Cognitive Science 31 (2007) 535–547

What Makes People Revise Their Beliefs Following Contradictory Anecdotal Evidence?: The Role of Systemic Variability and Direct Experience

Henry Markovits[1], Christophe Schmeltzer[2]

[1]*Université du Québec à Montréal, Montreal*
[2]*Université de Toulouse 2 – Le Mirail, France*

Received 14 April 2006; received in revised form 8 January 2007; accepted 10 January 2007

Abstract

The extent to which belief revision is affected by systematic variability and direct experience of a conditional (if A then B) relation was examined in two studies. The first used a computer generated apparatus. This presented two rows of 5 objects. Pressing one of the top objects resulted in one of the bottom objects being lit up. The 139 adult participants were given one of two levels of experience (5 or 15 trials) and one of two types of apparatus. One of these was completely uniform, while the other had an element that randomly alternated in its result. Following the testing of the apparatus, participants were asked to rate their certainty of the action of the middle element, which was always uniform (the AB belief). Then they were told of an observation inconsistent with this belief. Participants were then asked whether they considered the AB belief or the anecdotal observation to be more believable. Results showed that increased experience decreased the tendency to reject the AB belief, when the apparatus did not have any randomness. However, the presence of a single element showing random variation in the system strongly increased rejection of this belief. A second study looked at the effect of a single random element on a mechanical system as well as an electronic system using graphical representations. This confirmed the generality of the effect of randomness on belief revision, and provided support for the effects of embedding a belief into a system of relations. These results provide some insight into the complex factors that determine belief revision.

Keywords: Belief revision; Reasoning; Experience; Randomness

1. Belief revision, self-construction and systemic certainty

A belief can be roughly defined as an item of knowledge that a person considers a true description of the real world (Elio & Pelletier, 1997). Understanding the interplay between

Correspondence should be addressed to Henry Markovits, Département de psychologie, Université du Québec à Montréal, C.P. 8888, Succ Centre-Ville Montréal, Québec H3C 3P8. E-mail: henrymarkovits@gmail.com; markovits.henry@uqam.ca

beliefs and observation is an important component in understanding the way that people construct their models of the world. Particularly interesting are beliefs about relational knowledge, such as causal knowledge. Relational beliefs are hypothetical, since they postulate relations that may be consistent with experience, but can never be totally confirmed. The sum of people's beliefs constitutes a database that determines the state of a person's knowledge, similar to databases that are used in artificial intelligence (Gardenfors, 1988). One of the major problems faced by people and intelligent systems is determining the appropriate reaction when faced with evidence contradicting an existing belief.

Piaget (1974) did some of the earliest work examining this question in a cognitive-developmental perspective. Piaget looked at the way that children reacted to observations that contradicted predictions made on the basis of their internal causal models. Very young children whose causal models were internally contradictory tended to discard their relational beliefs. Older children, with more internally coherent causal models, tended to conserve their beliefs, and discounted the discordant observation.

The notion of coherence has been used in a broader sense to characterize reactions of systems to the introduction of inconsistent information (Thagard & Verbeurgt, 1998; Thagard, 2000) in several fields, including psychology and philosophy. Within this framework, it is assumed that it is possible to calculate the coherence between two specific elements. The level of coherence of a set of elements is then calculated by the sum of the coherence values of all combinations of two elements. One of the goals of any revision process must be to maximize the overall level of coherence of the system. Introducing an observation that is logically contradictory to existing beliefs will result in a partitioning of the extended set consisting of the beliefs and the observation into two sub-sets, one of which will be rejected, in the way that continues to maximize the overall coherence of the system.

Coherence is a formal concept, and in formal systems, such as those used in artificial intelligence, establishing the presence of contradictory information is determined principally by the formal properties of the underlying logic (Gardenfors, 1988). However, applying a coherence framework to human reasoning poses particular difficulties. Studies examining deductive reasoning have not shown a high degree of consistent logical responding (e.g., Cummins, Alksnis, & Rist, 1991; Quinn & Markovits, 1998; Thompson, 1994). Also, people's individual beliefs vary in strength. Understanding the processes of belief revision with humans could help in understanding how complex and ill-determined systems respond to contradictory information.

The first such study by Elio and Pelletier (1997) examined belief revision in the context of deductive reasoning. They used the following paradigm. A subject is presented with a conditional premise (If P then Q), which is described as being true. Then, P is presented as being true, which leads to the conclusion that "Q must be true." Subsequently, it is stated "Q is false." The ensuing contradiction is resolved in two ways. The first involves denying the conditional belief, by asserting that P does not necessarily lead to Q, the second involves denying the minor premise that "P is true" (i.e., the observation). Elio and Pelletier's (1997) results indicate that adults strongly tend to revise the conditional belief, with both familiar and symbolic contents (see also Elio, 1997; Dieussaert, Schaeken, De Neys, & d'Ydewalle, 2000; Revlin, Cate, & Rouss, 2001). Subsequently, Politzer and Carles (2001) observed that the tendency to reject the conditional belief increases when the initial strength of people's belief in the conditional relation decreases.

These studies provide some useful insights into the way that people cope with belief revision. They suggest that people may more easily abandon a relational belief than a contradictory observation, and that this tendency increases as the degree of belief in the conditional relation decreases. However, as was pointed out by Elio and Pelletier (1997), the method used in these studies does not provide any insight into how people reason with 'real, as opposed to verbally presented, beliefs. They suggested that one way of looking at self-constructed beliefs would be to use an artificial micro-world in order to examine what might underlie at least some of the processes of active belief construction.

In a recent study (Schmeltzer & Markovits, 2005), we used just such a micro-world to examine the effects of experience on belief revision. Participants were given a computer generated apparatus that had a row of 5 buttons on top, and a corresponding row of 5 'containers' on the bottom. They were allowed to test the apparatus without constraint. The key causal relation that was examined depended on the fact that pressing the top middle button always lit up a bottom middle container, referred to as the AB belief. In addition, the first four of the five buttons produced similarly consistent effects; pressing one of these buttons always lit up the container directly below. The only exception was the left-most button. The first time that this button was pressed, the container directly below lit up. The second time, the third container lit up. Subsequently, pressing on the left-most button systematically alternated between lighting up the first container and lighting up the third container. After testing the apparatus to their own satisfaction, participants were asked their level of certainty about the AB belief, which was uniformly close to 100%. They were then told that someone had used the apparatus yesterday and that when the middle button was pressed, the middle container did not light up. Participants were asked to choose between revising the AB belief, and rejecting the anecdotal observation. Results showed that participants who had tested the middle button more often revised the AB belief less often. In a second condition, participants were simply shown the apparatus and told that the middle button always lit up the middle container. In this case, 80% of participants revised the AB belief, compared to levels of around 45% when they were allowed direct experience with the apparatus. There was also an intriguing relation between belief revision and experience with the left-most (systematically alternating) button. Participants who tested this button more often, showed a stronger tendency to retain the AB belief, independently of their experience with the third button.

2. Study 1

The results of this initial study support the idea that direct experience makes a relational belief more resistant to subsequent revision, compared to one that is simply stated to be true. However, since the method used allowed participants to make their own decisions about how many times to test each component of the apparatus, the key relation between quantity of experience and belief revision remains uncertain. We thus redesigned the apparatus to provide a more controlled environment in order to examine two hypotheses. The first concerns the effects of degree of experience. Our previous results suggested that the quantity of experience with the AB belief makes this more resistant to revision, despite the uniformly high level of belief of all participants. This suggests that while people construct a high level of belief quite rapidly, the resistance of this belief to revision continues to increase with experience. In order to examine this systematically, we redesigned the apparatus so that all the elements produced uniform results. Participants were allowed two levels of experience with the apparatus, either 5 or 15 trials per element. We predicted first, that in both conditions, participants would have a uniformly high level of certainty about the AB relation and second, the tendency to revise the AB belief should be least in the 15 trials condition.

The second hypothesis was suggested by the strong relationship observed by Smeltzer and Markovits (2005), between experience with the systematically alternating left-most button and revision of the AB belief. Participants who explored this button more often rejected the AB belief less often, independently of experience with the third button. One explanation of this is to suppose that since all the other buttons showed a uniform effect, the behavior of the left-most button might be initially interpreted as random. Increasing experience with this button would allow participants to better understand that its behavior was in fact completely systematic. Thus we hypothesized that the presence of random variability in a system would in itself increase the tendency to revise a relational belief. In order to examine this, we constructed another version of our basic apparatus. In this, the four right-most buttons consistently lit up the four containers directly below them. But, repeatedly pressing the left-most button produced a random sequence of alternation between the fifth container and the third container. We predicted that the tendency to revise the AB belief concerning the action of the middle button (which was uniformly consistent in both versions) would be greater when the apparatus contained one element showing random variation.

2.1. Method

2.1.1. Participants

A total of 139 university students (66 females, 73 males; average age = 25 years, 10 months) were randomly assigned to one of four groups. All participants were students at the Université du Québec à Montréal, and were volunteers.

2.1.2. Materials

A computer program was created using Microsoft Visual Basic. The basic condition which, will be called *uniform* functioned as follows (see Fig. 1 for a diagram of the basic apparatus).

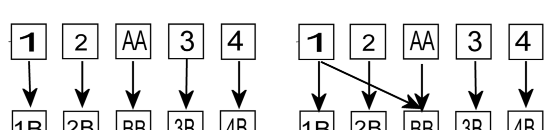

Fig. 1. Basic view of the computer program used for the experiment. In the uniform condition, clicking on any of the five boxes on the top resulted in the box directly below turning red for a short period. The random alternation condition was identical, except that clicking on the left-hand box randomly alternated between lighting up the box directly below and the BB box.

The opening screen asked for demographic information. Following this, a second screen showed a row of 5 grey buttons placed in one-to-one correspondence to a lower row of grey containers. The middle button in the top row was labeled AA, while the middle container in the bottom row was labeled BB. None of the other buttons or containers was labeled. On the top of the screen appeared the numbers 5 or 15, which indicated the number of trials allowed for each button. After an explanation by the experimenter, this number appeared on the top right button of the screen. When this button was clicked on using a computer mouse, the container directly below it turned red for about 1 second, and the number on the button was decreased by one. Clicking on any other button produced no effect. After the first button was fully tested, the number of trials appeared on the second button on the right. Clicking on this button resulted in the container directly below turning red for a short period, while the number on the button was decreased by one. This process was repeated for each of the other buttons, for which the result was always that the container directly below turned red. Following this, a third computer screen appeared. This showed the apparatus, with the following question (translated from the original French): "Suppose that someone clicks on AA, according to you, what is the probability that BB will light up (between 0 and 100)." An empty box appeared to the right of the question, and the participants used the keyboard to input their answer. After this, the question was replaced by a box in which the following description appeared: "John is a student who used this program last night. He clicked on AA, but he says that BB did not light up. According to you, which of these two statements is the more believable?" Directly opposite this box, appeared two other boxes, which were close to each other. The top one contained the statement: "It is not true that John clicked on AA." The bottom one contained the statement: "It is not always true that if one clicks on AA, then BB will light up." The participant could choose one of these by clicking on a space in the box, beside the statement that they wished to choose.

A second condition was also created which will be called the *random alternation* condition. This was identical to the *uniform* condition, with one exception. The first four, right most, buttons consistently lit up the container directly below. However, clicking on the left-most button alternated between lighting up the container directly below, which we refer to as B5, and the third container, BB. The sequence of alternation was a computer generated random sequence, which was applied to both the 5 and 15 trial conditions. The first 5 trials of the sequence were: B5, BB, BB, BB, BB. The subsequent trials in the 15 trials condition were: B5, B5, BB, B5, BB, BB, BB, BB, B5, B5.

2.1.3. Procedure

All participants were interviewed individually. The experimenter explained the functioning of the apparatus. Following this, participants followed the on-screen procedures with no subsequent interaction with the experimenter.

2.2. Results

The belief that clicking on the AA button will always light up the BB container is referred to as the AB belief. We first examined the relations between participants' certainty about the AB belief, their tendency to reject this belief, and the experimental conditions. Table 1 presents the mean level of certainty of the AB belief, and the mean number of revisions of the AB belief, as a function of condition (*uniform*, *random alternation*) and number of trials (5, 15).

An initial analysis examined the degree of certainty of the AB belief. An Anova was performed with Degree of certainty as the dependent variable and Condition and Number of trials as independent variables. This indicated a significant main effect of Condition, $F(1, 135) = 10.41$, $p < 0.001$. Level of certainty of the AB belief was higher in the *uniform* condition (M = 99.2) than in the *random alternation* condition (M = 91.5). There was no effect of Number of trials on degree of certainty.

A second analysis examined revision of the AB belief. A log-linear analysis was performed with revision or not of the AB belief as the dependent variable and Condition and Number of trials as independent variables. This indicated a significant main effect of Condition, $X^2(1) = 21.83$, $p < 0.001$ and a significant condition X number of trials interaction, $X^2(1) = 4.11$, $p < 0.05$. Overall, the proportion of participants in the *random alternation* condition who rejected the AB belief was much greater (M = 0.67) than the proportion in the *uniform* condition (M = 0.29). A chisquared analysis was used to examine differences in AB belief revision as a function of Number of trials. For the *uniform* condition, a significantly smaller proportion of participants revised the AB belief after 15 trials (M = 0.18) than after 5 trials (0.40), $X^2(1) = 4.19$, $p < 0.05$. No such difference was found in the *random alternation* condition.

These results indicate that degree of certainty in the AB belief decreases somewhat in the *random alternation* condition. We then analyzed the effect of condition on belief revision while explicitly controlling for degree of certainty. We first examined only the 117 participants who rated the certainty of the AB belief as 100%, and calculated the proportion who rejected the AB belief by condition (see Table 2). Inspection of Table 2 shows that the pattern of belief revision is very similar to that observed with the full set of participants. We then performed a log-linear analysis with revision or not of the AB belief as the dependent variable and Condition and Number of trials as independent variables. This indicated a significant main effect of Condition, $X^2(1) = 13.64$, $p < 0.001$ and a significant condition X number of trials interaction, $X^2(1) = 4.22$, $p < 0.05$. The pattern of results is identical to that obtained with the full sample.

As a final check on the potential effect of degree of certainty, we performed an Anova with the full sample, with AB revision as the dependent variable, and Condition and Number of trials as independent variables, and included Degree of certainty as a co-variate. This analysis showed a significant effect of Degree of certainty, $F(1, 134) = 5.96$, $p < 0.02$. In

Table 1
Mean percent certainty of the AB belief and mean number of belief revisions as a function of condition (uniform, random) and number of trials (5, 15)

Condition	Trials	N	Certainty of the AB belief	Belief revisions
Uniform	5	35	99.4	0.40
	15	34	99.0	0.18
Random	5	35	91.4	0.63
	15	35	91.5	0.71

Table 2
Mean number of belief revisions as a function of condition (uniform, random) and number of trials (5, 15) for participants who gave a maximal rating to the AB belief

Condition	Trials	N	Belief revisions
Uniform	5	34	0.38
	15	30	0.17
Random	5	26	0.54
	15	27	0.67

addition, there were significant main effects for Condition, $F(1, 134) = 17.27$, $p < 0.001$ and a significant condition X number of trials interaction, $F(1, 134) = 4.17$, $p < 0.05$. Thus, when Degree of certainty is explicitly entered into the full analysis, its effect makes no difference to the pattern of results linking belief revision to Context and Number of Trials.

2.3. Discussion

This study presented participants with an artificial system, which allowed them to directly experience and subsequently construct an internal belief about one of the elements of the system (the AB belief). Participants were then given anecdotal evidence that contradicted a direct inference that could be made on the basis of this belief.

Our first hypothesis was that increased experience with the AB belief should decrease the tendency to revise this belief when the entire apparatus produced consistent effects (the *uniform* condition). This was indeed confirmed. The proportion of participants who revised the AB belief was twice as high among those who were given five trials compared to participants who were given 15 trials. When asked to rate their initial expectations about the AB belief, both groups had equally high levels of certainty. However, when pressed by a conflicting report, direct experience had a large effect on belief revision, as predicted.

Our second hypothesis was that the presence of a single element showing random behavior would increase the tendency to reject the AB belief. This was also confirmed. The clearest indicator of the strength of this effect is the fact that 71% of participants who were given 15 trials in the *random alternation* condition rejected the AB belief; compared to only 18% of those were given the same number of trials on the *uniform* condition. Even with only 5 trials, there was a significant increase in belief revision between the *uniform* and the *random alternation* conditions (from 40–63%, Fisher's exact test, $p < 0.05$).

Introducing a random element also produced a small reduction in participants' level of certainty about the AB belief. The average degree of belief was around 91% for the 70 participants in the *random alternation* condition compared to the level of 99% in the *uniform* condition. Thus, the presence of a random element in the apparatus was sufficient to induce a 10% decrease in participants' confidence in their own observations of the AB relation. This did not account for the effect of randomness on belief revision, since the latter remained equally strong when degree of certainty was controlled.

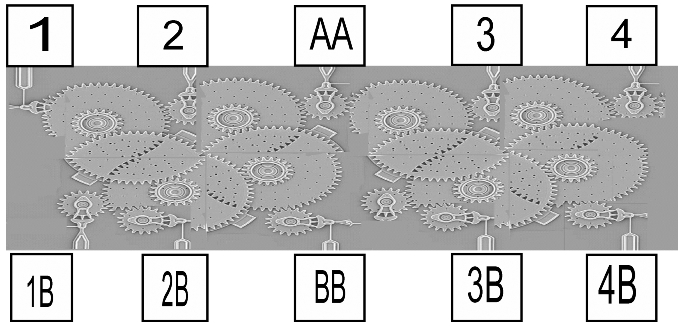

Fig. 2. Graphical representation of the mechanical device using marbles used in Study 2.

3. Study 2

The results of this first study concerning the effect of experience on belief revision are quite clear. However, observed effect of randomness, which is independent of experience, might be related to idiosyncratic beliefs about electronic systems. The aim of the second study was to generalize the effects of the presence of a random element in the context of a mechanical system for which the ordinary rules of causality should apply. Since, we did not have the technical ability to construct a working model of such a system, we decided to use a graphical representation of an apparatus presented on the sheet of paper as a substitute. A special effort was made to present the functioning of the apparatus in a way that was designed to maximize participants' attention to the workings of each component. Specifically, we constructed a picture of a device containing a row of containers on the top, and a second row on the bottom (see Fig. 2). The two sets of containers were linked by a complex set of gears and tubes. This was described to participants as a mechanical device in which marbles that were placed in one of the top containers would end up in one of the bottom containers. Participants were told that a group of students had tested this device for an entire day. The results of the testing for the first four, right-most, containers were indicated clearly on the paper. This showed that for each of these, a marble placed in the top container would always fall into the bottom container directly below it. The results of the testing of the fifth, left-most, container were hidden by a piece of paper. The device and the results of the testing of the first four containers were described to the participant by the experimenter. The experimenter then asked participants to predict the outcome of the testing of fifth container. Participants were then directed to look at the results of the testing of the fifth container, and then to answer questions about the strength of the AB relation, and the belief revision question which were presented on a second sheet of paper. Half of the participants were given a description that indicated that putting a marble into the fifth container always resulted in the marble falling into the corresponding container on the bottom. The other half were given a description that stated that a marble put into the fifth container would sometimes fall into the third container on the bottom and sometimes into the fifth container on the bottom, in a way that was totally unpredictable. In order to allow a more direct comparison with the results of the previous study, a second set of materials was prepared that described and showed an electronic device identical to that used in the first study.

3.1. Method

3.1.1. Participants

A total of 164 university students (101 females, 63 males; average age = 23 years, 3 months) were randomly assigned to one of four groups. All participants were students at the Université du Québec à Montréal, and were volunteers.

3.1.2. Materials

An initial two page booklet was prepared. On the top of the first page was a space for participants to indicate their age and sex. Directly below this, the following paragraph appeared:

"Below you will see the image of a mechanical device that works with marbles. It has two rows of five containers. When a marble is put into one of the containers on the top row, it falls into one of the containers in the bottom row. This device was tested for an entire day by some students."

Directly below this, was an image of the device (see Fig. 2), which showed all the containers with a complex pattern of gears between them. The containers on the top were labeled 1, 2, AA, 3, and 4; while the containers on the bottom were labeled 1B, 2B, BB, 3B, and 4B. Directly below each of the bottom containers was a rectangle, which briefly described the results of the testing. For container 4, this stated "each time a marble was put into 4, it always fell into 4B." For container 3, this stated "each time a marble was put into 3, it always fell into 3B." For container AA, this stated "each time a marble was put into AA, it always fell into BB." For container 2, this stated "each time a marble was put into 2, it always fell into 2B." A Post-it sticker hid the rectangle directly below container 1B. The text inside this rectangle stated "each time a marble was put into 1, the marble sometimes fell into 1B, sometimes into BB, in a totally unpredictable manner."

On the top of the second page of the booklet, appeared the following questions. The first question asked was "Suppose that you try the device and that you put a marble into AA. According to you, is it true that the marble will fall into BB." There was space for the participant to indicate either yes or no. Directly following this, the participant was asked to indicate their degree of certainty, on a scale between 0 and 100% in increments of 10%. Below these questions, appeared the following: "John is a student who used the device last night. He claims that he put a marble into AA, but that the marble did not fall into BB. According to you, which of the following statements is the most believable?" Directly below this were presented the following two statements: "It is not always true that if on clicks on AA, then BB will light up." and "It is not true that John clicked on AA." This was referred to as the *random alternation* condition.

A second booklet was prepared that was identical to the first, with one exception. The description of the outcome of the testing of the first container stated that "each time a marble was put into 1, the marble always fell into 1B." This was referred to as the *uniform* condition.

Two further booklets were also prepared. These were identical to the initial two with one major difference. The device that was represented was described as an electronic device for which one had to click on a top container in order to light up a bottom container, in this case electronic circuitry was shown instead of gears in the depiction of the device.

3.1.3. Procedure

The experimenter interviewed participants individually and showed them the top page of the booklet, read the paragraph explaining the device, and read the first four (visible) descriptions of the outcome of the testing of the four right-most containers. The participant was then asked to predict the outcome of the testing of the fifth container, and then was directed by the experimenter to uncover the fifth rectangle in order to determine the outcome of the testing of container 1. Following this, the participant was directed to turn the page and to respond to questions on the second page with no further intervention from the experimenter.

3.2. Results

All of the participants expected that using the final container would lead to the same one-to-one pattern as had been observed on the initial four containers. We then examined the relations between participants' certainty about the AB belief, their tendency to reject this belief, and the experimental conditions. Table 3 presents the mean level of certainty of the AB belief, and the mean number of participants who found that rejecting this belief was the more believable choice, as a function of condition (*uniform*, *random alternation*) and type of device (*mechanical*, *electronic*).

An initial analysis was performed on the degree of certainty of the AB belief. An ANOVA was performed with Degree of certainty as the dependent variable and Condition and Number of trials as independent variables. This indicated only a significant main effect of Condition, $F(1, 160) = 4.85$, $p < 0.05$. Level of certainty of the AB belief was higher in the *uniform* condition (M = 85.7) than in the *random alternation* condition (M = 78.9). A second analysis examined revision of the AB belief. A log-linear analysis was performed with revision or not of the AB belief as the dependent variable and Condition and Type of device as independent variables. This indicated only a significant main effect of Condition, $X^2(1) = 8.66$, $p < 0.01$. Overall, the proportion of participants in the *random alternation* condition who rejected the AB belief was greater (M = 0.70) than the proportion in the *uniform* condition (M = 0.47).

In order to separate the effects of belief revision from those related to the degree of certainty, we performed an ANOVA with AB revision as the dependent variable, and Condition and Type of device as independent variables, and included Degree of certainty as a covariate. This analysis showed a significant effect of Degree of certainty, $F(1, 159) = 24.84$, $p < 0.001$. In addition, there was a significant main effect for Condition, $F(1, 159) = 5.17$, $p < 0.05$. Thus,

Table 3
Mean percent certainty of the AB belief and mean number of belief revisions as a function of condition (uniform, random) and type of device (mechanical, electronic)

Condition	Device	N	Certainty of the AB belief	Belief revisions
Uniform	Mechanical	42	86.7	0.48
	Electronic	43	84.9	0.47
Random	Mechanical	40	75.8	0.73
	Electronic	39	82.1	0.67

when Degree of certainty is explicitly entered into the analysis, its effect makes no difference to the pattern of results linking belief revision to Condition.

3.3. Discussion

The results of this second study allow us to conclude that the effects of a random element on belief revision generalize to a system where the underlying causal relationships are mechanical, as opposed to electronic. The pattern of the effects obtained in the study is remarkably similar to those of the initial study. Specifically, the presence of a random element induced a relatively small decrease in the degree of certainty accorded to the AB belief, while resulting in much greater tendency to reject the AB belief, one that was independent of the effect on degree of certainty. These effects were the same for both the mechanical system and the electronic system.

It is also of interest to compare the results of the initial study with those of the present study, particularly with the electronic device. Overall, participants who were given a graphical representation of the electronic device had generally lower levels of certainty in the AB belief even in the *uniform* condition (87%) compared to participants with even the lowest level of direct experience (99%). These participants also tended to revise the AB belief with somewhat greater frequency (47%) than the latter (40%). In fact, the profile of the participants in this experiment resembles most that of the participants in the 5 trials condition in the first study.

Comparing these results to those obtained by Schmeltzer and Markovits (2005) and the results of the first study reveals an interesting pattern. When participants in the experiment of Schmeltzer and Markovits (2005) were shown the electronic apparatus and simply told that the AB relation was always true, the degree of belief revision was very high (around 80%). The degree of belief revision in the present case, when participants were told that all five containers showed consistent effects was substantially lower (47%). Thus, as would indeed be predicted by a coherence framework, the tendency to revise a conditional belief decreases when participants have additional information that suggests that the other elements in the system also are consistent in their effects. The results of the second study, which used a graphical representation but provided strong anecdotal evidence of the consistency of all five elements led to a profile of certainty and belief revision that is quite close to that observed in the 5-trial condition in the first study. Finally, 15 trials of direct experience with all five elements leads to an 18% degree of level revision, indicating that strong experience is more effective than strong anecdotal evidence in creating a belief that is resistant to revision.

4. Conclusions

The results of these studies allow us to modulate Elio and Pelletier's (1997) conclusion that people are generally willing to revise a conditional belief, in some very informative ways. The first factor concerns the effect of embedding a conditional belief in a system of consistent relations. When a conditional belief is embedded in such a system, the tendency to revise this belief when faced with contradictory evidence diminishes largely when people also have evidence that the rest of the system functions consistently. This is readily explicable by the coherence framework presented by Thagard and Verbeurgt (1998). This framework considers that the overall coherence of a system is given by the sum of the measured coherence between all possible combinations of two elements in the system. Thus, knowledge that the other elements are also consistent in their functioning will give a much higher level of coherence than that provided by direct knowledge about the AB relation. Although knowledge about the other elements of the system does not in fact affect the degree of belief in the AB relation, it does increase the costs to the overall level of coherence of rejecting the AB belief.

Another important dimension concerns the effects of direct experience on belief revision. The results of these studies clearly indicate that increasing the quantity of direct experience with an entirely consistent system results in a large decrease in the tendency to revise a strongly held belief when faced with contradictory evidence. Clearly, increased experience increases the strength of both the conditional belief and the strength of belief in the other relationships in the system. This would result in a very strong increase in the internal coherence of the system, which would in turn greatly increase the cost of rejecting the AB belief.

Although this basic effect can be easily integrated into a coherence framework, the lack of any effects of experience on participants' explicit certainty about the AB belief is somewhat puzzling. It appears that subjective certainty about a belief can best be considered as a multilayered construct. Clearly, as Politzer and Carles (2001) have shown, variable degrees of subjective certainty results in corresponding differences in belief revision when faced with contradictory evidence. However, the results of our current study, and those of the previous one, indicate that a very high level of subjective certainty can lead to very different levels of belief revision, which are then determined by the degree of direct experience consistent with the belief. In fact, it could be argued that once subjective certainty is very high (which occurs with relatively little experience), belief revision can be seen as a measure of the full extent to which a person adheres to a given belief.

The third factor that we have identified concerns the effects of randomness. The present results clearly show that the presence of a single random element in a system that is either electronic or mechanical in nature has two effects. Firstly, there is a relatively small decrease (about 10%) in participants' subjective certainty about the AB belief. One straightforward explanation of this effect might be that participants are modulating their prediction of the outcome of the AB relation by a general calculation of the overall consistency of the whole system. Since only one of the five elements showed any degree of inconsistency in all of these studies, a simple calculation would situate the overall level of inconsistency at around 10%, which corresponds nicely to the observed decrements in certainty.

In contrast, the effect of a single random element on belief revision is much greater than that on certainty. For example, participants who were given 15 trials with the electronic system showed a fivefold increase in belief revision with randomness. In fact, the level of belief revision shown with a single random element in all four conditions in these two studies (which varies between 63% and 73%) is fairly close to the level of 80% observed when participants are given anecdotal evidence only about the AB relation. This implies that the presence of a random element is sufficient to undo the effect of embedding a relational belief within a consistent system of relations.

What this analysis suggests is that the main effect of randomness is the uncoupling of information concerning an individual belief from the characteristics of other relations in the system. The accompanying increase in belief revision would be due more to a loss in overall coherence than to any decrement in people's direct certainty about the belief. This implies that people are particularly sensitive to the existence of randomness in a system, and that its presence is sufficient to undo much of the effect of embedding a relational belief into a system. This in turn suggests that an evaluation of the overall consistency of the apparatus as a system must be included in any calculation of coherence, over and above its effects on the levels of consistency of individual elements.

The results of this study provide some useful insights into the way that people revise relational beliefs that are embedded into a system involving multiple relations, and into the effects of direct experience and the presence of randomness. This analysis supports the idea that belief revision can be understood in a coherence framework (Thagard & Verbeurgt, 1998), but also suggests the usefulness of a system wide calculation that looks at whether individual components can be seen to share some degree of common characteristics as a component of any calculation of coherence.

Acknowledgments

Preparation of this manuscript was supported by a grant from the Natural Sciences and Engineering Research Council of Canada (CRSNG) to the first author.

References

Cummins, D. D., Lubart, T., Alksnis, O., & Rist, R. (1991). Conditional reasoning and causation. *Memory and Cognition, 19*, 274–282.

Dieussaert, K., Schaeken, W., De Neys, W., & d'Ydewalle, G. (2000). Intitial belief state as a predictor of belief revision. *Current Psychology of Cognition, 19*, 277–288.

Elio, R. (1997). What to believe when inferences are contradicted. In M. Shafto and P. Langley (Eds.), *Proceedings of the 19th Conference of the Cognitive Science Society* (pp. 211–216). Hillsdale, NJ: Erlbaum.

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

Gärdenfors, P. (1988). *Knowledge in flux*. Cambridge: Cambridge University Press.

Piaget, J. (1974). *Recherches sur la contradiction: 1. Les différentes formes de la contradiction* [Studies on contradiction: 1. The different forms of contradiction], Paris: Presses Universitaires de France.

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

Quinn, S., & Markovits, H. (1998). Conditional reasoning, causality, and the structure of semantic memory: Strength of association as a predictive factor for content effects. *Cognition, 68*, B93–B101.

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

Schmeltzer, C., & Markovits, H. (2005). Belief revision, self-construction and systemic certainty. *Current Psychology Letters, 17*(3).

Thagard, P. (2000). Probabilistic networks and explanatory coherence. *Cognitive Science Quarterly, 1*, 91–114.

Thagard, P., & Verbeurgt, K. (1998). Coherence as constraint satisfaction. *Cognitive Science, 22*, 1–24.

Thompson, V. A. (1994). Interpretational factors in conditional reasoning. *Memory and Cognition, 22*, 742–758.