do we take their source into account? This is the situation we will explore in this paper.

When confronted with two conflicting beliefs, we have to revise one of them. The mechanism used to do so is called “belief revision”, and it is notoriously hard to understand (e.g. Nichols & Stich, 2003, p. 30). It has been explored by researchers in artificial intelligence, who came up with such concepts as “epistemic entrenchment” (the fact that different kinds of propositions are more or less easily revised) or “minimal change” (we should do the smallest possible change to accommodate the new belief). Some experimental work tried to see if these concepts applied to humans as well. The protocol usually employed is akin to the following:

At time  $t$  you know that

*If P then Q*  
*P*

At time  $t+1$  you learn (and you are sure) that

*Non-Q*

Which belief do you revise?

And then different choices are offered to the participant: revising one premise, or the other, or both, or be undecided about one, etc. This kind of framework allows researchers to test the principles put forward in the field of AI, and the results also bear on some discussions in the psychology of reasoning (e.g. whether the mental model theory or the mental logic theory could best account for the findings) (Elio, 1997, 1998; Elio & Pelletier, 1997; Johnson-Laird, Girotto, & Legrenzi, 2004). However, in these experiments, the source of the belief is kept fixed: you are supposed to have come by the beliefs by yourself (as in Dieussaert, Schaeken, De Neys,

& d'Ydewalle, 2000). In the experiments of Politzer & Carles, 2001 some informant communicates the beliefs to you, but the source is also uniform since the informant communicates all the beliefs. Thus, earlier experiments do not tell us how the source of beliefs affects the degree to which these beliefs are entrenched. The two first experiments proposed here are designed to test the role of the source in a belief revision situation, with the ‘manipulation’ factor being studied in the second one. The third experiment will come back to simpler situations to see if there are some differences in the way we consider communicated and individually acquired beliefs when no conflict arises. Guided by evolutionary considerations, we hypothesise that individuals will give some importance to the source of beliefs, and will tend to favour individually acquired beliefs over communicated ones. Moreover, this tendency should be stronger if the source of this communicated belief is perceived as being potentially manipulative.

## Experiment 1

### Method

In the first experiment, participants were provided with short texts describing a situation in which two pieces of information coming from two different sources conflicted with each other. In one condition, the participant was first placed in a position to adopt a perspective in which she believes a piece of information (acquired on the basis of observations she has made very often). Then, the participant learns that an individual holds a view that contradicts her own position. This contradictory view is based on several observations (the ‘self-other’ condition). The participant is then asked whether she still gives credence to her initial belief. The second condition provided the same framework but with a switched perspective (the ‘other-self’ condition): i.e., the participant is informed that an individual has a belief supported by observations made very often. However, in this case, it is the participant who contradicts this belief, getting support from several observations. The participant is then asked whether she gives credence to the initial belief acquired by the other person. Here is an example of description participants received for each condition:

Self-Other condition (all the examples are translated from the French original version)

*You are a salesperson working in England. You work in the Peak District area and you have to make the trip from Leeds to Birmingham very often.*

*In a roadside café, you meet another salesperson who has just been transferred to the same area and who will have to make the trip from Leeds to Birmingham. He is asking you whether the fastest route is the one that passes through Walfham or the one that passes through Thetford. Here is what you answer: “The fastest route is the one that passes through Walfham”.*

*A few days later, you meet him again, and he tells you: “I have made the trip several times from Leeds to Birmingham, and the fastest route is the one that passes through Thetford”.*

*In the end, do you think that the fastest road is the one that passes through:*

Walfham                      Thetford

**Self-Other condition**

*You are a salesperson working in England. You have just been transferred to the Peak District area and you will have to make the trip from Leeds to Birmingham.*

*In a roadside café, you meet another salesperson who works in the same area and who very often makes the trip from Leeds to Birmingham. You are asking him whether the fastest route is the one that passes through Walfham or the one that passes through Thetford. Here is what he answers: "The fastest route is the one that passes through Walfham".*

*A few days later, you meet him again, and you tell him: "I have made the trip several times from Leeds to Birmingham, and the fastest route is the one that passes through Thetford".*

*In the end, do you think that the fastest road is the one that passes through:*

Walfham                      Thetford

We thus manipulate the source of both the initial and the conflicting beliefs (oneself vs. other) and we introduce an asymmetry regarding the number of observations supporting those beliefs (observations made very often for the initial belief vs. several observations for the conflicting belief). However, the difference between "very often" and "several" may not have been perceived as crucial by the participants, so we will focus on the role of the source. Our prediction is that more participants should stick with the initial belief in the 'Self-Other' condition than in the 'Other-Self' condition, since they acquired the initial belief individually in the former condition whereas it was communicated in the latter. Moreover, we predict that, overall, participants should favour individually acquired information over information coming from the peer.

Ninety-six French native speaking students from the University of Lyon II participated in this experiment. Forty-eight participants received a single description of the 'self-other' condition and 48 participants received a single description of the 'other-self' condition. Four different scenarios were used and involved a salesperson, a fisherman, a cook, or a motorcycle enthusiast. Participants were individually questioned in the lobby of the university. They were given written instructions with one description and they were requested to read the description and the question carefully.

**Results and discussion**

The four scenarios prompted comparable results and were thus combined in Table 1 below. More participants endorsed the initial assertion in the 'Self-Other' condition than in the "Other-Self" condition ( $\chi^2= 18.77, p<.001$ , see Table 1). Moreover, more participants gave credence to the belief individually acquired than to the communicated belief ( $\chi^2 = 18.38, p<.001$ ).

N=96	"Self-Other" condition	"Other-Self" condition
Initial belief	65%	21%
Conflicting belief	35%	79%

Table 1: Percentages of participants endorsing the initial and the conflicting beliefs.

Experiment 1 shows that the epistemic values attached to transmitted knowledge and to individually acquired knowledge differ. Participants favour information they obtained individually over information coming from others, at least when a conflict between these two pieces of information arises. This is even true in the face of objective evidence that runs against the choice of the participants (the number of observations should have favoured the communicated belief in the "other-self" condition). Participants are less likely to revise an initial belief acquired individually and challenged by a peer than when it is transmitted by a peer and challenged by information acquired individually. This may be in line with the hypothesis that cognitive mechanisms have been designed in order to avoid manipulation attempts arising with communication. But of course the extent to which those mechanisms weaken the epistemic value of socially transmitted information needs to be investigated. A speaker may indeed have the goal to provide his audience with false beliefs but he may also be genuinely sincere. Our cognitive mechanisms might therefore be tailored to the degree of manipulation conveyed by the speaker: The more an individual would be perceived as manipulative the more what he communicates would be likely to be rejected. In the next experiment we aimed at manipulating the degree of perceived manipulation.

**Experiment 2**

**Method**

Experiment 2 contained three conditions that differed according to the source of the initial belief. In the first condition, participants had to adopt a perspective in which they acquired a belief based on individual observation, which is then challenged by another individual observation. This is the 'oneself' condition. In the other two conditions, participants are presented with a piece of information communicated either by a) an individual – a neighbour – who does not have any obvious reason to manipulate his audience or b) an individual – a salesman – who does have obvious reasons to manipulate his audience (a salesman is typically seen as someone who provides customers with information in order to convince them to purchase a product). These are the 'neighbour' and the 'seller' conditions. Participants are subsequently presented with an individually acquired piece of information that challenges the belief based on what the neighbour or the seller communicated. In all three conditions, participants are then required to assess the degree of credence they give to the initial belief. We predict that participants should be less likely to revise the initial belief when it has been acquired individually than when it has been communicated (by a neighbour or a salesman). We also predict that they should be less likely to revise the initial belief when it has been communicated by a neighbour than

when it has been communicated by a salesman. Here is an example of the stimuli we used in this experiment (the three conditions only differ with respect to the first paragraph):

**‘Oneself’ condition:**

*A few days ago, you got a stain on your favourite pullover. In the past, you have used the “Fenrir” washing powder. You tell yourself: “I have often used it and it is effective on almost all stains”. You therefore decide to go and buy a drum of “Fenrir”.*

*Once you are back home, you try the washing powder on the stain of your pullover. Unfortunately, no matter how much you scrub the stain, you do not manage to get rid of it.*

*According to you, the “Fenrir” washing powder is:*

- Effective on all stains*
- Effective on almost all stains*
- Effective on many stains*
- Effective on few stains*
- Effective on very few stains*

**‘Neighbour’ condition (only the first paragraph):**

*You are talking with your neighbour. A few days ago, you got a stain on your favourite pullover. You talk to him about it. He then mentions the “Fenrir” washing powder. He tells you: “I have often used it and it is effective on almost all stains”. You therefore decide to go and buy a drum of “Fenrir”.*

**‘Seller’ condition (only the first paragraph):**

*You are doing the shopping at the supermarket. A few days ago, you got a stain on your favourite pullover. Among the salesmen, you hear one who sings praises over a washing powder: “Fenrir”. He tells you: “I have often used it and it is effective on almost all stains”. You therefore decide to go and buy a drum of “Fenrir”.*

A group of 237 undergraduate History students from the University of Lyon II participated in this experiment. Each participant received a single problem from one of the three conditions. Two different scenarios were used: a ‘washing powder’ scenario and an ‘insecticide’ scenario.

**Results and discussion**

The degree of belief revision was assessed according to the answer selected by the participant. The “all” answer (i.e. the washing powder was effective on *all* stains) means that the initial belief was reinforced (which is highly improbable here), the “almost all” answer indicates the absence of revision, the “many” answer indicates a weak revision, the “few” answer indicates a serious revision and the “very few” answer indicates a very strong revision. Not surprisingly, none of the participants choose the “all” answer. As shown in Figure 1, participants were more conservative in the 'oneself' condition than the in 'neighbour' and the 'salesman' conditions: More participants choose the “almost all” and “many” answers in the 'oneself' condition than in the 'neighbour' (36% vs. 21% for “all”,  $\chi^2 = 4.77, p < .03$ ; 21% vs.

7% for “many”,  $\chi^2 = 6.44, p < .02$ ) and 'salesman' conditions (36% vs. 13%, for “all”,  $\chi^2 = 9.92, p < .002$ ; 21% vs. 4% for “many”,  $\chi^2 = 9.68, p < .002$ ). In addition, more participants selected the “few” answer in the 'neighbour' condition than in the 'oneself' condition (58% vs. 32%,  $\chi^2 = 10.96, p < .001$ ) and more participants selected the “very few” answer in the 'salesman' condition than in the 'oneself' condition (36% vs. 11%,  $\chi^2 = 13.13, p < .0005$ ). Finally, more participants choose the 'very few' answer in the 'salesman' condition than in the 'neighbour' condition (36% vs. 14%,  $\chi^2 = 9.44, p < .003$ ). In brief, as for the previous experiment participants exhibited a bias towards the beliefs they acquired individually. They were also more inclined to revise beliefs received from an individual who is likely to be manipulative than from an individual who is likely to be neutral.




Figure 1: Percentages of “almost all”, “many”, “few” and “very few” answers in the three conditions of experiment 2.

Thus when we have to decide which of two conflicting beliefs we should revise, we take into account the source of these beliefs: We tend to favour one that was individually acquired over one that was communicated. However, it remains to be seen whether this apparent mistrust toward communicated information arises specifically in the context of belief revision or is a more general phenomenon.

**Experiment 3**

**Method**

In this third experiment participants were confronted with a belief, either individually acquired (‘oneself’ condition) or communicated by a neighbour or a salesman (‘neighbour’ and ‘seller’ condition). However, they were not given any contradictory belief; they just had to make a simple evaluation of the initial belief, as in the example that follows:

**‘Oneself’ condition:**

*A few days ago, you got a stain on your favourite pullover. In the past, you have used the “Fenrir” washing powder. You tell yourself: “I have often used it and it is effective on all stains”.*

*According to you, the “Fenrir” washing powder is:  
Effective on 100% of the stains*

Effective on 95% of the stains  
 Effective on 90% of the stains  
 Effective on 85% of the stains  
 Effective on 80% of the stains  
 Effective on 75% of the stains

**‘Neighbour’ condition (only the first paragraph):**

*You are talking with your neighbour. A few days ago, you got a stain on your favourite pullover. You talk to him about it. He then mentions the “Fenrir” washing powder. He tells you: “I have often used it and it is effective on all stains”.*

**‘Seller’ condition (only the first paragraph):**

*You are doing the shopping at the supermarket. A few days ago, you got a stain on your favourite pullover. Among the salesmen, you hear one who sings praises over a washing powder: “Fenrir”. He tells you: “I have often used it and it is effective on all stains”.*

These scenarios allowed us to test the credence given to the belief depending on its source, when no contradictory information was involved. A group of 149 undergraduate History students from the University of Lyon II participated in this experiment. Each participant received a single problem from one of the three conditions. As in experiment 2, two different scenarios were used: a ‘washing powder’ scenario and an ‘insecticide’ scenario.

**Results and discussion**

The degree of credence given to the belief was assessed according to the answer selected by the participant. The “100%” answer meant that the word “all” was given its logical meaning and that the belief was given maximum credence. However, pragmatic considerations lead us to expect lesser percentages: We often use the word “all” when we have less than a total fit in mind. But these pragmatic considerations should be the same in the three conditions: They depend on the type of items being quantified by “all” (e.g. “all” does not mean the same in “all the French like frog legs” and “all squares have four equal angles”). As the type of item quantified is the same in all three conditions, the pragmatic effects should also be the same. Thus, any difference in the distribution of answers between conditions will be due to the source of the information. It is possible to quantify the credence given to the source by observing the deviation from the “100%” answer.

The results are shown in figure 2. To analyse the results, the range of answer was cut in half, with the “lower part” being 75, 80 and 85%, and the “higher part” 90, 95 and 100%. With this distinction, the ‘oneself’ and the ‘neighbour’ condition were very close (66% vs. 60% for the “higher part”,  $\chi^2 = 0.343, p > .5$ ). By contrast, the ‘seller’ condition was significantly different both from the ‘oneself’ condition (40% vs. 66% for the “higher part”,  $\chi^2 = 7.09, p < .01$ ) and from the ‘neighbour’ condition (40% vs. 60% for the “higher part”,  $\chi^2 = 4.17, p < .05$ ). Thus in this experiment too the source of the information played a role, but the distinction was not between oneself and others: it was between a source that was overtly manipulative (the salesman) and source that was not (oneself or the neighbour)

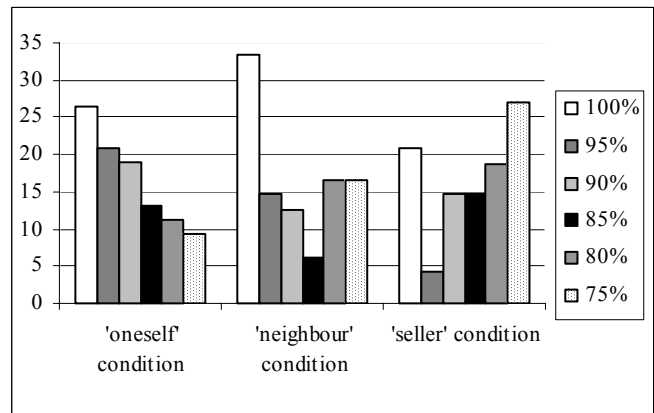


Figure 2: Percentages of the different answers in the three conditions of experiment 3.

This last experiment points to an interaction between the role of the belief revision context and the source of the belief. By itself, the source has an influence: Even when it does not contradict any of our beliefs, a communicated belief is given more or less credence depending on its source: less credence is given when the source is potentially manipulative. Taken alone, the results of this experiment might have indicated that we make at most a very slight difference between a belief individually acquired and a belief transmitted by a non-manipulative source. However, this conclusion would have been premature. As shown in the two first experiments, when we obtain a piece of information that contradicts one of our beliefs, we tend to grant it less status when it has been communicated (even by a non-manipulative source) than when it has been individually acquired.

**Conclusion**

Firstly, a point regarding the methodology. Given that our aim was to investigate a phenomenon that is at least partly social, it could be argued that our paper and pencil tasks are too artificial. They were designed in order to be in line with previous research on belief revision, and they had the advantage of allowing us to control a number of factors that are known to play a role in attitude change (see, e.g., Petty & Cacioppo, 1981). However, we plan to perform another set of experiments in a more naturalistic setting. Our prediction is that we would obtain qualitatively similar results than in the experiments which are reported here.

The theoretical starting point of this study was inspired by an evolutionary hypothesis. However, it would be quite premature to tell that our findings show that we are endowed with a mechanism specifically adapted to take into account the source of beliefs. It is possible that the bias observed is purely learnt: most people have been deceived or lied to in the past, and this may explain why they tend to favour individually acquired beliefs over communicated ones. To test further the evolutionary hypothesis, at least two roads could be followed: one is to investigate the ontogeny of this mechanism, and the other is to ascertain it is not a cultural particularity of the population under study. The first strategy

has been successfully used in a number of domains, from naïve physics to theory of mind, to show that we are innately endowed with special abilities. We do not know of any experiment with young children that would replicate the findings obtained here, but other related results may give us some clues as to what the result of such a replication would be. Experiments with children show that they are far from being wholly naïve: when given a piece of information that conflicts with their beliefs, they can take into account the informedness of the communicator and its previous truthfulness to decide whether they give up their own belief or if they stick with it (Clément, Koenig, & Harris, 2004; Koenig, Clément, & Harris, 2004; Robinson, Champion, & Mitchell, 1999). So it seems that we are able to take into account the source of beliefs and some of its relevant characteristics from very early on, a fact that can point to an innate basis for this ability.

If the mechanism under study is really an adaptation, it should not be dependent upon a specific cultural context: we should find it in other cultures, if not in all. However, it is possible that the effects observed in the present experiments result from cultural features of the Western world. This hypothesis finds some support in the work of Richard Nisbett and his colleagues. They used a large array of cognitive tasks to test for differences between Eastern and Western populations. Of particular relevance here is the finding that Easterners presented with conflicting beliefs are inclined to reconcile them instead of frankly favouring one of the beliefs, as Westerners tend to do (Nisbett, Peng, Choi, & Norenzayan, 2001). So the contrast between individually acquired and communicated belief may be less stark in an Eastern cultural setting than in the Western population studied here. In order to disentangle these views cross-cultural studies ought to be carried out in the future.

## References

- Byrne, R. W., & Whiten, A. (Eds.). (1988). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. New York: Oxford University Press.
- Clément, F., Koenig, M. A., & Harris, P. (2004). The ontogeny of trust. *Mind and Language*, 19(4), 360-379.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. H. Barkow, L. Cosmides & J. Tooby (Eds.), *The Adapted Mind* (pp. 19-136). Oxford: Oxford University Press.
- Cosmides, L., & Tooby, J. (2000). Consider the source: The evolution of adaptations for decoupling and metarepresentations. In D. Sperber (Ed.), *Metarepresentations: A Multidisciplinary Perspective*. Oxford: Oxford University Press.
- Dawkins, R., & Krebs, J. R. (1978). Animal signals: Information or manipulation? In J. R. Krebs & N. B. Davies (Eds.), *Behavioural Ecology* (pp. 282-309). Oxford: Basil Blackwell Scientific Publications.
- Dieussaert, K., Schaeken, W., De Neys, W., & d'Ydewalle, G. (2000). Initial belief as a predictor of belief revision. *Current Psychology of Cognition*, 19(3), 277-288.
- Ekman, P. (2001). *Telling Lies*. New York: Norton.
- Elio, R. (1997). What to believe when inferences are contradicted. The impact of knowledge type and inference rule. In *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 211-216). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Elio, R. (1998). How to disbelieve p-->q: Resolving contradictions. In *Proceedings of the Twentieth Meeting of the Cognitive Science Society* (pp. 315-320). Mahwah, NJ: Lawrence Erlbaum Associates.
- Elio, R., & Pelletier, F. J. (1997). Belief change as propositional update. *Cognitive Science*, 21(4), 419-460.
- Johnson-Laird, P. N., Girotto, V., & Legrenzi, P. (2004). Reasoning from inconsistency to consistency. *Psychological Review*, 111(3), 640-661.
- Koenig, M. A., Clément, F., & Harris, P. (2004). Trust in testimony: Children's use of true and false statements. *Psychological Science*, 15(10), 694-698.
- Nichols, S., & Stich, S. (2003). *Mindreading*. Oxford: Oxford University Press.
- Nisbett, R. E., Peng, K., Choi, I., & Norenzayan, A. (2001). Culture and Systems of Thought: Holistic vs. Analytic Cognition. *Psychological Review*, 108, 291-310.
- Petty, R. E., & Cacioppo, J. T. (1981). *Attitudes and Persuasion: Classic and Contemporary Approaches*. Dubuck: WCB.
- Politzer, G., & Carles, L. (2001). Belief revision and uncertain reasoning. *Thinking and Reasoning*, 7(3), 217-234.
- Robinson, E. J., Champion, H., & Mitchell, P. (1999). Children's ability to infer utterance veracity from speaker informedness. *Developmental Psychology*, 35(2), 535-546.
- Sperber, D. (1997). Intuitive and reflexive beliefs. *Mind and Language*, 12(1), 67-83.
- Sperber, D. (2000). Metarepresentations in an evolutionary perspective. In D. Sperber (Ed.), *Metarepresentations: A Multidisciplinary Perspective* (pp. 117-137). Oxford: Oxford University Press.
- Sperber, D. (2001). An evolutionary perspective on testimony and argumentation. *Philosophical Topics*, 29, 401-413.
- Whiten, A., & Byrne, R. W. (Eds.). (1997). *Machiavellian Intelligence II: Extensions and Evaluations*. Cambridge: Cambridge University Press.