

This is an Author's Original Manuscript of an article submitted for consideration in the Journal of Cognitive Psychology [copyright Taylor & Francis]; Journal of Cognitive Psychology is available online at

<http://www.tandfonline.com/10.1080/20445911.2011.589381>

Running Head: WHEN LOGIC AND BELIEF COLLIDE

When Logic and Belief Collide: Individual Differences in Reasoning Times
Support a Selective Processing Model

Edward J. N. Stupple

University of Derby, UK

Linden J. Ball*

Lancaster University, UK

Jonathan St. B. T. Evans

University of Plymouth, UK

Emily Kamal-Smith

University of Derby, UK

Abstract

When the validity of a deductive conclusion conflicts with its believability people often respond in a belief-biased manner. This study used response times to test the selective processing model, which views belief-bias effects as arising from the interplay between superficial heuristic processes and more rigorous analytic processes. Participants were split into three response groups according to their propensity to endorse logically normative conclusions. The low-logic, high belief-bias group demonstrated rapid responding, consistent with heuristic processing. The medium-logic, moderate belief-bias group showed slower responding, consistent with enhanced analytic processing, albeit selectively biased by conclusion believability. The high-logic, low belief-bias group's relatively unbiased responses came at the cost of increased processing times, especially with invalid-believable conclusions. These findings support selective processing claims that distinct heuristic and analytic processing systems underpin reasoning, and indicate that certain individuals differentially engage one system more than the other. A minor amendment is proposed to the current selective processing model to capture the full range of observed effects.

When Logic and Belief Collide: Individual Differences in Reasoning Times

Support a Selective Processing Model

During deduction a conclusion's logical status may conflict with background beliefs. For example, the conclusion to the following syllogism is logically valid despite being unbelievable: "All mammals can walk. Whales are mammals. Therefore, whales can walk". In contrast, the conclusion to the syllogism below is logically invalid despite its apparent believability: "All flowers need water. Roses need water. Therefore, roses are flowers" (both examples are taken from De Neys & Franssens, 2009). When logic and belief collide, as in these previous syllogisms, people often respond on the basis of their prior knowledge, giving rise to a "belief-bias" effect (see Evans, 2007b, for a review). For example, in their seminal study, Evans, Barston, and Pollard (1983) observed a 71% conclusion acceptance rate for invalid-believable conflict problems versus a 10% conclusion acceptance rate for invalid-unbelievable non-conflict problems. The conclusions to both of these invalid problems should, in fact, be rejected. Evans et al. likewise observed a conclusion acceptance rate of 56% for valid-unbelievable conflict problems versus 89% for valid-believable non-conflict problems. The conclusions to both of these valid problems should be accepted.

The pattern of acceptance rates observed by Evans et al. (1983) gave rise to three statistically significant effects: (1) a main effect of a logical validity; (2) a main effect of belief; and (3) a logic-by-belief interaction – since belief had a greater impact on invalid than valid problems. All three effects have been replicated many times (e.g., Klauer, Musch, & Naumer, 2000; Quayle & Ball, 2000; Stuppel & Ball, 2008; Thompson, Striener, Reikoff, Gunter, & Campbell, 2003), with arguably the most comprehensive explanation of such findings being the "selective processing

model” (e.g., Evans, 2000, 2007a; Klauer et al., 2000). This dual-process model (Figure 1) posits the involvement of two distinct types of processes in reasoning: relatively superficial, associative “heuristics” driven by prior knowledge and beliefs, and more rigorous “analytic” processes involving rule-based inference.

(Figure 1 about here)

According to selective processing theorists, belief bias is determined by the operation of both the heuristic *and* the analytic components of the model depicted in Figure 1. The default heuristic response is to accept believable and reject unbelievable conclusions. This explains why belief bias arises with both valid and invalid inferences, thereby accounting for the main effect of belief. However, the analytic component of the model is needed to explain the logic-by-belief interaction. If the analytic system intervenes during reasoning then a mental simulation process attempts the construction of a single “mental model” of the premises. This analytic process is, however, itself biased by the believability of conclusions such that reasoners are viewed as operating in a “satisficing” manner (Evans, 2007a). Thus, for a believable conclusion a satisficing search is initiated for a single mental model that *supports* the conclusion, whereas for an unbelievable conclusion a satisficing search is initiated for a single mental model that *refutes* the conclusion (Figure 1).

These latter assumptions provide a clear rationale for the emergence of a logic-by-belief interaction. When conclusions are valid, despite unbelievable content motivating a search for a counterexample model, such a model cannot be found, thus limiting the influence of belief-bias. When conclusions are invalid, however, models exist that both support and refute such conclusions, thus leading to high levels of erroneous acceptance of invalid-believable items and high levels of correct rejection of invalid-unbelievable items.

Despite the appeal of the selective processing model, recent studies measuring the time taken to evaluate presented conclusions appear to challenge the account since they reveal increased processing times for conflict problems relative to non-conflict problems, particularly on invalid-believable items (Ball, Wade, Phillips, & Quayle, 2006; Stupple & Ball, 2008; Thompson et al., 2003). The model in Figure 1 provides no obvious reason for why conflict problems – especially invalid-believable ones – should be subjected to such increased processing effort. According to the model, an element of time-consuming analytic processing should be equally likely for *all* problem types.

To reconcile the selective processing model with chronometric evidence this paper addresses the possibility that response-time effects reflect individual differences [in people's tendency to respond normatively](#) to belief-oriented problems. In particular, we hypothesised that the relatively long response latencies evident in aggregate data for invalid-believable problems may primarily reflect the performance of a sub-group of individuals who are frequently able to understand the underlying logic of these problems (leading to the rejection of fallacious conclusions), but whose reasoning strategy requires increased processing effort, leading to prolonged response latencies for these problems relative to other ones.

Supportive evidence for these proposals comes from a study by Sá, West, and Stanovich (1999), who demonstrated that high cognitive ability, skills in cognitive decontextualisation and dispositions toward active, open-minded thinking are all markers [for the tendency](#) to avoid responding in terms of the real-world content of syllogistic conclusions. Sá et al.'s research therefore attests to the existence of reasoners who are motivated and able to decontextualise their reasoning so as to avoid what Stanovich (1999) refers to as the “fundamental computational bias” (i.e., the bias

to reason according to content rather than form). As Figure 1 indicates, *no-one* should be able to respond logically to invalid-believable problems via either the heuristic or the analytic routes within the current version of the selective processing model. The heuristic route would lead to conclusion acceptance (as a response bias), while the analytic route would permit the discovery of a confirmatory model (via a satisficing process) that would also warrant conclusion acceptance. So, participants who *do* manage to respond logically by rejecting invalid-believable conclusions must presumably have a unique ability to avoid falling foul of both heuristic *and* analytic biases. In effect, this latter sub-group of participants would fall outside the explanatory reach of the currently formulated selective processing model, which captures the behaviour of the majority of reasoners who demonstrate a degree of logically erroneous responding when tackling belief-oriented problems. The apparent inability of the present version of the model to account for the behaviour of *all* reasoners is clearly problematic, and tackling this issue formed the central motivation for the study presented here.

Method

Participants

The study involved 130 participants: 70 undergraduates from the University of Derby and 60 Oxfam employees (age range: 18 to 60 years). All were untrained in logic and the psychology of reasoning.

Design

A standard belief-bias paradigm was used involving a 2 x 2 repeated-measures design that systematically manipulated the logic of conclusions (valid vs. invalid) and their belief status (believable vs. unbelievable). Dependent measures were conclusion acceptance rates and response times to register accept/reject decisions.

Materials and Procedure

Sixteen syllogisms were presented: eight conflict problems (four valid-unbelievable; four invalid-believable) and eight non-conflict problems (four valid-believable; four invalid-unbelievable). Standard experimental controls were implemented relating to syllogistic forms (AB-BC vs. BA-CB figures), logical quantifiers (IEO vs. EIO moods) and preferred conclusion orders (A-C vs. C-A conclusions). These controls neutralise the effect of various response biases that can have an impact on performance (Evans et al., 1983; Stupple & Ball, 2007).

Problem content was based on that used by Quayle and Ball (2000), where unbelievable conclusions were false by definition (e.g., *Some snakes are not reptiles*), while believable conclusions were true by definition (e.g., *Some reptiles are not snakes*). There were equal numbers of valid and invalid problems and equal numbers of believable and unbelievable conclusions. The presentation order of syllogisms was counterbalanced using a balanced Latin square design, with contents rotated through the 16 problems. Authorware 6.5 for Windows was used to present syllogisms and record responses and response times.

A logic index was calculated for each participant based on responses to all 16 syllogisms, enabling participants to be classified as high-logic, medium-logic or low-logic responders. This logic index ($VB + VU - IB - IU$) ranged from -8 to +8 and measured the difference between an individual's acceptance of valid and invalid conclusions; the bigger the index, the more normative the individual's responding.

Results

Aggregate Data: Manipulation Check

The data for conclusion acceptance rates are depicted in Figure 2a, and follow the standard pattern established by Evans et al. (1983) and replicated in subsequent

studies. More valid conclusions were accepted than invalid conclusions (28% difference), more believable conclusions were accepted than unbelievable conclusions (25% difference), and the effect of belief was more marked on invalid than valid conclusions (nearly double the magnitude). A 2 x 2 repeated-measures ANOVA on the conclusion acceptance data confirmed the significance of the main effect of logic, $F(1, 129) = 153.20$, $MSE = 1.07$, $p < .001$, $\eta_p^2 = 0.54$, the main effect of belief, $F(1, 129) = 87.37$, $MSE = 1.47$, $p < .001$, $\eta_p^2 = 0.40$, and the belief-by-logic interaction, $F(1, 129) = 18.11$, $MSE = 0.82$, $p < .001$, $\eta_p^2 = 0.12$.

(Figure 2 about here)

Response-time data (Figure 2b) also replicated previous findings: conflict problems took longer to process than non-conflict problems, with particularly long times for invalid-believable items. Response time data were subjected to a log transformation to eliminate positive skew prior to inferential analysis. A 2 x 2 repeated-measures ANOVA revealed longer processing times (mean difference: 2.40s) for problems with invalid as opposed to valid conclusions, $F(1, 129) = 22.61$, $MSE = .014$, $p < .001$, $\eta_p^2 = 0.15$, but no difference between problems with believable versus unbelievable conclusions, $F < 1$. However, as seen in previous chronometric studies, the belief-by-logic interaction was highly reliable, $F(1, 129) = 25.83$, $MSE = .015$, $p < .001$, $\eta_p^2 = 0.17$. Exploring this interaction using simple main effects analyses indicated significant differences between the processing times of the conflict and non-conflict problems for three out of four comparisons ($F_s > 9.5$; $p_s < .005$), with the one exception being the comparison between valid-unbelievable and invalid-unbelievable problems ($F < 1$). The present chronometric data confirm that the invalid-believable conflict items are unique in terms of the longer time required to evaluate their validity status.

Response Time Analyses for Response Groups

Our core prediction was that the long response latencies evident in aggregate data for invalid-believable problems primarily reflect the performance of a sub-group of individuals whose dominant response to belief-oriented problems is normative. To test this prediction, participants were assigned to groups on the basis of their logic index. This enabled partitioning of individuals into three response groups: high-logic responders (logic index ≥ 4 , $N = 34$), medium-logic responders (logic index from 2-3, $N = 48$) and low-logic responders (logic index ≤ 1 , $N = 48$). These partitions were driven not only by the desire to identify three distinct response groups but also by the need to preserve the integrity of subsequent statistical analyses.

We suggest that categorising people into different groups on the basis of their responses to all 16 syllogisms represents an ideal way to demarcate individual differences in reasoning abilities and strategies for the purpose of our study. This is because a participant's logic index represents a relatively pure measure of their success in reaching a normative conclusion to belief-oriented problems, rather than a proxy measure based on an assessment of working memory capacity or intelligence. Although these latter measures are certainly correlated with logical reasoning (Kyllonen & Christal, 1990) they also embody additional variance unrelated to [such reasoning](#), which could undermine the opportunity to detect clear-cut differences in the chronometric profiles of high-, medium, and low-logic responders within the present paradigm. It was only by grouping participants according to these latter three categories that we could be certain to provide definitive evidence regarding our prediction that high-logic responders have longer response latencies for invalid-believable problems relative to other response groups and problem types.

Figure 3 presents response-time data (log-transformed but converted back into original units) for the four problem types broken down by response group. A mixed-design ANOVA revealed a significant main effect of problem type, $F(3, 381) = 17.64$, $MSE = .015$, $p = .001$, $\eta_p^2 = 0.12$, with valid-believable problems being processed faster than all other problem types (all $ps < .001$) and invalid-believable problems being processed slower than all other problem types (all $ps < .001$). Valid-unbelievable and invalid-unbelievable problems did not differ ($p = .81$). There was also a significant main effect of response group, $F(2, 127) = 11.33$, $MSE = .152$, $p < .001$, $\eta_p^2 = 0.15$, with post-hoc tests (Bonferroni corrected) showing that: (1) the high-logic group was 8.9s slower at responding than the low-logic group ($p < .001$); (2) the medium-logic group was 6.74s slower at responding than the low-logic group ($p = .001$); and (3) the high-logic group was 2.23s slower than the medium-logic group, although this latter difference was not significant ($p = .99$).

(Figure 3 about here)

Importantly for our predictions, this ANOVA also revealed the presence of a significant interaction between problem type and response group, $F(6, 381) = 2.45$, $MSE = .015$, $p = .025$, $\eta_p^2 = 0.04$. Simple main effects analyses were conducted to unpack this interaction. Response group impacted on response times for valid-unbelievable conflict problems, $F(2, 205.9) = 4.94$, $MSE = .049$, $p = .008$, with Games-Howell tests indicating that the high-logic and medium-logic groups processed these problems more slowly (by 6.60s and by 5.39s, respectively) than the low-logic group ($p = .027$; $p = .029$). The high-logic and medium-logic groups did not differ ($p = .91$). The group differences were even more pronounced for the invalid-believable problems, $F(2, 205.9) = 16.59$, $MSE = .048$, $p < .001$, with the high-logic and medium-logic groups being slower (by 14.33s and by 9.33s, respectively) than the

low-logic group (both $ps < .001$), and with the high-logic group being 5.00s slower than the medium-logic group, but not reliably so ($p = .24$).

The same pattern emerged for the non-conflict problems, with a reliable simple main effect for valid-believable problems, $F(2, 205.9) = 6.27$, $MSE = .049$, $p = .002$, underpinned by significant differences between the low-logic group and the other groups (both $ps < .01$), while the latter groups did not differ ($p = .76$). Invalid-unbelievable problems also demonstrated this pattern, $F(2, 205.9) = 9.49$, $MSE = .049$, $p < .001$ (high-logic vs. low-logic, $p = .001$; medium-logic vs. low-logic, $p = .002$; medium-logic vs. high-logic responders, $p = .93$). Although these data perhaps indicate a general “cautiousness” in reasoning as normative performance improves, in the following simple main effects analyses conducted at the level of response groups it is clear that normative responders also vary in their response times across problem types, which militates against a pure cautiousness effect.

Simple main effects analyses conducted at each level of logic group (with a Bonferroni corrected alpha of $p < .007$ for follow-up tests) indicated that although there was a difference in response times across problem types for the low-logic group, $F(3, 141) = 3.65$, $MSE = .012$, $p = .014$, there were no reliable follow-up comparisons. For the medium-logic group the contrast in response times across problem types was more pronounced, $F(3, 141) = 7.76$, $MSE = .016$, $p < .001$, but only the comparison between invalid-believable and valid-believable items was reliable ($p < .001$). Critically, however, the uniqueness of the invalid-believable items was particularly salient in the high-logic group, $F(3, 99) = 8.97$, $MSE = .017$, $p < .001$, where these items displayed increased response times compared with valid-believable ($p < .001$), valid-unbelievable ($p = .002$) and invalid-unbelievable ($p = .002$) items, with no contrasts between other items nearing significance ($ps > .50$).

Overall, the interaction between problem type and response group supports our core hypothesis that the relatively long response latencies evident in aggregate data for invalid-believable problems primarily reflect the performance of normatively-responding individuals whose assiduous analysis of these problems leads to prolonged response latencies.

Regression Analysis for all Problem Types

A further strength of the methodology employed is that it permits an examination of whether only response times for invalid-believable items are uniquely predictive of normative responding (which would be congruent with our overarching hypothesis) or whether response times for each problem type predicts normative responding. To examine this issue we treated normative responses (as measured by the logic index) as a continuous variable rather a categorical one so as to increase the sensitivity of the analysis; we then pursued a multiple regression analysis to examine the extent to which the log-transformed response times for the differing problem types predicted normative responding.

The model that included all four problem types was highly reliable, $R = .45$, adjusted $R^2 = .18$, $F(4, 125) = 7.93$, $p < .001$. Standardised regression coefficients for each problem type indicated that an increase in response times for invalid-believable problems was associated with an increase in logic index, standardised Beta = .58, $t(125) = 4.06$, $p < .001$. However, none of the other problem types made a significant independent contribution toward accounting for the variance in the logic index: valid-believable, standardised Beta = -.10, $t(125) = -0.75$, $p = .46$; valid-unbelievable, standardised Beta = -.05, $t(125) = -0.40$, $p = .69$; and invalid-unbelievable, standardised Beta = -.03, $t(125) = -0.25$, $p = .81$. These regression findings confirm that the most important factor in accounting for levels of overall normative

responding was the processing effort devoted by participants to the more difficult invalid-believable conflict problems. To put it another way, those participants who worked the hardest at trying to reason out a solution to the invalid-believable problems were those who also demonstrated the highest logic indices.

Item-Based Analysis

As a final analysis we checked whether the response-time effect observed for invalid-believable problems generalised across all four of the invalid-believable problem forms used in the experiment or whether the effect was restricted to a few invalid-believable items. The four invalid-believable problems forms varied in mood and figure as follows: (1) AB-BC, IEO; (2) BA-CB, EIO; (3) BA-CB, IEO; and (4) AB-BC, EIO. There were no significant differences in response times across these four forms, and neither was there a significant form-by-group interaction (both $F_s < 1$). We also note that these problem forms were presented in eight different counterbalanced versions and were rotated across all problem contents. We are confident, therefore, that response-time effects for invalid-believable problems were not driven by a sub-set of items, but instead generalised across all items.

Discussion

Our findings reveal clear individual differences in response times for belief-oriented syllogisms, with invalid-believable conflict items showing particularly marked variations in response times relative to other problems types as well as across response groups. These findings, moreover, are broadly interpretable in terms of the assumptions of the selective processing model of belief bias presented in Figure 1 (Evans, 2007a). Our low-logic group (overall accuracy $\leq 56\%$) responded reliably faster to all problems than either the medium-logic or the high-logic groups, and response times for different problem types within this group did not differ. As such,

individuals in the low-logic group appeared to be operating primarily on the basis of a response bias as depicted as the default mode of responding in the selective processing model (i.e., they mostly evaluated conclusions by means of a rapid, low-effort, heuristic route that was driven by the belief status of presented conclusions).

Individuals in the medium-logic group (overall accuracy 57%-74%) demonstrated significantly longer response times for all problems compared with the low-logic group, and also took longest to process the invalid-believable items, although only the time difference between invalid-believable and valid-believable problems was reliable. In terms of the standard selective processing model (Figure 1) these participants appear to be reasoning via the analytic route (e.g., conducting a satisficing search for a mental model that supports invalid-believable conclusions and refutes valid-unbelievable conclusions). This analytic search results in increased response times relative to the low-logic group and allows these individuals to demonstrate a modicum of logical competence, especially when dealing with the valid-unbelievable conclusions where a refuting model does not exist.

Finally, individuals in the high-logic group (overall accuracy $\geq 75\%$) again demonstrated significantly longer response times for all problems compared to the low-logic group, and also revealed significantly increased response times for the invalid-believable problems relative to *all* other problem types. This latter finding suggests that these individuals are doing more than merely searching for a supporting model for a presented invalid-believable conclusion, since such a model is available and would encourage conclusion acceptance rather than the conclusion rejection that was typically observed. We propose that normatively responding reasoners are often uniquely able to avoid analytic processing biases so as to engage in a diligent search for counterexample models even when the presented conclusions are believable and

consistent with possible models of the premises. However, this rigorous search process, coupled with the need to overcome analytic processing biases, incurs the observed processing cost in terms of additional time on task.

Overall, our individual differences analysis of response times enables us to reconcile the selective processing model of belief bias (Evans, 2007a) with chronometric evidence that, at first sight, appeared to be troublesome for this model when viewed at an aggregate level (Ball et al., 2006; Stuppel & Ball, 2008). Our more penetrating analysis of the data indicates that people of differing logical ability have different response-time profiles, with these profiles providing good evidence to support key selective processing assumptions. Nevertheless, we acknowledge that the current version of the selective processing model (Figure 1) does not explicitly capture the reasoning behaviour of those participants whose dominant response to belief-oriented syllogisms is normative, and we therefore propose a slightly modified model that can accommodate this sub-set of reasoners who conduct a more exhaustive search for counterexample models (Figure 4).

(Figure 4 about here)

The reconciliation between chronometric evidence and the selective processing model is important since this model arguably provides the most comprehensive extant account of belief-bias phenomena (Ball, 2010). Moreover, the present findings resonate well with the broader dual-process literature in which similar chronometric patterns have been shown in judgment and decision-making (JDM) tasks, where participants demonstrate inflated response times to conflict problems (e.g., Bonner & Newell, 2010; De Neys & Glumicic, 2008). [An interesting aspect of these previous studies is that while the response-time effect for conflict problems was particularly marked for high-ability reasoners it was nevertheless still](#)

present, though diminished, for low-ability reasoners. At first sight this finding seems to depart from our data (Figure 3), which appear to demonstrate that the increased latencies for conflict versus non-conflict problems are restricted to the more normative responders and do not arise in the low-logic group. However, the predicted trend is hinted at in our low-logic group, and when we analysed the latency data with a direct planned contrast the evidence for significantly longer latencies for conflict versus non-conflict problems was reliable, $t(47) = 1.91, p = .031$, one-tailed. In this way the present data seem compatible with findings from the JDM literature and suggest that even in our low-logic group there is some degree of sensitivity to logic-belief conflicts, perhaps arising from the processing approach of a sub-set of these individuals.

We finally note that we have recently become aware of another contemporary study by Thompson, Morley, and Newstead (2011) that also takes an individual differences stance in examining chromomeric data within a belief-bias paradigm. Thompson et al.'s study employed an analysis strategy directly equivalent to that advanced in the present paper, whereby they split their sample into three groups based on their syllogistic reasoning performance and analysed response times as a function of group. Their findings, broadly speaking, were comparable to ours, with long latencies for invalid-believable problems being most marked for the highest performing group (i.e., individuals possessing demonstrable analytic ability).

We give full credit to Thompson et al. (2011) for their important and timely observations but we also note methodological reasons for why our study represents an important replication and extension of their research. First, their analysis was based on data from five experiments that involved noticeable variations in syllogistic figure, materials, testing conditions (groups vs. individuals), delays (simultaneous vs. non-

simultaneous presentation of premises and conclusion) and truth types (definitionally vs. empirically believable conclusions). The resulting lack of control arising from this combined data analysis could have increased Type I and Type II error rates. Second, Type I and Type II errors may also have arisen from the way Thompson et al. removed extreme data-points, which involved eliminating the top and bottom 2% of responses. The data transformation approach that we applied to normalise response times and stabilise variances is arguably superior in counteracting violations of parametric test assumptions. In sum, we contend that the close correspondence between our response time findings and those of Thompson et al. (2011) presents a compelling independent corroboration of their data. Where we differ markedly, however, is in our theoretical interpretation of findings. Thompson et al. claim that a dual-process theory such as the selective processing model cannot accommodate evidence for increased times arising for invalid-believable problems relative to other problem types. As we have demonstrated, this claim is not sustainable since a modified selective processing model can readily accommodate the full pattern of processing times seen across different response groups.

Thompson et al. (2011) further argue that the only existing explanation for their findings is the one presented by Thompson et al. (2003), which accounts for long processing times for invalid-believable problems as arising from two factors: (1) people's motivation to work harder with *believable* conclusions, which are claimed to be more palatable; and (2) people's difficulty in evaluating invalid conclusions, which have been shown to require more mental-model construction relative to valid conclusions (e.g., Quayle & Ball, 2000). Interestingly, Thompson et al.'s assumptions predict that invalid-believable problems will have long response times regardless of whether a participant responds "yes" (non-normatively) or "no" (normatively). **They**

claimed support for this prediction based on trends evident in latencies across problem types, whereby participants appeared to take longer to respond to invalid-believable items compared to other items whether analyses were restricted to “no” responses or to “yes” responses. However, no inferential statistics were presented to validate these trends.

To address this latter issue in our dataset we pursued two separate one-way ANOVAs comparing mean latencies across problem types, the first analysis focusing on participants ($N = 75$) who responded “yes” to at least one problem of each type and the second analysis doing the same for participants ($N = 38$) who responded “no” to at least one problem of each type. Neither analysis provided evidence in support of Thompson et al.’s prediction, although we acknowledge that our data may be limited by the focus on “inconsistent” responders. Thompson et al.’s prediction contrasts with the modified selective processing view, which is that normatively correct “no” responses to invalid-believable problems should be slower (i.e., more effortful) than non-normative “yes” responses. As a test of this prediction we conducted a t-test on the latency data for participants with inconsistent response profiles to invalid-believable problems (i.e., where they responded “no” to some of these problems and “yes” to others). We observed that the rejection of invalid-believable conclusions took significantly longer than their acceptance (25.77s vs. 21.18s), $t(85) = 3.36$, $p < .001$, $d = .30$, fully supporting the modified selective processing account.

To provide further support for dual-process accounts of belief bias such as the selective processing model, research could benefit from in-depth exploration of the patterns of neurological activation associated with different responses on logic-belief conflict problems (see De Neys, Vartanian, & Goel, 2008; Goel & Dolan, 2003, for preliminary research). Goel and Dolan (2003), for example, have provided evidence

that different reasoning strategies underpin normative versus non-normative responses to conflict problems, with normative responses engaging the right lateral prefrontal cortex (a region concerned with cognitive monitoring), and belief-based responses engaging the ventral medial prefrontal cortex (implicated in affective processing). As proposed in this paper, the capacity to generate normative responses to certain conflict problems presumably comes at the cost of having to devote extra time not only to inhibiting default, heuristic responses but also to avoiding selective-processing biases associated with analytic intervention. Both inhibitory mechanisms may well require intensive cognitive monitoring mediated by the ventral medial prefrontal cortex. Future neuroimaging studies should ideally pay particular attention to individual differences in reasoning strategies in order to provide further insights into the mechanisms involved in conflict-detection and response inhibition.

Acknowledgements

We thank the University of Derby for funding part of this study through their Research Inspired Curriculum Fund. We are grateful to Andrew Conway, Thomas D. Griffin and two anonymous reviewers for providing valuable critical comments on previous versions of this paper.

References

Ball, L. J. (2010). The dynamics of reasoning: Chronometric analysis and dual-process theories. In K. I. Manktelow, D. E. Over, & S. Elqayam (Eds.), *The Science of Reason: A Festschrift for Jonathan St. B. T Evans* (pp. 283-307). Hove: Psychology Press.

Ball, L. J., Phillips, P., Wade, C. N., & Quayle, J. D. (2006). Effects of belief and logic on syllogistic reasoning: Eye-movement evidence for selective processing models. *Experimental Psychology*, *53*, 77-86.

Bonner, C., & Newell, B. R. (2010). In conflict with ourselves? An investigation of heuristic and analytic processes in decision making. *Memory & Cognition*, *38*, 186-196.

De Neys, W., & Franssens, S. (2009). Belief inhibition during thinking: Not always winning but at least taking part. *Cognition*, *113*, 45-61.

De Neys, W., & Glumicic, T. (2008) Conflict monitoring in dual process theories of thinking. *Cognition*, *106*, 1248-1299.

De Neys, W., Vartanian, O., & Goel, V. (2008). Smarter than we think: When our brains detect that we are biased. *Psychological Science*, *19*, 483-489.

Evans, J. St. B. T. (2000). Thinking and believing. In J. García-Madruga, N. Carriedo, & M. J. González-Labra (Eds.), *Mental Models in Reasoning* (pp.41-56). Madrid: UNED.

Evans, J. St. B. T. (2007a). *Hypothetical Thinking: Dual Processes in Reasoning and Judgement*. Hove, UK: Psychology Press.

Evans, J. St. B. T. (2007b). On the resolution of conflict in dual process theories of reasoning. *Thinking & Reasoning*, *13*, 378-395.

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

Goel, V., & Dolan, R. J. (2003). Explaining modulation of reasoning by belief. *Cognition*, *87*, B11-B22.

Klauer, K. C., Musch, J., & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review*, *107*, 852-884.

Kyllonen, P. C., & Christal, R. E. (1990). Reasoning ability is (little more than) working-memory capacity? *Intelligence*, *14*, 389-433.

Quayle, J. D., & Ball, L. J. (2000). Working memory, metacognitive uncertainty and belief bias in syllogistic reasoning. *Quarterly Journal of Experimental Psychology*, *53A*, 1202-1223.

Sá, W. C., West, R. F., & Stanovich, K. E. (1999). The domain specificity and generality of belief bias: Searching for a generalizable critical thinking skill. *Journal of Educational Psychology*, *91*, 497-510.

Stanovich, K. E. (1999). *Who is Rational? Studies of Individual Differences in Reasoning*. Mahwah, NJ: Lawrence Erlbaum Associates.

Stuppel, E. J. N., & Ball, L. J. (2007). Figural effects in a syllogistic evaluation paradigm: An inspection-time analysis. *Experimental Psychology*, *54*, 120-127.

Stuppel, E. J. N., & Ball, L. J. (2008). Belief-logic conflict resolution in syllogistic reasoning: Inspection-time evidence for a parallel-process model. *Thinking & Reasoning*, *14*, 168-181.

Thompson, V. A., Morley, N. J., & Newstead, S. E. (2011). Methodological and theoretical issues in belief-bias: Implications for dual process theories. In K. I. Manktelow, D. E. Over, & S. Elqayam (Eds.), *The Science of Reason: A Festschrift for Jonathan St. B. T. Evans* (pp. 309-338). Hove: Psychology Press.

Thompson, V. A., Striener, C. L., Reikoff, R., Gunter, R. W., & Campbell, J. D.
(2003). Syllogistic reasoning time: Disconfirmation disconfirmed. *Psychonomic
Bulletin & Review*, *10*, 184-189.

Figure 1:

The selective processing model of belief bias. This figure was previously published in Evans (2007a; Figure 4.3, p. 91) and is reproduced here by permission of Taylor & Francis Group.

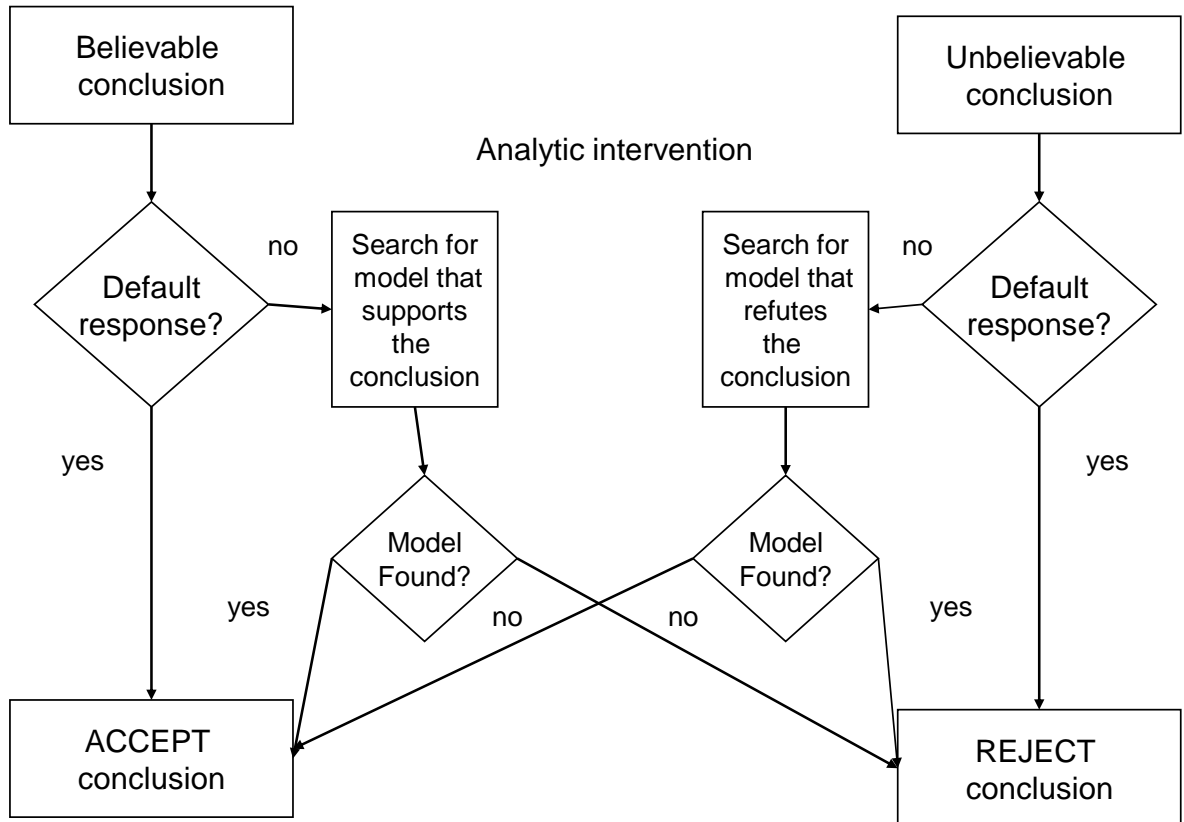



Figure 2:

(a) Mean percentage conclusion acceptance rates (+/- SE) and (b) mean response time in seconds (+/- SE) for all problem types as a function of logic and belief

(a)



(b)

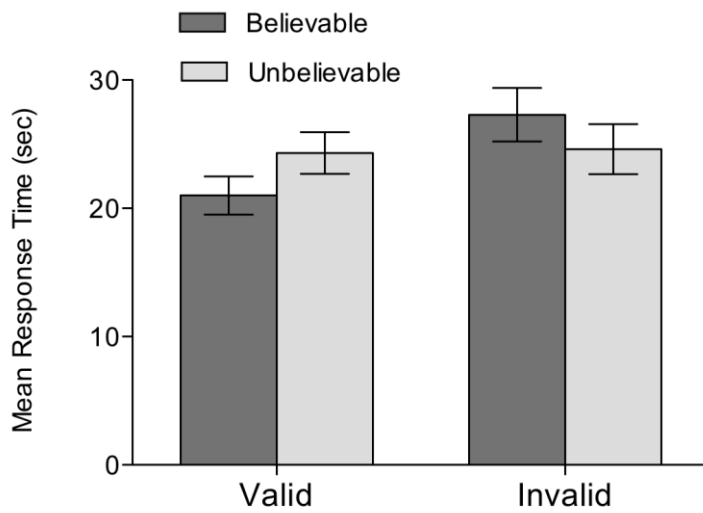


Figure 3:

Mean response time in seconds (transformed data converted to original units) for all problem types as a function of response group

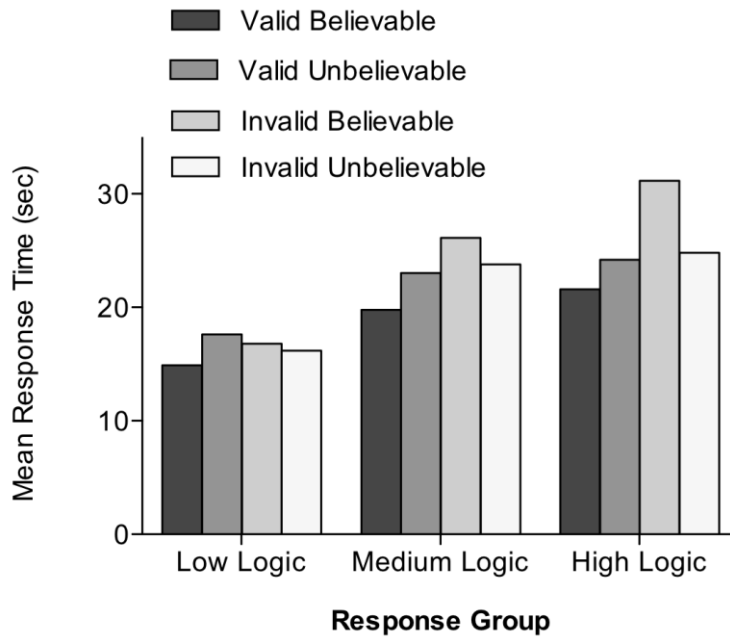


Figure 4:

A modified selective processing model of belief bias that captures the behaviour of three sub-groups of reasoners: (1) non-normative responders, who do not engage in analytic inhibition, instead operating on the basis of a response bias; (2) satisficing responders, who reason via the analytic route and conduct a satisficing search for a single mental model that supports a believable conclusion or that refutes an unbelievable conclusion; and (3) normative responders, who reason via the analytic route and conduct an exhaustive search for counterexample models.

