UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Subjective Confidence and the Belief Bias Effect in Syllogistic Reasoning

Permalink
https://escholarship.org/uc/item/84g7z1pj

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors
Quayle, Jeremy D.
Ball, Linden J.

Publication Date
1997

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

Subjective Confidence and the Belief Bias Effect in Syllogistic Reasoning

Jeremy D. Quayle and **Linden J. Ball**
Cognitive and Behavioral Sciences Research Group,
Division of Psychology, University of Derby,
Mickleover, Derby, DE3 5GX, UK.
j.d.quayle@derby.ac.uk

Abstract

An experiment is reported in which participants were asked to record how confident they felt about the correctness of their responses as they assessed the validity of deductive arguments whose conclusions varied in prior believability. The results showed that participants were more confident of their responses to valid problems than invalid problems irrespective of believability status, providing support for the idea that invalid problems are more demanding to process than valid problems. Effects of belief, logic on conclusion acceptance rates and a logic×belief interaction are also demonstrated, and evidence is provided to suggest that belief bias principally reflects a tendency to reject unbelievable arguments. A theory is proposed in which belief bias effects are accounted for by the variations in the processing demands of valid and invalid syllogisms.

General Introduction

First devised by Aristotle, the syllogism is one of several reasoning problems used by psychologists when investigating deductive reasoning performance. Although it was once believed that syllogistic reasoning formed the basis of all rational thought (see Evans, Newstead & Byrne, 1993), real-world examples of reasoning problems whose underlying structures resemble the syllogism are hard to find. Instead, the syllogism is perhaps best characterized as a useful tool for use in the experimental study of deduction. The following is an example of the type of syllogism typically used in reasoning experiments:

No politicians are drinkers
All drinkers are judges
Therefore,
Some judges are not politicians

As can be seen, a syllogism has two premises and a conclusion. The premises feature three terms (politicians, drinkers, and judges in the above example). There are two 'end terms' (judges and politicians) and a 'middle term' (drinkers). The items featured within a syllogism can be letters of the alphabet (*abstract* content) or any words that refer to sets of individuals or objects (*realistic* or *thematic* content). Participants in syllogistic reasoning experiments are required either: (1) to produce their own valid conclusions (ones that link together and describe the logical relationship of the two end terms to one another in a way that is *necessarily* true given that the premises are true); or else (2) evaluate the validity of presented conclusions.

Traditionally, four different quantifiers are featured within the premises and conclusions of English-language syllogisms. These are: *all*, *no* (or *none*), *some*, and *some...are not*. Conventionally these quantifiers are referred to by letters of the alphabet: *all* = A, *no* = E, *some* = I, *some...are not* = O. It should be noted that a logician's definition of the quantifier *some* is: 'At least one and possibly all'.

The main aim of research into syllogistic reasoning has been to identify the factors that contribute to the variations in difficulty between different problem types. It has been found that logical errors made by participants carrying out syllogistic reasoning tasks are not random (see Evans et al. 1993 for a recent review). Instead, the difficulty of abstract syllogisms has been found to vary according to two main factors: (1) the arrangement of the three terms within the two premises - this factor is termed *figure*; and (2) the different combinations of quantifiers that may be featured within syllogisms - this factor is termed *mood*.

In addition to the effects of mood and figure, the effect of thematic content upon syllogistic inferences has been a key focus of interest. Studies have revealed that people are affected by beliefs and prior knowledge when they judge whether a putative conclusion follows from presented premises (see, for example, Garnham & Oakhill, 1994, for a review). There are in fact three basic findings that derive from studies of 'belief bias' which manipulate the validity of presented arguments alongside the believability of given conclusions (cf. Evans et al., 1993). First, believable conclusions are more readily accepted than unbelievable ones. Second, logically valid conclusions are more readily accepted than invalid ones. Third, there is an interaction between logical validity and believability such that the effects of believability are stronger on syllogisms leading to invalid conclusions than on those leading to valid conclusions. In essence, people seem greatly more inclined to reject invalid-unbelievable arguments than invalid-believable ones. A good illustration of this interaction and the basic belief bias effect is shown in Table 1 from data produced by Evans, Barston and Pollard (1983).

The findings of belief bias research suggest that in real-world situations we may not always draw conclusions or base our decisions upon valid logical arguments, but instead

Table 1: Overall Mean Percentage Frequency of Acceptance of Conclusions as a Function of Logical Validity and Believability in the Study of Evans et al. (1983) (taken from Evans et al 1993, p.245)

	Believable	Unbelievable	Mean
Valid	89	56	72
Invalid	71	10	40
Mean	80	33	

upon the *a priori* believability of the information available to us. It is not difficult to comprehend how such a behavior might have developed in humans for reasons of cognitive economy, and how it might be an effective strategy most of the time (see Evans & Over, 1996, for relevant arguments). However, the possibility that this behavior may result in the drawing of erroneous and seemingly irrational conclusions raises a number of interesting and important questions which have formed the focus of research into belief bias. For example, under what conditions do we reason logically and apparently ignore believability, and under what conditions do we not? Also, does the believability of a putative conclusion determine whether we test a logical argument, or does success in evaluating a logical argument determine whether we respond in accordance with belief?

This paper will proceed with a brief review and evaluation of three current theoretical explanations for belief bias effects. From this review it will be concluded that the account offered by the mental models theory (e.g., Johnson-Laird & Byrne, 1991) has the greatest amount of empirical support, although this explanation falls down with the addition of *ad hoc* assumptions that have been proposed to explain unpredicted findings. It will be argued that there is scope for alternative mental models accounts of belief bias which explicitly consider the notion of working memory capacity limitations. Although central to the mental models explanation of deductive competence, this notion is curiously overlooked in the existing mental models account of belief bias.

Next, an experiment will be reported in which subjective confidence in the correctness of syllogistic inferences was measured in order to test the idea of a processing distinction between valid and invalid arguments. The data presented show how participants were more confident of the correctness of their responses to valid problems than invalid problems, thus providing support for this distinction. An example of the three main findings of belief bias research is also provided, as well as evidence to suggest that belief bias principally reflects a tendency to reject unbelievable arguments rather than a tendency to accept believable ones (cf. Evans & Pollard, 1990). Based upon the suggested valid-invalid processing distinction, plus the empirical support gained for this notion in the reported experiment, a new theory of belief bias will then be proposed.

Three Theories of belief bias

Over the past decade, three main theories have been proposed in order to account for the belief bias effect and the interaction between belief and logical validity: (1) the *selective scrutiny model* (Barston, 1986; Evans, 1989); (2) the *misinterpreted necessity model* (Evans et al., 1983; Evans, 1989); and (3) the *mental models* account (e.g., Oakhill & Johnson-Laird, 1985; Oakhill & Garnham, 1993).

Each of the three accounts places the locus of the belief bias effect at a different point in the reasoning process. The selective scrutiny model places belief bias at a point before any logical reasoning takes place. According to this model, reasoners accept believable conclusions and only go on to reason logically with unbelievable ones. The misinterpreted necessity model is based on the assumption that reasoners have a fragile understanding of the principle of *logical necessity*, and places the locus of the belief bias effect at a point after an attempt at logical reasoning has failed. In this model, people attempt to reason, but when a conclusion is indeterminately invalid, i.e. the conclusion might *possibly* be true but is not *necessarily* true, they base their response on the believability of the conclusion. The mental models account places the locus of the belief bias effect immediately after an initial mental model of a syllogism's premises has been constructed. At this point, reasoners are motivated to accept believable conclusions, and only go on to construct alternative (and potentially falsifying) models of the premises if a conclusion is unbelievable.

The three accounts of belief bias vary in terms of what aspects of the belief bias effect they successfully explain and what predictions they make. The selective scrutiny model explains the effect of belief and the logic×belief interaction, but cannot explain the effect of logical validity on believable conclusions. The misinterpreted necessity model readily explains the logic×belief interaction. This theory also explains how the interaction is exhibited with *indeterminately* invalid syllogisms, but disappears when *determinately* invalid conclusions are presented (Newstead, Pollard, Evans & Allen, 1992). However, this model does not explain how the effect of belief is exhibited with syllogisms other than those which lead to indeterminately invalid conclusions.

A number of systematic manipulations have lead to a good degree of support for the mental models account (cf., Evans & Over, 1996). Still, a number of unsatisfactory elements remain within this account (e.g., the notion of 'conclusion filtering' as an explanation of belief effects on one-model syllogisms). It is maintained that whilst the framework offered by the mental models theory is a valuable one, there is scope for alternative mental models accounts of belief bias.

A Processing Distinction Between Valid and Invalid Multiple-Model Syllogisms

The notion of limited working memory capacity and how this factor impacts upon our ability to construct adequate problem representations that show a putative conclusion to be true or untrue is central to the mental models theory of

syllogistic reasoning (e.g. see Johnson-Laird & Byrne, 1991). This theory states that people begin reasoning with syllogisms by constructing a mental model in which the minimum amount of information concerning the logical relationships between the terms in the premises is made explicit. Putative conclusions that are true in this model are tested against 'fleshed out' mental models that make more of this information explicit. If a falsifying model cannot be constructed, then the conclusion is valid, otherwise, the conclusion is invalid. Syllogisms that require the consideration of more than one mental model in order to test a conclusion ('multiple-model' syllogisms) place great demands on working memory resources. These syllogisms, therefore, are more difficult than syllogisms requiring a single model to be considered ('one-model' syllogisms)(see Johnson-Laird & Bara, 1984).

The mental models account of belief bias distinguishes between the way in which one-model syllogisms and multiple-model syllogisms are mentally processed, making different belief bias predictions for each. However, this theory does not consider the varying processing demands that different multiple-model syllogisms place upon working memory resources and how this factor might contribute to the belief bias effect and the widely reported logic×belief interaction. It is argued that a clear distinction exists between the processing demands of valid and invalid multiple-model syllogisms, and that this distinction may be the cause of belief bias effects and the logic×belief interaction. The argument hinges on the observation that with valid multiple-model syllogisms correct responses can be given after the construction of a single mental model (cf., Hardman & Payne, 1995). As the second term in a valid conclusion to a multiple-model syllogism is represented exhaustively in all accurate mental models, the relationship between the end-terms in the model that shows the conclusion to be true remains unchanged when the model is 'fleshed out' (see below).

Valid three-model syllogism:

Some A are B
No B are C
Therefore,
Some A are not C

```
a  [b]        a  [b]        a  [b]
a  [b]        a  [b]        a  [b]
      [c]     a     [c]     a     [c]
      [c]           [c]     a     [c]
  . . .         . . .         . . .
```

(The exhaustive representation of a term is signified in mental models notation by square brackets (see above and Johnson-Laird and Byrne, 1991, pp. 107-110)).

Hence, participants may be aware that the consideration of more than one model is unnecessary. With valid problems, therefore, participants may feel certain of the correctness of their responses. However, as the second term in an indeterminately invalid conclusion to a multiple-model syllogism is *not* represented exhaustively in any accurate mental model, it will typically be necessary to flesh out initial models before a response can be given with certainty (see below).

Invalid three-model syllogism:

No A are B
Some B are C
Therefore,
Some A are not C

```
[a]           [a]           [a]     c
[a]           [a]     c     [a]     c
  [b]  c        [b]  c        [b]  c
  [b]  c        [b]  c        [b]  c
  . . .         . .           . . .
```

Due to working memory constraints this process may be unsuccessful, and participants will be uncertain of whether to accept or to reject a conclusion.

The Experiment

It has been discussed how the differing processing demands of valid and invalid syllogisms might cause feelings of certainty to vary between these problem types. An experiment was carried out to test this idea by asking participants to record how confident they felt about the correctness of the responses they gave in a syllogistic conclusion evaluation task. Syllogisms were presented with thematic contents designed to produce believable, unbelievable, and belief-neutral conclusions. It was predicted that there would be an effect of logic upon confidence ratings such that participants would be more confident in their responses when evaluating valid conclusions than invalid conclusions. Also, based upon the findings of earlier studies of the belief bias effect (e.g., Evans et al., 1983; Newstead et al., 1992) it was predicted that there would be standard effects of belief and logic, and a logic×belief interaction such that the effect of belief would be greater on the evaluation of invalid arguments than valid ones.

It was difficult to predict whether making a response based upon the believability of a conclusion would affect feelings of confidence. It may be that rejecting an unbelievable conclusion or accepting a believable conclusion would somehow boost or reduce confidence levels. Alternatively, confidence levels may be unaffected by the application of a belief heuristic such that confidence levels would be low with invalid problems and high with valid problems irrespective of believability. The idea that feelings of certainty/confidence determine the extent to which participants will respond in accordance with belief does not allow a prediction to be made on this matter. However, the selective scrutiny model - in which it is assumed that no consideration is given to the premises of syllogisms leading to believable conclusions - clearly predicts that there will be no effect of logic upon confidence ratings given for syllogisms with believable conclusions.

Similarly, the mental models account also predicts that there will be no effect of logic upon confidence ratings given for believable syllogisms, as conclusions to these problems are simply accepted after the construction of an initial mental model of the premises. By recording confidence ratings with the believable and unbelievable problems as well as the belief-neutral problems it was possible to test these predictions.

At this point it would be useful to discuss briefly the nature and distinction between believable, unbelievable and belief-neutral statements of the type typically used as conclusions in studies of belief bias. Oakhill and Johnson-Laird used the terms "rich people" and "millionaires" in believable and unbelievable conclusions. The term "millionaires" is a sub-category of the term "rich people". This relationship is a necessary function of the definitions of these terms, i.e., people know that rich people exist that are not millionaires, and that millionaires are rich people. As a consequence of this knowledge, the statement "Some millionaires are not rich people" is unbelievable and the statement "Some rich people are not millionaires" is believable. In the present study, the statement "Some teachers are not poets" will be used as a belief-neutral conclusion. Although "teachers" *may* belong to the category of "poets", and "poets" *may* belong to the category of "teachers", these relationships are not necessary functions of the definitions of these words. Therefore, the statement "Some teachers are not poets" should be neither particularly believable nor particularly unbelievable.

Evans and Pollard (1990) also used believable, unbelievable and belief-neutral materials. In the first of two experiments, the thematic materials used were *subjects of study* and *leisure interests*, e.g., "One of the leisure activities of mechanical engineering students is car maintenance" (believable), "One of the leisure activities of Social Work students is home electronics" (unbelievable), "One of the leisure activities of physical science students is sailing" (neutral). As can be seen, these statements concern student stereotypes and are, therefore, empirically rather than definitionally true or false. In the second experiment definitionally as well as empirically true and false statements were used, e.g. "All crocodiles are reptiles" (believable), "All cats are birds" (unbelievable), "All horses are traxites" (the neutral statements used invented nonsense categories). In both experiments participants were required to rate conclusions on a five point scale ranging from "Definitely False" through to "Definitely True" Evans and Pollard found significant evidence only for a tendency to reject unbelievable conclusions, i.e., unbelievable conclusions were given lower 'truth' ratings than believable and neutral conclusions, whilst the difference in ratings between believable and neutral statements was not significant. It was concluded from this finding that belief bias reflects a tendency to reject unbelievable conclusions rather than a tendency to accept believable conclusions, i.e., belief bias is a predominantly *negative* effect. And so, in addition to testing the predictions detailed above, the use of equivalent materials provided an opportunity to replicate Evans and Pollard's findings.

Method

Participants. Sixty participants took part in the experiment. The participants were all undergraduate psychology students at the University of Derby. None of the participants had taken formal instruction in logic. The participants were tested together in one group.

Materials. Two forms of three-model syllogism were used. Both syllogisms had the same B-A, C-B figure and both yield valid conclusions. The same "Some C are not A" form of conclusion was presented with both syllogisms. For one syllogism (in the EI mood) this conclusion is logically valid, but for the other syllogism (in the IE mood) this conclusion is indeterminately invalid.

A set of potential conclusions which were false by definition, (e.g. "Some kings are not men") were chosen, together with a set of believable conclusions, e.g., "Some animals are not cats". The conclusions were devised so as to appear believable when the terms were presented in one order, but unbelievable when the terms were reversed. In order to assess believability, the potential conclusions were rated by a group of 30 participants on a seven point scale ranging from "totally unbelievable" (-3) to "totally believable" (+3). Those conclusions which received the most extreme and consistent ratings were used in this study. A set of potential conclusions that were neither particularly believable nor unbelievable were also chosen, (e.g., "Some artists or not beekeepers").

The valid and invalid syllogisms were presented with three different sets of thematic contents leading to conclusions that were: (1) believable; (2) unbelievable by definition; and (3) belief-neutral. In addition to these six three-model syllogisms there were six filler problems (i.e., valid, one-model syllogisms featuring arbitrary thematic content). These were used to distract the participants from the form of the syllogisms of interest.

Design. A within participants design was used. All of the participants received the six three-model syllogisms together with the six one-model filler syllogisms (12 syllogisms in total). The syllogisms were presented one to a page, together with their potential conclusions, and stapled together in booklets. The order of the problem types was varied using a 6×6 balanced Latin square design; with the restriction that the filler items appeared in the same position in each booklet: in 1st, 3rd, 5th, 7th, 9th and 11th places. The thematic contents of the syllogisms were rotated over the different problem types, producing six different sets of materials. The six sets of materials were distributed evenly and randomly amongst the participants.

Procedure. The following instructions were presented on the second page of the test booklets and were also read aloud to each participant:

"This is an experiment to test people's reasoning ability. You will be given twelve problems. On each page, you will be shown two statements and you are asked if certain conclusions (given below the statements) may be logically deduced from them. You should answer this question on the

assumption that the two statements are, in fact, true. If, and only if, you judge that the conclusion necessarily follows from the statements, you should tick the "yes" box, otherwise the "no" box.

Beneath the yes/no response boxes there is a scale ranging from 1 (not at all confident) to 6 (very confident). Please indicate how confident you feel about each answer you give by circling a number on the scale.

Please take your time and be sure that you have the right answer before moving on to the next problem. You must not make notes or draw diagrams to help you in this task. Thank you very much for participating."

The participants were allowed as much time as they required to complete the book of syllogisms.

Results

Conclusion Acceptances. The mean percentage frequencies of participants accepting conclusions (i.e. deciding that a conclusion was logically valid) are presented in Table 2 for each type of syllogism. In order to test for negative and positive aspects of belief bias, acceptances for each of the three problem types were compared with each of the others in turn using sign tests. The difference in acceptances between believable and neutral conclusions was not significant. Acceptances of unbelievable conclusions were significantly lower than acceptances of believable conclusions ($p<.01$; one-tailed sign test) and neutral conclusions ($p<.001$; one-tailed sign test). There was a significant effect of logic ($p<.01$; one-tailed sign test) with participants accepting more valid conclusions than invalid conclusions. There was also a significant interaction between logic and belief such that the effects of belief (believable or unbelievable - excluding neutral conclusions) were stronger on syllogisms leading to invalid conclusions than on syllogisms leading to valid conclusions ($p<.05$).

Confidence Ratings. Participants were generally confident in the responses they gave in the evaluation task (see Table 3). The confidence ratings data were subjected to a multi-factorial analysis of variance. The factors were Logic (two levels) and Believability (three levels). The Logic factor was significant ($F = 3.38$; $p < .05$, 1-tailed) with participants giving higher confidence ratings to valid syllogisms than to invalid syllogisms. Despite a slight tendency to be more confident in responses to unbelievable and neutral problems than believable problems the Believability factor was not significant. The interaction between the Logic and Believability factors was not significant.

Table 2: Overall Mean Percentage Frequency of Acceptance of Conclusions as a Function of Logic and Believability Status.

	Believable	Neutral	Unbelievable	Mean
Valid	72	80	63	72
Invalid	65	58	32	52
Mean	68	69	48	

Table 3: Mean Confidence Ratings as a Function of Logic and Believability Status.

	Believable	Neutral	Unbelievable	Mean
Valid	4.15	4.27	4.32	4.24
Invalid	3.90	4.07	4.05	4.01
Mean	4.03	4.17	4.19	

Discussion

Conclusion Acceptances. The observed effect of logic provides further evidence for deductive competence in syllogistic reasoning. However, the generally high acceptance rates for both valid and invalid arguments support the idea that participants are often unable to flesh out initial mental models in which invalid conclusions are true, and so accept the result of an incomplete evaluation.

The three basic effects observed in previous studies of belief bias are also evident in the data: (1) an effect of logic; (2) an effect of belief; and (3) an interaction between logic and belief. The use of belief-neutral conclusions in addition to believable and unbelievable conclusions has allowed the investigation of whether belief bias is, in fact, a tendency to accept believable conclusions or to reject unbelievable conclusions. Acceptances of believable conclusions were not significantly different from acceptances of neutral conclusions, whilst significantly lower levels of acceptance were observed with unbelievable conclusions than with believable and neutral conclusions. Although it is acknowledged that caution should be applied when drawing conclusions from negative results, these observations suggest that belief bias reflects a tendency to reject unbelievable conclusions rather than a tendency to accept believable conclusions. This is consistent with the earlier findings of Evans and Pollard (1990) who argue that such a finding casts doubt upon the validity of the selective scrutiny model. In this model it is assumed that belief bias reflects a tendency to accept believable conclusions and to subject only unbelievable conclusion to tests of logical validity. They suggest that changing the first question in the model from "Is the conclusion believable?" to "Is the conclusion unbelievable?" can account for the equivalent treatment of believable and neutral conclusions. However, this alteration would mean that no effect of logic would be expected with syllogisms that lead to belief-neutral conclusions. Clearly, an effect of logic *is* observed with

such problems - both in the present experiment and in others.

Confidence Ratings. The confidence ratings data show that participants were more confident in their evaluations of valid conclusions than invalid conclusions. The relationship between confidence ratings and levels of belief bias observed provides strong support for the uncertainty/confidence hypothesis, i.e., low confidence levels and high belief bias were observed with invalid problems, whilst high confidence levels and low belief bias were observed with valid problems. Interestingly, the effect of logic upon confidence ratings was present irrespective of whether a conclusion was believable, unbelievable or neutral, suggesting that feelings of confidence in the correctness of a response are not strongly affected by the *a priori* believability of a conclusion. This finding cannot be explained by either the selective scrutiny model or the mental models account, both of which predict that no effect of validity upon confidence ratings will be found with syllogisms that have believable conclusions.

General Discussion

The idea behind the uncertainty hypothesis is not entirely a novel one. For example, Markovits and Nantel (1989) have suggested that participants may "use both a logical strategy and a nonlogical strategy in the same context, with the latter being used in order to accommodate a degree of internal conflict". However, no current theories of belief bias explicitly consider the possibility that participants make syllogistic inferences without feeling absolutely certain of their correctness. A new theory of belief bias based around the notion of uncertainty and upon the findings of the experiment presented in this paper will now be proposed.

An Uncertainty Theory of Belief Bias

Set within the general framework of the mental models theory of syllogistic reasoning, the uncertainty theory asserts that whether belief has an effect is determined by the ability to construct a mental model in which a conclusion is false or true without the possibility of there being alternative falsifying models. People are more successful in constructing such 'decisive' mental models with valid and determinately invalid problems than with indeterminately invalid ones (the representational effect that gives rise to this distinction has been discussed earlier in this paper). When successful, participants will be certain of the correctness of their logical evaluations, and ignore the believability status of the conclusion. However, if people are unable to construct such a mental model, then they cannot be certain of the correct response. In this situation they will fall back upon a belief heuristic in order to resolve the "internal conflict" that exists.

The uncertainty theory resembles the misinterpreted necessity model inasmuch as belief bias is said to have its effect upon responses after attempts at logical reasoning have taken place. However, the uncertainty theory suggests that people are often *unable* to determine the falsity of indeterminately invalid conclusions, and not that they *misinterpret* the notion of logical necessity - although the theory does not rule out this possibility. Also unlike the misinterpreted necessity model, the uncertainty theory shows how belief may bias responses with *all* types of syllogism, and not just those with indeterminately invalid conclusions. In this way the new theory overcomes the faults that also exist with the selective scrutiny model (which cannot account for the effect of logic on believable problems) and the mental models account (which, without *ad hoc* modifications, cannot account for the effect of belief on one-model syllogisms).

References

Barston, J.I. (1986), *An investigation into belief bias in reasoning.* Unpublished Ph.D. thesis, Plymouth Polytechnic.

Evans, J.St.B.T. (1989). *Bias in human reasoning: causes and consequences.* Hove: Lawrence Erlbaum Associates.

Evans, J.St.B.T., Barston, J.L. & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition*, 11, 295--306.

Evans, J.St.B.T. Newstead, S.E. & Byrne, R.M.J. (1993). *Human Reasoning: The Psychology of Deduction.* Hove: Lawrence Erlbaum Associates.

Evans, J.St.B.T., & Over, D.E. (1996). *Rationality and reasoning.* Hove: Psychology Press.

Evans, J.St.B.T., & Pollard, P. (1990). Belief bias and problem complexity in deductive reasoning. In J.-P. Caverni, J.-M. Fabre & M. Gonzales (Eds.), *Cognitive Biases.* North Holland: Elsevier.

Garnham, A. & Oakhill, J. (1994). *Thinking and reasoning.* Oxford: Blackwell.

Hardman, D.K. & Payne, S.J. (1995). Problem difficulty and response format in syllogistic reasoning. *The Quarterly Journal of Experimental Psychology*, 48A (4), 945--975.

Johnson-Laird, P.N. & Bara, B.G. (1984). Syllogistic inference. *Cognition*, 16, 1--62.

Johnson-Laird, P.N. & Byrne, R.M.J. (1991). *Deduction.* Hove: Lawrence Erlbaum Associates.

Markovits, H. & Nantel, G. (1989). The belief bias effect in the production and evaluation of logical conclusions. *Memory and Cognition*, 17, 11--17.

Newstead, S.E. Pollard, P. Evans, J.St.B.T., & Allen, J.L. (1992). The source of belief bias effects in syllogistic reasoning. *Cognition*, 45, 257--284.

Oakhill, J. & Garnham, A. (1993). On theories of belief bias in syllogistic reasoning. *Cognition*, 46, 87--92.

Oakhill, J. & Johnson-Laird, P.N. (1985). The effect of belief on the spontaneous production of syllogistic conclusions. *Quarterly Journal of Experimental Psychology*, 37A, 553--570.