

Familiarity Effects and Questioning Biases in Human Belief Revision

Ann G. Wolf (ann.g.wolf@psychol.uni-giessen.de)

Experimental Psychology and Cognitive Science, Otto-Behaghel-Strasse 10F
D-35394 Giessen, Germany

Markus Knauff (markus.knauff@psychol.uni-giessen.de)

Experimental Psychology and Cognitive Science, Otto-Behaghel-Strasse 10F
D-35394 Giessen, Germany

Abstract

Belief revision is the process by which one alters his or her belief state in the face of contradicting evidence. The contradictions typically arise from a set of statements in which not all propositions can be true at the same time. Despite widespread agreement that people have little difficulty finding such inconsistency, we still lack sufficient knowledge on how people revise their beliefs to resolve the inconsistencies. In the following paper, we report two experiments that were concerned with this research topic. In experiment 1, we explored how familiarity with the content of the statements (familiar vs unfamiliar) affects peoples' belief revision choices. In experiment 2, we investigated whether different ways of questioning (what do you believe "more" vs. "less") affect belief revision. The results show that both factors have a significant effect on peoples' belief revision choices. Our results fit the predictions of the mismatch principle of the mental model theory of human reasoning (e.g. Johnson-Laird, Girotto, & Legrenzi, 2004).

Keywords: Belief revision; mental model theory; mismatch principle; familiarity; task instruction.

Introduction

Belief revision can broadly be defined as the process by which one alters his or her belief state in the face of contradicting evidence. This revision process refers to both abandoning and retaining beliefs to achieve consistency in the belief state. The study of belief revision originates from the areas of artificial intelligence and logic (Gärdenfors, 1988, Harman, 1986), but has recently also received attention in the psychology of reasoning. Although little disagreement exists about the concept of belief revision as being a cognitive process in humans, the mechanisms that underlie this process are fairly unknown. Studies on belief revision have thus far fostered the deductive conditional reasoning paradigm. The two most studied deductive reasoning problems are the modus ponens (MP) and the modus tollens (MT). MP is the logic inference rule of the sort 'if p then q' where conclusion q is inferred from the categorical statement p. Similarly, MT is the logic inference rule of the sort 'if p then q' where conclusion not-p is inferred from the categorical statement not-q. In belief revision tasks, the conclusion for the MP problems is instead not-q and the conclusion for the MT problems is instead p. Both these conclusions cause an inconsistency between the conclusion on the one hand and the conditional and categorical premises on the other hand. The reasoner, then, is typically

instructed to assume that the conclusion is true but that the truth status of the premises is uncertain and decide thereupon whether to revise his or her belief in the conditional or the categorical premise. Elio and Pelletier (1997) performed a set of pioneering experiments on human belief revision and concluded from these that people have a preference to revise their belief in the conditional. This finding was corroborated by other studies, although there exists some debate as to what caused this preference (Dieussaert, Schaecken, De Neys, & d'Ydewale, 2000; Elio, 1997; Elio & Pelletier, 1997; Politzer & Carles, 2001; Revlin, Cate, & Rouss, 2001). Elio and Pelletier, for example, ascribe it to the syntactical form of the conditional. Politzer and Carles argue that it is not due to the conditional nature of the major premise, but instead its status quo of being the major premise; its compound nature makes it more likely to be the source of error. These studies set out a path toward the understanding of how people revise their beliefs. However, the lack of a firm understanding of human belief revision calls for a systematic exploration of factors that might influence the choice of which belief to revise. In the following paper, we report two experiments that were concerned with this research topic. In experiment 1, we explored how familiarity with the content of the statements (familiar vs unfamiliar) affects peoples' belief revision choices. In experiment 2, we investigated whether different ways of questioning (what do you believe "more" vs. "less") affect belief revision. In the following, we first describe the mental model theory of belief revision followed by an outline of a probabilistic approach to human reasoning. Then we report some of our earlier findings after which we proceed with our new experiments. Finally, we draw some general conclusions on the cognitive processes involved in belief revision.

The Mental Model Theory of Belief Revision

According to the mental model theory of reasoning (Johnson-Laird & Byrne, 2002), people construct a set of mental models of the possibilities that the situation embedded in the premises might represent. The key assumption of the theory is that mental models represent only what is true according to the premises, and not what is false, which is called '*The principle of truth*' (Johnson-Laird, et al., 2004). Initially, people construct only one possible mental model, the conjunctive p & q, called the

explicit model. They do however make ‘mental footnotes’ of further implicit models that if fully fleshed out represent the remaining true possibilities, which are “not-p and q” and “not-p and not-q”. Johnson-Laird and his colleagues posit that the first step in reasoning is detecting an inconsistency among a set of premises (Johnson-Laird et al., 2004). According to the principle of *models of consistency*, people search for a mental model that holds a possibility in which all premises are true (Legrenzi, Girotto, & Johnson-Laird, 2003). If they find such a model then the set of premises is judged as consistent, otherwise it is regarded inconsistent.

The model theory introduced a so-called *mismatch principle* to explain how people resolve an inconsistency between statements. According to this principle, the statement that will be revised is that statement, whether it be conditional or categorical, that has a mental model that mismatches and conflicts with the mental model of the contradicting fact (Johnson-Laird, 2006; Johnson-Laird, et al., 2004). With MP problems, the contradicting fact not-q mismatches and conflicts with the mental model of the conditional (p q) and the conditional premise is therefore discarded. With MT problems the contradicting fact p matches the model of the conditional and therefore people would revise the categorical instead. Johnson-Laird and his colleagues (Johnson-Laird, 2006; Johnson-Laird et al., 2004) demonstrate the strength of their theory by harmonizing the results of former studies with the mismatch principle. For example, Elio and Pelletier (1997) found that the belief revision was a function of which counter fact followed the belief set. When it was a negation of the consequent, then subjects tended to reject the conditional and believe the categorical statement more. However, when the counter fact was of the form p then they believed the conditional statement more.

Probability Effects in Reasoning

In more recent years, a probabilistic approach to reasoning has emerged (e.g. Oaksford & Chater, 2001). This approach holds that people interpret a conditional “if p then q” as ‘what is the probability of q given p’ and this seems to affect their reasoning strategy. For example, in a study by Liu, Lo, and Wu (1996), subjects had to decide whether the conclusion followed logically from the premises on valid and invalid inference problems. They found that the higher the perceived sufficiency of the problems, the higher the correct responses. In a follow-up test, they also found a positive relationship between the probability judgment of the conditional and the endorsement of the set of premises. Evans, Handley, & Over (2003) have shown in a similar fashion that people are less likely to endorse a conditional when the antecedent has a low probability. The influence of probability of conditionals on reasoning came to be known as the *conditional probability hypothesis*, first implied by Marcus and Rips (1977) and later further developed by other researchers (e.g. Over, Hadjichristidis, Evans, Handley, & Sloman, 2007). The conditional probability hypothesis is grounded on the conditional subjective probability, $P(q|p)$,

which is known as the *Ramsey test* (Edgington, 1995, as cited in Oberauer & Wilhelm, 2003). This test implies that people judge the conditional the same as the conditional probability. Findings from several studies on reasoning that tested the mental model theory against the conditional probability hypothesis were in favour of the conditional probability hypothesis (Evans, et al., 2003; Oberauer & Wilhelm, 2003; Over et al., 2007).

Former Studies

In two earlier experiments, using MP and MT inference problems, we incorporated the conditional probability hypothesis into the research on belief revision and tested it against the Mismatch Principle of the Mental Model Theory (Wolf & Knauff, 2008). We found support for the conditional probability hypothesis but less convincing support for the mismatch principle. In the first experiment we used high- and low-probability problems. When the conditional expressed a high probability (e.g. *If Carl goes to work, then he takes the car*), people tended to believe it more than the categorical premise. In contrast, when the conditional contained a low probability (e.g. *If Carl goes to work, then he will take a hot air balloon*), people were more inclined to believe the categorical instead. This was the case for both inference problems. These findings extended those of Politzer and Carles (2001), who found that the level of probability of conditionals affected revision choice with MP problems. In a second experiment, we included problem sets that contained a neutral-probability conditional (*If Carla takes a rest, then she puts on classical music*). Here we expected people to revise their beliefs according to the mismatch principle. That is, with MP problems we expected people to retain their belief in the categorical and with MT problems to believe the conditional more. However, this is not what we found. For both MP and MT problems, the conditional was believed more. Several explanations can account for this finding. The conditionals that were meant to express either a high or low probability also did so quite clearly. Therefore it could have been that the participants picked up on this difference and reasoned that these were the only two probability levels. Since events expressed in the neutral-probability problems were still about 50% likely to occur (as measured in a norming study), the participants could have judged them to belong more to the high probability problems and chose accordingly to believe the conditional more for these problems, irrespective of logic inference. Another line of argumentation for the finding, related to the first one, is that it can be argued that people embrace the conditional statement when they can imagine the event it conveys and at the same time judge the event reasonably likely to occur. In this regard, the role of familiarity comes into play.

Current Research Questions

The results from our previous studies motivated us to explore two new research questions. With the first experiment, we wanted to investigate whether familiarity

underlies the lack of support for the mismatch principle in the earlier study. We reason that if indeed, as was found in the former study, people have a tendency to believe the conditional more when it is familiar and likely to occur, then people choose the conditional more with familiar sets of problems. Likewise, if the problem content is unfamiliar, then people should show a tendency to shy away from believing the conditional and instead choose to believe the categorical statement. If, instead, familiarity has no effect then in both cases people would revise their beliefs according to the mismatch principle. With the second experiment we wanted to explore the role of task instruction. Here we hypothesized that the phrasing of questioning has an effect on peoples' belief revision choice.

Experiment 1

Method

Participants Forty students (23 females, 17 males) aged 20 to 37 from different faculties of the University of Giessen participated in this study in exchange for monetary compensation. The participants had no pre-knowledge of basic logic and reasoning.

Materials and design Conditional statements with familiar and unfamiliar content were created for the experiment. The unfamiliar conditionals referred to the domains of microbiology, physics, music instruments, and archaeology. These topics were chosen because they represent real-life domains but at the same time are familiar and meaningful only to people who are experts in the respected fields. The conditionals were checked for correctness by researchers working in the faculties related to the above-mentioned domains. Familiar statements referred to everyday courses of action the participants would be familiar with, either by engaging in these activities themselves or knowing about them irrespective of expertise. Booklets were created consisting of 48 conditionals each printed on a separate page; 24 familiar and 24 unfamiliar (6 of each unfamiliar domain). A group of 20 students rated each conditional for probability and familiarity. The instructions on the first page were as follows, translated from German:

"On each of the following pages, a statement is presented that is uttered by a person. Under each statement is printed a rating scale. Please rate how probable it is, that this person is speaking the truth. 0% means "very unlikely" and 100% means "very likely". Please also mark for each statement, whether you are familiar with the domain presented in the statement. Please rate the statements in the order that they are presented to you. Please do not turn back pages."

Descriptive statistics were run to find eight conditionals rated closest to 50% from both the familiar and unfamiliar conditionals. The eight unfamiliar conditionals consisted of two conditionals of each of the four expertise domains. A Friedman Test showed that the conditionals in each group did not significantly differ from one another with respect to probability, $\chi^2(7) = 8.701, p = .275$ and $\chi^2(7) = 3.498, p =$

.835 for the eight familiar and eight unfamiliar conditionals, respectively. A Wilcoxon's Signed Rank Test showed that the eight familiar conditionals were rated as significantly more familiar than the eight unfamiliar conditionals, $Z = -3.689, p < .001$. Also, the two sets were on average rated equally probable (53.71% and 49.31% for the familiar and unfamiliar conditionals respectively), $Z = -1.065, p = .287$. Thus, the manipulation checks were successful. From both eight familiar and eight unfamiliar conditional statements, MP and MT sets of inference problems were created which resulted in a total of 32 problem sets for the experiment. Below are examples of a familiar MP problem and an unfamiliar MT problem:

If Kerstin visits a friend, Then she brings flowers (p q) Kerstin visits a friend (p) Kerstin does not bring flowers ($\neg q$)	If Paul activates genes, then he mutates specific promoters (p q) Paul does not mutate specific promoters ($\neg q$) Paul activates genes (p)
--	---

A 2 x 2 factorial design was used. The independent variables were Logic Inference (MP vs. MT) and Familiarity (Familiar vs. Unfamiliar). The dependent variables were revision choice and decision time.

Procedure The participants were tested individually in a quiet laboratory room equipped with a computer. Instructions were shown on the computer screen. They were explained that they would repeatedly be presented with sets of three statements, one at a time and that they had to imagine each statement being uttered by a different person. It was further stressed that the truth of the first two statements was uncertain but that the third statement, causing an inconsistency, was certainly true. Their task was then to resolve this inconsistency by choosing which of the two preceding statements they believed more. Four practice items preceded the actual experiment. The participant could read the statements in a self-paced manner and press a space-bar like button to run through the experiment. After the third statement, the participants made their choice by pressing a left or right button. The designation of these two buttons was counterbalanced across subjects. Furthermore, all 32 problem sets were randomized within subjects. After the experiment, participant were set aside and asked whether they were actually familiar with any of the 'unfamiliar' problems. This was not the case.

Results and Discussion

Mean percentages of belief revision choices in the four conditions are depicted in Figure 1 (the bars represent the percentages of what is retained). A repeated measures ANOVA revealed a main effect of Logic Inference, $F(1,39) = 43.877, p < .001, MSE = .118$. The conditional was more often believed with MT problems than with MP problems. No main effect appeared for Familiarity, $F(1,39) = .134, p = .716, MSE = .026$. We did find a Logic Inference x Familiarity interaction effect, $F(1,39) = 6.973, p = .012, MSE = .020$. For MP problems, the belief in the conditional was higher for familiar than for unfamiliar

problems. Conversely, for MT problems the belief in the conditional was greater for unfamiliar than familiar problems.

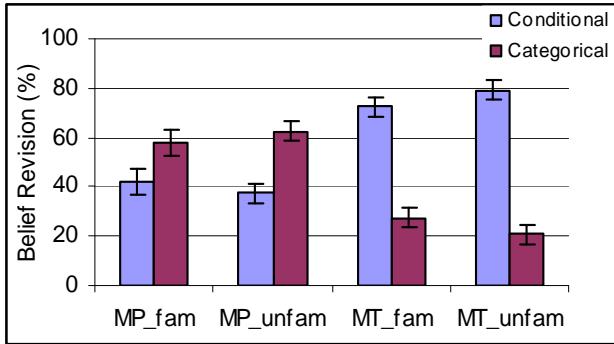


Figure 1: Belief revision choices in percentages. Notes: MP_fam = Modus ponens familiar problems; MT_fam = Modus tollens familiar problems; MP_unfam = Modus ponens unfamiliar problems; MT_fam = Modus tollens unfamiliar problems.

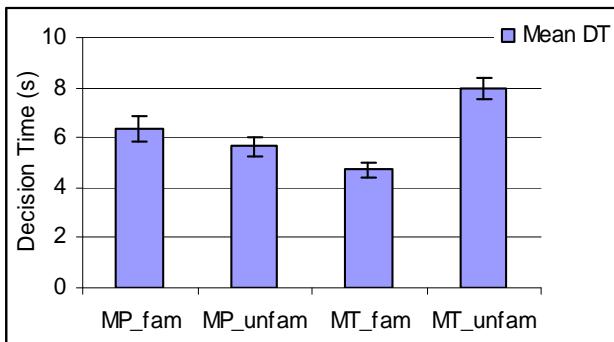


Figure 2: Mean decision times for the 4 conditions

Figure 2 depicts the mean DTs for the 4 conditions. To investigate the effect of Familiarity and Logic Inference on DTs, a Repeated Measures ANOVA was performed. The analysis did not reveal a main effect of Logic, $F(1,39) = .135$, $MSE = 2.030$. There was a main effect of Familiarity, $F(1,39) = 69.214$, $p < .001$, $MSE = 0.949$. In overall, it took subjects significantly longer to decide which statement to believe more ($M = 6.82$ s) with the unfamiliar problems than with the familiar problems ($M = 5.54$ s). However, a significant interaction between Familiarity and Logic Inference accounted for this effect, $F(1,39) = 81.272$, $p < .001$, $MSE = 1.962$. For MP problems, the DTs for the familiar problems were higher than the DTs for the unfamiliar problems. In contrast, with MT problems it took subjects longer to revise their belief with unfamiliar problems than with familiar problems. With respect to the belief revision choices, they were not revised differently with familiar and unfamiliar problems. With both sets of problems, participants chose to believe the categorical statement more with MP problems and the conditional more with MT problems. Moreover, this trend was stronger for unfamiliar problems, as was revealed by the interaction

effect. Most likely, the content of unfamiliar problems have no informational value, with the consequence that people have no other choice than to use mental models to derive at their belief revision choice. These findings strongly support the predictions made by the mismatch principle. Based on this, the failure to find support for the mismatch principle in an earlier study by us (Wolf & Knauff, 2008) was most likely due to the difference in probability levels. Furthermore, the current findings contradict a study by Byrne and Walsh (2005) in which they found a preference to disbelief the categorical with familiar problems and to disbelief the conditional with unfamiliar problems. However, their familiar conditionals were inherently high in probability and their unfamiliar problems referred to an imaginary world (similar to Elio & Pelletier, 1997; Politzer & Carles, 2001). With respect to the DTs, these show the interesting finding that choice of revision does not coincide with ease of revision. In short, the findings of the current experiment clearly demonstrate the strategy of the mismatch principle.

Experiment 2

In the first experiment, we found support for the mismatch principle. However, we used only one kind of task instruction; we asked the participants to choose the statement they believe more. We were interested in finding out whether phrasing of task instruction influences the belief revision process. Task instructions have not been consistent across studies. Participants have been asked to give a degree of belief (Elio, 1997), to select from different options (e.g. negation vs. doubt) (Politzer & Carles, 2001; Elio & Pelletier, 1997), and to choose which statement is true (Byrne & Walsh, 2005). We opine that the influence of task instruction warrants investigation. In the next experiment we asked the participants either to choose which statement they believe more or which they believe less. By this, we wanted to investigate what influence task instructions might have on model construction and in turn belief revision. Of additional interest was to find out which cognitive process is easier to perform for people.

Method

Participants Eighty students (63 female, 17 men) aged 18 to 48 from varying faculties from the University of Giessen participated in the experiment in exchange for a monetary incentive. All students had no basic knowledge of logic and deductive reasoning.

Materials and design The materials for the current experiment were 12 conditionals with a near 50% probability of occurrence taken from two former norming studies, performed by students from the same population. The conditionals did not significantly differ from one another with respect to probability, $\chi^2(11) = 8.751$, $p = .645$. A 2 (Logic Inference: MP vs. MT) by 2 (Task Instruction: 'believe more' vs. 'believe less') between-within subject design was used. The participants were randomly assigned to one of two conditions; one group

received the instruction ‘Which of the first two statements do you believe more?’ (hereafter called the ‘more’ condition) and the other group the instruction ‘Which of the first two statements do you believe less?’ (‘less’ condition).

Procedure The procedure for the current experiment was similar to the first, with the only difference being the task instruction prior to the start of the experiment.

Results and Discussion

Mean percentages of belief revision choices in the four conditions are depicted in Figure 3 (the bars represent the percentages of what is retained). The behavioral data were submitted to a mixed between-within subjects ANOVA. This revealed a highly significant main effect of Logic Inference, $F(1,78) = 41.073, p > .001, MSE = 0.046$. There was no main effect of Task Instruction, $F(1,78) = .905, p = .344, MSE = 0.085$. However, we did obtain a significant Logic Inference x Task Instruction Interaction effect, $F(1,78) = 5.484, p = .022, MSE = 0.046$.

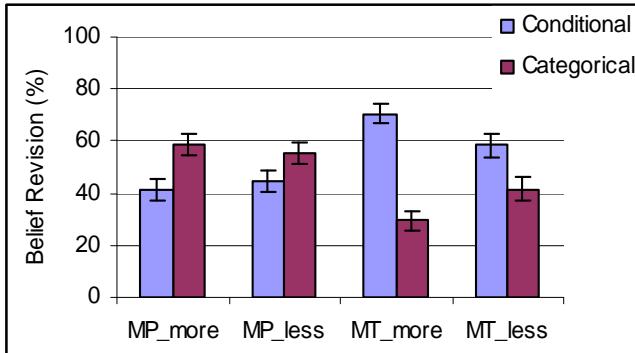


Figure 3: Mean belief revision choices in percentages. Notes: MP_more = ‘more’ condition with modus ponens problems; MT_more = ‘more’ condition with modus tollens problems; MP_less = ‘less’ condition with modus ponens problems; MT_less = ‘less’ condition with modus tollens problems.

In both conditions, the revision choice pattern resonates with the mismatch principle; the categorical was more often chosen with MP problems and the conditional more frequently with MT problems. However, with MP problems the percentage of choosing to believe the conditional more was higher in the ‘less’ condition than in the ‘more’ condition. Conversely, with MT problems, the percentage of choosing to believe the conditional more was higher in the ‘more’ than in the ‘less’ condition. As a result, the mismatch principle reveals itself stronger in the ‘more’ condition. Figure 4 depicts the mean DTs for the 4 conditions. Similar to the first experiment, the analyses on the DT data did not elicit a main effect of Logic Inference, $F(1,78) = .095, p = .759, MSE = 1.973$. Also, no main effect appeared for Task instruction, $F(1, 78) = 1.844, p = .178, MSE = 1.978$. However, we found a significant Logic Inference x Task instruction interaction, $F(1,78) = 4.266, p = .042, MSE = 1.973$. In the ‘more’ condition, faster decision times were

obtained with the MT problems, whereas in the ‘less’ condition subjects made faster decisions with MP problems.

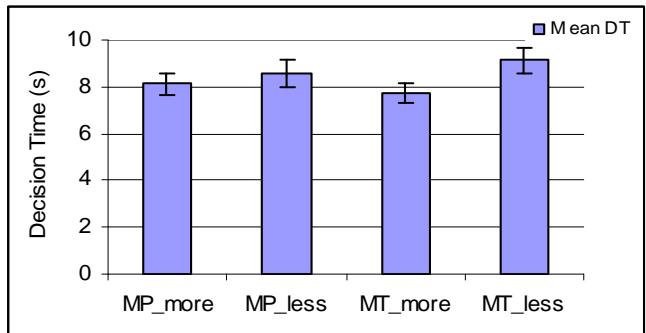


Figure 4: Mean decision times for the 4 conditions

General Discussion

Findings from both experiments demonstrate strong support for the mismatch principle; the categorical is believed more with MP problems and the conditional is believed more with MT problems. With respect to the first experiment, this pattern was more robust for the unfamiliar than for the familiar problems, which only strengthens support for the mismatch principle. Presumably, people use mental models as a guide to resolve an inconsistency between statements that lack clear probability and this becomes an even more preferred strategy when the problems are unfamiliar (but nevertheless known to occur in real-life). The DTs offer further interesting insights; the DTs do not parallel the revision choices. Intuitively you would expect unfamiliar problems to be more difficult than familiar problems. But based on the revision choices you would actually in this case expect the DTs to be faster for unfamiliar problems. However, this was the case only for MP unfamiliar problems. A possible explanation for this could be that although matching of mental models demonstrated to be the prevailing belief revision strategy, psychological processes still exerted its influence, albeit to a lesser degree. To explain, with MT unfamiliar problems, people would prefer to believe the conditional due to a clear match but nevertheless might find it a difficult task to perform because they possess no knowledge of the content. In contrast, the fast DTs for the MT familiar problems presumably result from the increased ease of choosing the conditional due to both its familiarity and the matching of mental models. Thus, the time to reach a belief revision choice depends not only on logic inference, but also on the content of the problems. Furthermore, this non-equivalent pattern between revision choice and the underlying DTs show that DTs on their own can not be taken as an indication of belief revision.

With regard to the second experiment, task instructions revealed a further interesting finding. In both conditions, the participants applied the mismatch principle. However, this seemed more pronounced in the ‘more’ condition than in the ‘less’ condition, where the belief revision choices hovered

around 50%. A possible explanation for this might be that in real life one focuses more readily on what to believe more since that represents the statement one chooses to adopt; such thinking sets out a more straightforward cognitive path. With regard to the DTs, again they don't seem to parallel the pattern of belief revision choices. Since the belief revision choices in the 'less' condition were not as robust as in the 'more' condition, it would seem as if the participants were more careful with their choices. However, although the DTs in the 'less' condition were higher than those in the 'more' condition, this difference did not reach significance. Furthermore, the interaction effect could be explained in a similar fashion as was done for the DTs in the first experiment. Namely, with MT problems, asking people which statement to believe more would presumably be easier than asking people which statement to believe less because in the first case they are directed toward the match between the counterpart and the conditional.

In summary, taking into account the findings of the previous study, when the conditionals do not contain a clear probability that people can use to direct their belief revision process, they use the matching versus mismatching of mental models to resolve an inconsistency between statements. Furthermore, although familiarity appears to influence the strength of using mental models in belief revision, it does not seem to be a factor on its own in belief revision. And finally, the results from the second experiment suggest that directing people to choose the belief they favor instead of disfavor is a more solid approach. Thus, so far we have demonstrated that probability and the use of mental models function as factors that influence belief revision. Which other factors exist is an interesting research question awaiting further investigation.

Acknowledgments

We thank Susann Ruis for her assistance with the data collection. We also express our gratitude to the people from the departments of music, archaeology, physics, and biology for their helpful comments on the material.

References

Byrne, R.M.J., & Walsh, C.R. (2005). Resolving contradictions. In V. Girotto & P.N. Johnson-Laird (Eds.), *The shape of reason: essays in honour of Paolo Legrenzi*. Hove: Psychology Press.

Dieussaert, K., Schaeken, W., Neys, W. De, & d'Ydewalle, G. (2000). Initial belief state as a predictor of belief revision. *Current Psychology of Cognition*, 19, 277-286.

Elio, R. (1997). What to believe when inferences are contradicted: The impact of knowledge type and inference rule. *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 211-216). Mahwah, NJ: Lawrence Erlbaum Associates.

Elio, R., & Pelletier, F.J. (1997). Belief change as propositional update. *Cognitive Science*, 21, 419-460.

Evans, J.St.B.T., Handley, S.J., & Over, D.E. (2003). Conditionals and conditional probability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 321-335.

Gärdenfors, p. (1988). *Knowledge in flux*. Cambridge, MA: MIT Press.

Harman, G. (1996). *Change in view*. Cambridge, MA: MIT Press.

Johnson-Laird, P.N. (2006). Beliefs, heresies, and changes in mind. In P.N. Johnson-Laird (Ed.), *How we reason*. New York: Oxford University Press.

Johnson-Laird, P.N., & Byrne, R.M.J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, 109, 211-228.

Johnson-Laird, P.N., Girotto, V., Legrenzi, P. (2004). Reasoning from inconsistency to consistency. *Psychological Review*, 111, 640-661.

Legrenzi, P., Girotto, V., & Johnson-Laird, P.N. (2003). Models of consistency. *Psychological Science*, 14, 131-137.

Liu, I.-M., Lo, K.-C., & Wu, J.-T. (1996). A probabilistic interpretation of "If-Then". *The Quarterly Journal of Experimental Psychology*, 49A, 828-844.

Marcus, S.L., & Rips, L.J. (1979). Conditional reasoning. *Journal of Verbal Learning and Verbal Behaviour*, 18, 199-223.

Oaksford, M., & Chater, N. (2001). The probabilistic approach to human reasoning. *Trends in Cognitive Science*, 5, 349-357.

Oberauer, K., & Wilhelm, O. (2003). The meaning(s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29, 680-693.

Over, D.E., Hadjichristidis, C., Evans, J.St.B.T., Handley, S.J., & Sloman, S.A. (2007). *The probability of causal conditionals*, 54, 62-97.

Politzer, G., & Carles, L. (2001). Belief Revision and uncertain reasoning. *Thinking and Reasoning*, 7, 217-234.

Revlin, R., Cate, C.L., & Rouss, T.S. (2001). Reasoning counterfactually: Combining and rending. *Memory & Cognition*, 29, 1196-1208.

Wolf, A.G., & Knauff, M. (2008). The strategy behind belief revision: A matter of judging probability or the use of mental models. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 831-836). Austin, TX: Cognitve Science Society.