Belief Bias, Logical Reasoning and Presentation Order on the Syllogistic Evaluation Task

Nicola J. Lambell (n.lambell@derby.ac.uk)
Institute of Behavioural Sciences; University of Derby
Mickleover, Derby, DE3 5GX, UK.

Jonathan St. B.T Evans and Simon J. Handley
(j.evans@plymouth.ac.uk, s.handley@plymouth.ac.uk)
Department of Psychology; University of Plymouth
Drake Circus, Plymouth, PL4 8AA, UK.

Abstract

Evans, Barston and Pollard, (1983) found that on the syllogistic evaluation task participants tended to endorse believable conclusions as being valid but reject unbelievable conclusions as invalid. A phenomenon known as "Belief Bias". Additionally, they collected verbal protocols from participants and established that this influence of belief was primarily associated with initial reference to the conclusions of these syllogistic arguments. In contrast, better logical reasoning was associated with initial reference to the premises. This experiment was designed to try to direct participants' attention to either the conclusion or the premises of a syllogistic argument with the intention of manipulating participants' logical reasoning ability and susceptibility to belief. The results reflected an inability to alter the influence of beliefs, but in one condition where the conclusion was presented prior to the premises, there was a successful reduction in participants' reasoning ability. The results are discussed with respect to the current theories of belief bias.

General Introduction

Traditional syllogisms consist of three statements, containing one of the four logical quantifiers *all, some, no and some...not* (see Example 1).The first two statements, the premises, each specify the relationship between an end term (A and C in Example 1) and a middle term (B). The conclusion, on the other hand, specifies a relationship between the two end terms.

All A are B
All B are C
Therefore,
All A are C

Example 1

There are three ways that psychologists utilise syllogisms in order to investigate logical reasoning. Firstly, in the evaluation task, the participant is asked to determine whether a given conclusion necessarily follows from the information given in the premises. Secondly, in the production task, the participant is required to produce a conclusion which necessarily follows from the information given in the premises. Finally, in the multiple choice task, the participant is required to choose a conclusion which necessarily follows from the premises from amongst a set of alternatives.

Early research into the multiple choice paradigm soon established that reasoning performance was poor and that participants demonstrated a systematic pattern of errors (c.f Chapman & Chapman, 1959; Dickstein, 1975; Erickson, 1974). Early research also established that the believability of the conclusion influenced the extent to which people endorsed conclusions (c.f Kaufman & Goldstein, 1967; Revlin, Leirer, Yopp & Yopp, 1980). This phenomenon has been termed belief bias. However, this early research into the influence of belief was heavily criticised on methodological grounds (see Evans, 1989; Evans, Barston & Pollard, 1983).

The first experiments to take account of these possible extraneous influences were the evaluation task studies performed by Evans, Barston and Pollard (1983). In these studies participants were presented with a conclusion for evaluation which was either valid (necessitated by the premises) or invalid (possible but not necessitated by the premises), believable or unbelievable. Their results clearly demonstrated that participants were more willing to accept believable than unbelievable conclusions as being valid. However, the greater acceptance of valid conclusions than invalid conclusions additionally demonstrated clear evidence of deductive capabilities.

What was surprising was that the influence of belief was more marked on invalid problems than valid problems. Whilst there was a small difference in acceptance of believable and unbelievable valid problems, there was a large difference in the acceptance of believable and unbelievable invalid problems. To facilitate understanding of their methodology and to demonstrate the large influence of belief bias observed on invalid problems consider the syllogisms presented in Examples 2 and 3. In Example 2 the conclusion to this syllogism is invalid as the class of highly trained dogs could include all police dogs. This would contradict the given conclusion. Yet a staggering

71% of participants across three experiments erroneously endorsed this conclusion as being valid.

No highly trained dogs are aggressive
Some police dogs are aggressive
Therefore,
Some highly trained dogs are not police dogs.

Example 2

No police dogs are aggressive
Some highly trained dogs are aggressive
Therefore,
Some police dogs are not highly trained.

Example 3

Now consider Example 3; this problem is logically equivalent to the previous example so this presented conclusion is also invalid. The only difference is that the terms have been re-arranged. However, only 10% of participants across three experiments erroneously endorse this conclusion as being valid (Evans et al., 1983). The difference between these two example syllogisms is the believability of the presented conclusion. Example 2 has a believable conclusion which participants tended to erroneously endorse as being invalid, whilst Example 3 has an unbelievable conclusion which participants tended to correctly reject as being invalid.

The pattern of results observed by Evans et al. led them to posit two possible explanations of how beliefs might influence reasoning. The first account, the Selective Scrutiny Model, proposes that participants initially scan the believability of the presented conclusion. When the conclusion is believable the model suggests that participants are likely to accept it without any consideration of its logical validity. On the other hand, when the conclusion is unbelievable the model proposes that some logical analysis takes place to determine whether the conclusion necessarily follows from the premises. The term "Selective Scrutiny" is derived from the fact that only unbelievable conclusions promote any attempts at reasoning. The model explains the belief bias phenomenon by proposing that participants do not reason when presented with believable conclusions. Instead they tend to unequivocally accept them as valid. This accounts for the high erroneous acceptance of invalid-believable conclusions and thus enables the model to explain the belief by logic interaction demonstrated by Evans et al. (1983).

The second account, the Misinterpreted Necessity Model, was motivated by the claims of Dickstein (1981) that participants often misunderstand logical necessity. The model proposes that participants initially engage in logical reasoning but tend to fall back on beliefs when reasoning is inconclusive (in other words, it fails to establish that the conclusion is falsified or necessitated by the premises). The model explains the belief by logic interaction as the invalid problems have conclusions which are consistent, but not necessitated, by the premises. Participants will, therefore, fall back on the believability of the conclusion to determine their response.

The robustness of the belief bias findings was subsequently demonstrated by Newstead, Pollard, Evans and Allen, (1992) and Evans, Newstead, Allen and Pollard (1994). Not only did they successfully replicate the pattern of results observed by Evans et al., they also found that increasing the logical nature of the instructions had little influence on the robustness of the belief bias findings. Their additional inclusion of belief-neutral conclusions also allowed them to investigate the direction of the belief bias effect. They found that only the unbelievable conclusions significantly differed in acceptance from the neutral conclusions, which enabled them to posit that belief bias is primarily associated with the rejection of unbelievable conclusions.

Other research into belief bias has primarily focused on the production task and has demonstrated that beliefs can also influence the production of conclusions (Oakhill & Johnson-Laird, 1985; Oakhill, Johnson-Laird & Garnham, 1989). Oakhill and her colleagues adapted the Mental Model theory of reasoning (Johnson-Laird & Byrne, 1991) to account for their findings. The Mental Model account proposes that participants construct an initial model of the premises and then produce a conclusion which is consistent with this model. The final deductive stage involves participants searching for alternative models of the premises in which this conclusion doesn't hold and positing other potential conclusions for similar evaluation. Oakhill hypothesised that believable conclusions influenced this process by curtailing a participants' willingness to search for alternative models of the premises. The model, therefore, explains the belief bias phenomenon by positing that all believable conclusions are unequivocally accepted without an attempt to search for falsifying models. This would explain the high erroneous acceptance of invalid believable conclusions found by Evans et al. (1983).

Some potentially important data which are often ignored are the concurrent and retrospective verbal protocol reports collected by Evans et al., (1983) during their experiments. Coding of these protocols revealed that when participants initially focused on the conclusion of a syllogistic argument they were more susceptible to the influences of belief. In contrast, when participants initially focused on the premises of a syllogistic argument they tended to show much better levels of logical responding.

This research is controversial because it suggests that belief bias arises due to the initial consideration of the conclusion, yet the production task findings demonstrated that belief bias occurred in the absence of a given conclusion. Thus, these findings appear inconsistent. Perhaps more importantly, these results are only consistent with the explanation of beliefs as proposed by the Selective Scrutiny Model. The other two models propose that participants initially attempt to reason from the premises thus reducing the extent to which belief bias should arise on the evaluation task. Thus, if belief bias on the evaluation task is due to initial focusing on the conclusion the Misinterpreted Necessity Model and Mental Model accounts are not adequate explanations of how beliefs influence performance.

The Experiment

The primary aim of this experiment was to direct participants attention to either the premises or the conclusion of a syllogistic argument in an attempt to encourage participants to adopt methods of responding akin to the verbal protocol findings. By presenting the conclusion of a syllogistic problem prior to its premises (CP condition) initial attention should be focused on the conclusion. Alternatively by presenting the premises of a syllogistic argument prior to its conclusion (PC condition) initial attention should be focused on the premises. However, Evans, et al. (1983) have already established from their protocol work that under this standard presentation condition some participants still give initial consideration to the conclusion. Therefore in order to enhance the focus of attention on the premises or conclusion some participants were given a delay (D condition) after the first piece of information was displayed whilst others received the standard simultaneous presentation of the problem information (N condition).

Two issues can be addressed by using this manipulation. First, if belief bias is a result of the initial consideration of the believability of the conclusion, then encouraging participants to focus on the conclusion should increase the influence that prior beliefs have on performance. Secondly, if logical performance is a result of the initial consideration of the premises, then encouraging participants to focus on the conclusion should reduce logical competence, whilst encouraging participants to focus on the premises should increase logical performance. These issues give rise to a number of testable predictions:

i) There should be more evidence of belief bias in the CP condition than the PC condition.

ii) There should also be less evidence of an effect of validity in the CP condition compared to the PC condition.

iii) There should be more evidence of belief bias in the CPD condition than the CPN condition.

iv) There should be more evidence of an effect of validity in the PCD condition than the PCN condition.

Evans, Handley and Buck, (1998) employed a similar methodology to investigate whether the order of information would affect performance on a conditional inference task. Using the Mental Model account of the reasoning process as a guideline of how participants were responding on the inference task they proposed that presenting the conclusion first should facilitate the building of models in which the conclusion held. They tested this proposal by suggesting that there should be a boosting of acceptance rates in the CP condition, especially for the more difficult Modus Tollens inference. They also explored whether the inclusion of a delay would exaggerate this effect. These hypotheses were not, however, confirmed. The absence of an increase in acceptance rates was explained by the boosting of the acceptance of affirmative conclusions for this condition in the presence of a general reluctance to endorse any conclusion which was presented first. They rejected their initial proposal in favour of the notion that participants' natural mode of reasoning was from the premises to the conclusion. This claim was supported by their additional finding that it took participants longer to respond to problems where the conclusion was presented first.

A secondary aim of this experiment was to compare our findings to the results observed by Evans et al. (1998). In addition to capturing the actual response that participants gave, a second dependant measure of how long participants took to respond to each problem was also included for comparison.

Method

Design. A four-way mixed design was used incorporating two between participant variables and two within participant variables. The first between participants variable was Presentation Order; this was whether they received the conclusion prior to the premises (CP) or the premises prior to the conclusion (PC). The PC condition acted as the control condition as it is the traditional syllogistic format used in the study of belief bias. The second factor was Delay; this was whether participants received the whole syllogism simultaneously (N condition) or whether a three second delay was introduced between the presentation of the first part of the syllogism and the rest of the problem (D condition). Participants were randomly assigned to one of four experimental groups; CPN, PCN, CPD, or PCD.

Each participant received eight syllogisms, half of which were valid (the conclusion presented followed logically from the premises) and half of which were invalid (the conclusion presented did not follow logically from the premises). This was the within participants variable of Validity. Half of these syllogisms had conclusions which were believable and half had conclusions which were unbelievable. This was the within participants variable of Belief. In all participants were given two valid-believable syllogisms, two valid-unbelievable syllogisms, two invalid-believable syllogisms and two invalid-unbelievable syllogisms.

Two dependent measures were taken during the experiment. The first dependent variable was a measure of whether participants accepted a conclusion as being valid or rejected a conclusion as being invalid. The second dependent variable was a measure of the time that participants took to evaluate a syllogistic problem.

Participants. Eighty undergraduate students from the University of Plymouth acted as paid volunteers in this experiment. None of them had any previous experience of syllogistic reasoning or any formal training in logic.

Materials. The EIO-2 form of the syllogisms and the materials employed in this experiment were identical to those employed by Evans et al. (1983). In order to control for any material differences two lists of materials were created. Conclusions which supported valid and believable arguments in list 1 supported invalid and unbelievable arguments in list 2. The syllogisms were presented to

participants using a computer program which controlled for the presentation of the premises and the conclusion.

Procedure. Participants were initially presented with a set of instructions to read on the computer screen. The bracketed information denotes additions to the instructions for conditions where the conclusion was presented first and a delay was introduced.

"This experiment is designed to find out how people solve logical problems. On the screen there will be a series of reasoning problems presented one at a time. You will be shown two premises which you should assume to be true and a conclusion which may or may not follow from these premises.

*In each case the premises are printed prior to the conclusion (**or the conclusion is printed prior to the premises**). You have to evaluate the conclusion in respect to the premises. **(There will be a short delay after the presentation of the conclusion before the premises are presented).***

A logical conclusion is one which has to be true, if the premises are true. If you believe that the conclusion must follow from the premises answer YES, otherwise NO. You must give your answer to each problem by pressing either the left or right hand mouse button as follows:

LEFT button- answer YES, the conclusion must follow from the premises.

RIGHT button- answer NO, the conclusion need not follow from the premises.

Please take your time and be sure that you have the logically correct answer before deciding.

If you have any questions please ask them now as the experimenter cannot answer any questions once you have begun the experiment. You are free to leave the experiment now or at any time during the presentation of the reasoning problems. Thank you very much for participating."

Participants were then shown the syllogisms one at a time on the computer screen. The latency measure was taken from when the last piece of information was shown on the screen to when the appropriate choice had been made. The program fully randomised the presentation of the eight syllogisms for each participant.

Results

Acceptance Responses. It was initially necessary to try to establish whether, in general, across all conditions there was a replication of the findings of Evans et al. (1983). Table 1 presents the overall percentage acceptance rates for four problem types.

Analysis revealed an effect of believability (sign test, 10/56, with 14 ties; p<.001, one-tailed), with substantially more believable (72%) than unbelievable conclusions (44%) being accepted by participants. A sizeable effect of logic was also established whereby 67% of participants accepted the conclusion of valid problems compared to 49% of participants who erroneously accepted the conclusions of invalid problems (sign test, 14/46, with 20 ties; p<.001, one tailed). This demonstrated strong evidence of logical competence. Sign test analysis also revealed a significant interaction between Logic and Belief (sign test, 42/17, with 21 ties; p<.001, one-tailed). Comparison of the effects of Belief on the valid and invalid problem types revealed that there was an effect of Belief on valid problems (sign test, 7/29, with 44 ties; p<.001, one-tailed) but the effect of Belief on invalid problems was slightly greater (sign test 6/51, with 23 ties; p<.001, one-tailed). It is interesting to note that like the Evans et al. (1983) study there was an effect of belief bias on valid problems.

Table 1: The overall percentage acceptance rates for the four problem types, collapsed across the four experimental conditions.

	Believable	Unbelievable	Combined
Valid	74	59	67
Invalid	70	29	49
Combined	72	44	58

Having successfully replicated the belief bias findings of Evans et al. (1983), it was then necessary to consider whether there was any general influence of the variables of Presentation Order and Delay on the acceptance rates of conclusions. Analysis revealed that there was a main effect of Presentation Order (U=1261.0; p<.001, two tailed) such that more conclusions were accepted in the PC condition (68%) than the CP condition (48%). Analysis of the Delay variable (U=1717.0; n.s, two tailed) revealed that there was no difference in the level of acceptance of conclusions in the delay condition (55%) compared to the no delay condition (61%).

Table 2: The percentage acceptance rates for the four problem types, divided according to the variable of Presentation Order and collapsed across Delay.

		Believable	Unbelievable	Combined
CP	Valid	59	50	54
	Invalid	63	23	43
	Combined	61	36	48
PC	Valid	89	69	79
	Invalid	78	35	56
	Combined	83	52	68

Analysis of the data with respect to the four predictions produced the following results:

i) There should be more evidence of belief bias in the CP condition than the PC condition. (see Table 2). For the CP condition there was a greater acceptance of believable (61%) than unbelievable conclusions (36%) (sign test, 7/24, with 9 ties; p<.002, one tailed). However, for the PC condition there was also a greater acceptance of believable (83%) than unbelievable conclusions (52%) (sign test 3/32, with 5 ties; p<.001, one tailed). A comparison of the effects of belief revealed similar levels of acceptance for both

conditions. (U=1535.5; n.s, one tailed). Thus, these results reflect that presenting the conclusion prior to the premises does not make participants more susceptible to the influences of belief. It is also interesting to note that significant belief by logic interactions were observed in both conditions.

ii) There should be less evidence of an effect of validity in the CP condition compared to the PC condition. (see Table 2). For the CP condition there was a slightly greater acceptance of valid (54%) than invalid conclusions (43%), however, analysis revealed the difference was not significant (sign test 8/17, with 15 ties; n.s, one tailed). In contrast, the greater acceptance of valid (79%) than invalid conclusions (56%) in the PC condition was significant (sign test, 6/29, with 5 ties; $p<.001$, one tailed). A comparison of the effects of validity across these two conditions revealed a significant difference (U=435.5; $p<.034$, one tailed). These findings, are therefore in clear support of prediction ii) and reflect that presenting the conclusion prior to the premises in some way disrupts logical responding.

Table 3: The percentage acceptance rates for the four problem types, divided according to the four experimental conditions.

		Believable	Unbelievable	Combined
CPN	Valid	68	53	60
	Invalid	65	20	43
	Combined	66	36	51
PCN	Valid	95	73	84
	Invalid	80	35	58
	Combined	88	54	71
CPD	Valid	50	48	49
	Invalid	60	25	42
	Combined	55	36	46
PCD	Valid	83	65	74
	Invalid	75	35	55
	Combined	79	50	64

Note : CPN =conclusion first with no delay, PCN = premises first with no delay CPD = conclusion first with delay, PCD = premises first with delay.

iii) There should be more evidence of belief bias in the CPD condition than the CPN condition. (See Table 3). In contrast to this prediction, the acceptance of believable (55%) compared to unbelievable conclusions (36%) in the CPD condition was slightly smaller than the acceptance of believable (66%) compared to unbelievable conclusions (36%) in the CPN condition. These results therefore reflect that the introduction of a delay had no influence on the effect of belief bias.

iv) There should be more evidence of an effect of validity in the PCD condition than the PCN condition. (See Table 3). In contrast to this prediction, the acceptance of valid (74%) compared to invalid conclusions (55%) in the PCD condition was slightly smaller than the acceptance of valid (84%) compared to invalid conclusions (58%) in the PCN condition. These results therefore reflect that the introduction of a delay had no influence on the effect of logical validity.

In summary, the order of presentation of information had a clear effect on logical performance but not susceptibility to belief bias. Encouraging focus of attention by the introduction of a delay had no effect on logical performance.

Latency Responses. Whilst no predictions were made concerning presentation order and delay variables, inspection of these latencies may provide insight into the possible differences between the presentation conditions. For the Presentation Order variable it was found that participants took significantly longer to respond to problems in the CP condition (M=19.4 seconds) than in the PC condition (M=15.6 seconds), ($F(1,76) = 4.39$, $p<.039$). A similar pattern of responding was reported by Evans et al. (1998). It was also found that when participants were given a delay during the presentation of a problem they took significantly less time to reach a decision (M=14.4 seconds) than when participants were given no delay (M=21.0 seconds), ($F(1,76) = 12.63$, $p<.001$). Again, this was reported by Evans et al. (1998). This result reflects the fact that some evaluation of the task is occurring during the delay.

In summary, the latency findings are consistent with the findings of Evans et al. (1998) that Presentation Order and Delay have clear influences on how long it takes participants to evaluate a given conclusion.

Discussion

The acceptance results reflect that belief bias does not necessarily arise from initial consideration of a syllogistic conclusion, as clear evidence of belief bias was apparent across all conditions. These findings are in contrast to the verbal protocol findings of Evans et al. (1983) and are more consistent with the findings of Oakhill et al. that belief bias occurs even in the absence of a conclusion for evaluation. It seems that whilst participants might adopt different methods of responding (as demonstrated by Evans et al.'s protocols) it has not been possible either to encourage people to adopt these methods or to demonstrate that focusing on the conclusion is primarily responsible for increasing susceptibility to belief bias. In hindsight, one plausible explanation for the absence of an increase in the effect of beliefs in the CP condition is that belief bias is already a very strong and robust phenomenon. It might not be possible to increase the levels of belief bias as we may have reached a ceiling effect in terms of the influence of beliefs.

What is slightly puzzling is that in the absence of any reduction in belief bias there was a clear reduction in logical performance when the conclusion was presented prior to the premises. This disruption in logical performance was clearly due to the order in which the information was presented as there was no additional effect of introducing a

delay. There was also a general suppression of acceptance of conclusions to all problems when the conclusion was presented prior to the premises. Both of these findings are consistent with the theoretical interpretation that Evans et al. (1998) posited to explain their findings.

Using the Mental Model theory as a framework for their explanation they proposed that participants' natural mode of reasoning was from the premises to the conclusion. This is consistent with the superior logical performance observed in the current experiment when the conclusion was presented first. Secondly, they proposed that presenting the conclusion first facilitates its inclusion in the initial model that participants construct. This attempt to include the conclusion as part of an initial model, could increase the difficulty that participants have in constructing any model of the premises. This would account for the finding in the current experiment of a suppression of acceptance of conclusions when they are presented first. Further support for this hypothesis comes from the latency findings which reflect the increased time taken to respond to problems when the conclusion is presented first. Whilst these findings should not be taken as clear support for the Mental Model account, the account does provide a useful framework in which to propose a possible explanation of the logical disruption caused by presenting the conclusion first.

The robust evidence of belief bias on all conditions fails to distinguish between the accounts of how beliefs influence performance on the syllogistic evaluation task. However, the superior reasoning performance demonstrated when the premises are presented first is consistent with the Mental Model theory and the Misinterpreted Necessity Model's notion that participants reason from the premises.

Perhaps the best way in which to view these findings is to adopt the distinction between our belief system and logical reasoning system proposed by Evans and Over (1996). They argue that when presented with a logical reasoning task participants attempt explicitly to comply with the logical instructions of the task but are unable to ignore the implicit influences of our beliefs. This would not only explain why it was possible to alter participants' logical reasoning performance by presenting the conclusion prior to the premises, it also suggests why the manipulations had no influence on the levels of belief bias observed.

Acknowledgments

This work was funded by a studentship grant awarded to Nicola Lambell by the Biotechnology and Biological Sciences Research Council (Reference: 94305619).

References

Chapman, L.J., & Chapman, J.P. (1959). Atmosphere effect re-examined. *Journal of Experimental Psychology, 58,* 220-226.

Dickstein, L.S. (1975). Effects of instructions and premise order on errors in syllogistic reasoning. *Journal of Experimental Psychology: Human Learning and Memory, 104,* 376-384.

Dickstein, L.S. (1981). Conversion and possibility in syllogistic reasoning. *Bulletin of the Psychonomic Society, 18,* 229-232.

Erickson, J.R. (1974). A set analysis theory of behaviour in formal syllogistic reasoning tasks. In R.L. Solso (Ed.), *Theories of cognitive psychology: The Loyola Symposium.* Hillsdale, NJ: Lawrence Erlbaum Associates.

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

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., Handley, S.J., & Buck, E. (1998). Ordering of information in conditional reasoning. *British Journal of Psychology, 89,* 383-403.

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

Evans, J.St.B.T., Newstead, S.E., Allen, J.L., & Pollard, P. (1994). Debiasing by instruction: The case of belief bias. *European Journal of Cognitive Psychology, 6* (3) 263-285.

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

Kaufman, H., & Goldstein, S. (1967). The effects of emotional value of conclusions upon distortions in syllogistic reasoning. *Psychonomic Science, 7,* 367-368.

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

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.

Oakhill, J., Johnson-Laird, P.N., & Garnham, A. (1989). Believability and syllogistic reasoning. *Cognition, 31,* 117-140.

Revlin, R., Leirer, V.O., Yopp, H. & Yopp, R. (1980). The belief bias effect in formal reasoning: The influence of knowledge on logic. *Memory and Cognition, 8,* 584-592.