

Order Effects in Human Belief Revision

Hongbin Wang (Hongbin.Wang@uth.tmc.edu)

Jiajie Zhang (Jiajie.Zhang@uth.tmc.edu)

Todd R. Johnson (Todd.R.Johnson@uth.tmc.edu)

Department of Health Informatics

University of Texas - Houston Health Science Center

7000 Fannin, Suite 600

Houston, TX 77030

Abstract

The order effect, a phenomenon in which the final belief is significantly affected by the temporal order of information presentation, is a robust empirical finding in human belief revision. This paper investigates how order effects occur on the basis that human belief has a coherence foundation and a probability/confidence distinction. Both the experimental results and the UEcho modeling suggest that confidence play an important role in human belief revision. Order effects in human belief revision occur where confidence is low and disappear when confidence increases. UEcho provides a computational model of human belief revision and order effects

Introduction

Uncertainty is a specter that haunts human beings. On the one hand, the environment around us is full of uncertainties, ranging from simple situations like tossing a coin, to much more complex ones like natural weather patterns and world peace. On the other hand, uncertainty inevitably arises in one's mental processes as well. People, as John Locke put it, are most of the time condemned to live in the twilight of uncertainty rather than the midday sun of certainty.

While pursuing *absolute knowledge* is the ultimate goal of human life, uncertainty is a pervasive fact with which life must be prepared to contend. It has long been believed, through epistemological inquiries, that forming and maintaining various *beliefs* (e.g., scientific theories and religions) is one way, if not the only way, that human beings try to overcome uncertainty and embrace the ultimate goal of human life.

It is generally agreed that one constantly conducts *belief revision* – a process in which one revises one's beliefs in the light of new information, with a goal to maintain a reasonably consistent and up-to-date belief system. The ultimate reason for belief revision is that belief, unlike knowledge, can be wrong. Since the environment changes dynamically, as time passes new pieces of relevant information constantly become available. These new pieces of information may suggest that one's current beliefs are incomplete, inappropriate, incorrect, inconsistent, or out-of-date. How to sufficiently and efficiently incorporate the new information into the current belief system to reflect the changed reality is a big challenge that one faces. It is of great philosophical and psychological interest to investigate whether one is able to achieve such a goal and what the underlying regularities are.

Psychological investigations of human belief revision have revealed an important finding -- *order effects* (e.g., Hogarth & Einhorn, 1992; Schlottmann & Anderson, 1995; Zhang, Johnson, & Wang, 1997). Generally speaking, the order effect refers to the phenomenon that the temporal order in which information is presented affects the final judgment of an event. Undoubtedly, the temporal order of incoming evidence often carries some important information about the true meaning of an event. However, robust order effects have been found even in situations where the temporal order of incoming evidence seems not meaningful. It is these cases that make the order effect a very interesting phenomenon.

This paper aims to investigate how order effects occur in human belief revision, both empirically and computationally. It consists of four sections. In the first section, some previous studies on human belief and belief revision, uncertainty, and order effects are briefly reviewed. Then a psychological experiment and its UEcho modeling (see Wang, Johnson, Zhang; 1998; Wang, 1998) are reported in the next two sections. The final section provides general discussions and conclusions.

Human Belief Revision

Modern philosophical analysis regards belief as a sort of mental disposition that guides behaviors (Quine & Ullian, 1978). Belief is often contrasted with knowledge in the sense that belief is an inferior substitute of knowledge since believing a proposition does not entail that the proposition is true while knowing it does (see Ackermann, 1972).

There are two main views regarding how an unconvinced belief could be justified (e.g., Gärdenfors, 1990). According to the *foundations* approach, a rational individual derives beliefs from reasons for these beliefs. In other words, a belief is justified if and only if it possesses some satisfactory and "hard" underlying reasons. The *coherence* approach, in contrast, maintains that a belief may be held independent of its supporting reasons. An individual holds a belief as long as it logically coheres with the individual's other beliefs. Therefore, coherent beliefs can mutually justify each other, no belief is more fundamental than another.

How beliefs are justified has a direct implication on how beliefs should be revised when new information becomes available. Based on the foundations view, one should give up those beliefs that lose their underlying reasons and ac-

quire new beliefs that become well supported. An example is the *Truth Maintenance System* developed by Doyle (1979). In contrast, the coherence view emphasizes consistency and conservatism. Therefore, in belief revision one should retain as many of one's beliefs as possible while accommodating any new evidence. In other words, as long as the coherence of the resulting state is maintained, the change to the current state should be minimal. The so-called AGM theory of belief revision (Alchourron, Gardenfors and Makinson, 1985; Gardenfors, 1990) is one well-known example of the coherence approach.

The coherence approach to human belief revision is generally preferred (see Gardenfors, 1990; Thagard, 1989). It has been argued that the foundational approach involves excessive computational cost. It is intellectually very costly to keep track of the reasons of beliefs. Moreover, it has been shown that the foundational approach conflicts with observed human behavior. For example, the *belief preservation effect* (e.g., Ross & Lepper, 1980) suggests that people are reluctant to give up some beliefs even when the original evidential bases of these beliefs are completely destroyed.

Uncertainty

The ultimate reason for human belief and its revision is uncertainty, which is inevitable at all levels of human-environment interactions.

It is well agreed that there are two general types of uncertainty (see Walley, 1991). On the one hand, due to various information deficiencies, one may not completely determine the truth of a proposition. However, the average proportion of that proposition being true in a long run can be precisely specified. The indeterminacy involved in this case is called *uncertainty*. Consider the tossing of an ordinary (fair) coin. One is unable to predict which side the coin will land facing upwards. Thus one is uncertain about the outcome. However, if the coin is tossed many times, one knows that approximately 50% of the times it will land with head (or tail) facing upwards.

On the other hand, in some cases, neither can one completely determine the truth of a proposition nor can one precisely specify the average proportion of that proposition being true in a long run. This type of uncertainty – the indeterminacy of the average behavior – is usually called *imprecision*. Consider a guessing game, in which one is to guess which one of the two boxes contains a ball. Each time the ball is put into one box by a magician, and the outcome is completely random and without any predictability even in a long run. In this case, if one wants to make a judgment that the ball is in a specific box, one is in a situation that is both imprecise and uncertain. Interestingly, if one only wants to say that the ball is either in box 1 or box 2 but not to specify which one, one is in a situation that is imprecise but certain.

The distinction between imprecision and uncertainty is such a fundamental issue that it causes the “holy war” in the field of uncertainty management. On the one hand, *probability theory* (along with Bayes' Theorem for belief revision), the best-established formal method for belief management, has long been criticized that it has difficulty to

handle imprecision. It has been suggested that while a probability number is sufficient to summarize the uncertainty dimension, a confidence measure is needed to handle the imprecision dimension, with a high confidence measure representing precise and a low confidence measure representing imprecise (see Almond, 1995). On the other hand, *fuzzy sets* and the *possibility theory* (see Zadeh, 1978) often deal with imprecision but not uncertainty. The *theory of belief functions* (see Shafer, 1976) deals with both imprecision and uncertainty. Along with Dumpster's rule for evidence combination, it thus provides a more complete picture of formal belief management.

Order Effects

How people actually manage belief is empirically studied by psychology. Quite interestingly, the dominant findings in empirical research suggest that observed human behaviors are not consistent with the predictions of normative postulates, such as formal logic and Bayesian probability rules. A large number of studies have demonstrated that people often systematically deviate from these normative postulates. With the assumption that these normative postulates prescribe how a reasonable individual should behave, these systematic deviations are often labeled as *cognitive illusions, biases, or fallacies* (e.g., Kahneman, Slovic, & Tversky, 1982). Several well-known biases are: *base rate fallacy, conjunction fallacy, overconfidence, and belief bias* (for a recent review, see Kahneman & Tversky, 1996).

The order effect in human belief revision is yet another robust finding in empirical research (see Hogarth & Einhorn, 1992). By the similar standard, the order effect should also be called a bias since the normative postulates, in particular Bayes' Theorem, have no room for it – it simply violates commutativity. However, as many researchers have already pointed out, calling it a bias is nothing more than giving it a label, which provides no help to understand how and why the order effect occurs.

Several theories have been proposed to explain order effects (see Wang, 1998 for a review). Viewed from a broader angle all these theories fall into the same category. That is, they all view human belief revision as a serial process of incorporating new evidence. Therefore, the primary focus is always on how to assess the impact of new evidence. This impact is usually represented in terms of adjustment weights assigned to each new evidence item. However, there is no general agreement regarding what is important to determine the weights.

Miller and Campbell (1959) link belief revision and memory and argue that order effects in belief revision represent order effects in memory. Due to memory decay, previous evidence items get weighted less as time goes. Later studies show that this view is problematic since direct comparisons suggest that beliefs are largely independent of recall of evidence items (e.g., Anderson & Hubert, 1963).

The serial integration model (e.g., Schlottmann & Anderson, 1995), proposed in the framework of information integration theory (Anderson, 1981), link order effects with attention. It is argued that people pay less attention to successive items of evidence due to attention decrement. Attention decrement results in different weights being as-

signed to different evidence items, which in turn results in order effects. Unfortunately, this model does not specify what underlying factors affect the attention decrement, thus leaving the model in a non-testable status.

Hogarth and Einhorn (1992) suggested an anchoring and adjustment model to explain order effects. According to this model, belief revision is a sequential anchoring-and-adjustment process in which people adjust the current belief (the anchor) on the basis of how strongly new information confirms or disconfirms this belief. In addition, the adjustment weight is a function of both the anchor and the new evidence. More specifically, when the impact of the new evidence is smaller than the reference point, the adjustment weight is proportional to the anchor. And when the impact of the evidence is larger than the reference point, the adjustment weight is inversely proportional to the anchor. It is this kind of contrast effect that results in order effects. The model further adopts two parameters (α and β) to regulate the weight assignment process. It claims that the two parameters represent people's sensitivity toward negative and positive evidence, respectively. In particular, the model argues that some individuals tend to view negative (or positive) evidence more seriously than others. Therefore, in terms of the underlying factors that regulate the weight assignment, the model actually points to unidentifiable individual differences.

Summary

The above review reveals two important findings in human belief and its revision. First, human belief has a coherence foundation. A belief can survive without solid foundational evidence. Beliefs hold each other as a coherent system. As a result, in terms of belief revision, new positive evidence does not necessarily reinforce a belief, and new negative evidence does not necessarily discredit a belief. As long as one's beliefs are consistent and coherent with each other, they are justified. Second, human belief has a multi-component structure. The uncertainty/imprecision (or probability/confidence) distinction indicates that uncertainty has many faces, as does human belief. A single probability number simply cannot capture all the important aspects of a belief. A confidence component is necessary.

Unfortunately, previous theories of order effects hardly take these findings into consideration. These theories all attempt to explain order effects by a weight assignment mechanism that weighs members of the evidence sequence differentially. However, they encounter great difficulties in fully explaining why weights have to be assigned in a particular way at a particular time. Consequently, in some cases, one or more task characteristics are particularly emphasized (e.g., memory decay, or attention decrement), which of course often only account for a fragment of the order effect. In some other cases, arbitrary parameters are adopted in the weight assignment to summarize unidentifiable sources.

The probability/confidence distinction suggests that the impact of the new evidence cannot be fully understood without the nature of the current beliefs being sufficiently appreciated. More specifically, the confidence component of a belief, mainly determined by the amount of previous

experience, represents how easily this belief can be revised. A belief with no previous experience has very low confidence and is easiest to change. Consider a coin tossing example. With no experience about the fairness of the coin, getting 5 heads in a row easily makes one to suspect its fairness. On the other hand, a belief established by significant previous experience is committed with a high confidence level and thus is hard to change. For example, after tossing the coin 1000 times and observing 500 heads and 500 tails, 5 heads in a row will hardly make one to suspect the fairness of the coin.

In the context of order effects, the above analysis implies that the order effect pattern may change with different levels of experience. The rationale is as follows. As one keeps interacting with the environment, one gains more and more experience. As a result, beliefs are gradually tuned to the statistical structure of the environment (see Anderson, 1990). And confidence increases as more experience is gained. Both factors will make one react to any new evidence more realistically rather than over-react or under-react. Since over-reacting and under-reacting are the fundamental causes of the order effect, this is to say that when one gains more and more experience about the environment, the order effect in belief revision tends to diminish and disappear.

The experiment presented next is designed to test this hypothesis.

Experiment

Design and Procedure

A modified version of the CIC (Combat Information Center, Towne, 1995; see also Wang, Johnson, & Zhang, 1998; Zhang, Johnson, & Wang, 1997) computer simulator was used as the task domain. In this CIC task, the goal of the participant, acting as a commanding officer of a naval ship, was to maintain the safety of his/her own ship by collecting two pieces of information sequentially about an aircraft in the radar area and accurately identifying its intention.

One piece of information was about the route (R), which indicates the target is either on or off a commercial air route. The other piece of information was self-identification (Self-ID), which indicates the target's response after being warned. The participant had to report the degree of belief (on a 0-100 scale) that the target is friendly before any evidence and after each piece of evidence. Finally, the participant was forced to make a two-alternative (i.e., friendly or hostile) judgment about the identity of the target. After the decision was made, the participant could request the true identity of the target if available.

The experiment adopted a 3x(4) factorial design. The between-subject independent variable was the ratio of total friendly targets to total hostile targets in the training samples. The ratio was 1:1 (equal number of friendly and hostile targets), 3:1 (friendly targets are three times as frequent as hostile targets), or 1:3 (hostile targets are three times as frequent as friendly targets). The purpose of this factor was to create environments with different statistical structures

and test if participants could gradually tune their beliefs to capture these structures.

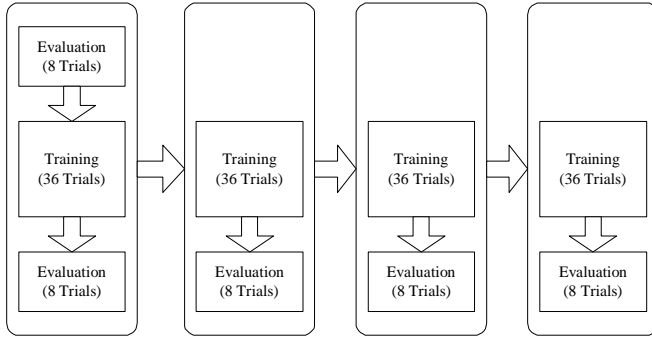


Figure 1. The experimental design

The experiment attempted to investigate how the patterns of order effects changed with training. The training was organized in four blocks, which is the within-subject variable (see Figure 1). Five evaluation blocks were inserted in the process to provide a way to easily evaluate the pattern changes of order effects. The major difference between training trials and evaluation trials is that no true identity feedback was provided at the end of each evaluation trial.

Each evaluation block had eight evaluation trials in it. The eight evaluation trials were constructed in the following way. There were two pieces of evidence (Route and SelfID), each of which had two possible values (“on” and “off” for Route, and “friendly response” and “no response” for SelfID), so there were 4 kinds of trials; and each piece of evidence could be collected first (this was different from the training trial, in which the evidence collection order was a between-subject variable), thus we had a total of eight trials. Participants were instructed to summarize their training experience in order to perform these evaluation trials.

Each training block consisted of 36 trials. The trial distribution is dependent on the friendly-hostile ratio and is shown in Table 1. Since a value of “on” for Route and a value of “friendly response” for SelfID are regarded as positive evidence for a friendly target, they are represented by “+”s in Table 1. Similarly the opposite values are represented by “-”s.

Table 1. The trial distribution


Route	SelfID	1:1		3:1		1:3	
		F	H	F	H	F	H
+	+	8	2	12	1	4	3
+	-	4	4	6	2	2	6
-	+	4	4	6	2	2	6
-	-	2	8	3	4	1	12
Total		18	18	27	9	9	27

140 undergraduate students participated in the experiment. They were randomly assigned to the 3 friendly-hostile-ratio treatment groups. The trials in each block were completely randomized for each participant.

Results

The five evaluation blocks, distributed in the critical positions in the training, are the focus of our analysis. In addition, for the purpose of easily examining order effects, only the data from the two critical evidence sequences (“+-“ and “-+”) are reported. The results are shown in Figure 2.

Three major findings are described. First, the effect of friendly-to-hostile ratio is evident. While the average initial belief judgment (i.e., before any evidence) tends to increase with training in the 3:1 group (56.7, 59.0, 67.1, 68.6, from block1 to block4, respectively), it tends to decrease with training in the 1:3 group (47.8, 43.0, 41.3, 40.2, from block1 to block4, respectively). Note that it is largely unchanged with training in the 1:1 group (50.0, 50.0, 49.5, 53.3, from block1 to block4, respectively). This pattern of result suggests that the initial belief judgments were gradually tuned



to more closely reflect the built-in friendly-to-hostile ratios.

Figure 2. The belief revision patterns in all three friendly-to-hostile ratio conditions. The initial evaluation block before any training is labeled as block0, which also combines data from all three ratio groups. The evaluation blocks after each training block are labeled as block1 to block4, respectively

Second, the belief revision patterns change significantly across the whole training session. A recency effect is evident in block0 (the final judgment is 41.9 for the “+-“ sequence vs 50.8 for the “-+” sequence), and this recency effect tends to disappear in the later blocks. More specifically, recency effects exist in block1 and disappear in block3 and block4 in all three ratio groups. This pattern is consistent with our prediction that order effects diminish and disappear with training.

Finally, it is interesting to note that the areas inside the fish-tail like order effect patterns tend to become systematically smaller as the training progresses. A rough estimation procedure is created to quantitatively measure the areas and it shows that area decrement is statistically significant in the 1:1 and 3:1 groups, though not in the 1:3 group. This pattern of area decrement indicates that participants fluctuated less in their belief judgments as more experience was

gained, which further suggests that participants tended to be less sensitive to new evidence as confidence goes up.

In summary, the experiment results show that the recency effect disappeared as more training trials were performed. The disappearance of the recency effect suggests that instead of over-reacting in the light of new evidence, participants made more proper and more realistic reactions. As suggested previously, as more experience was acquired during training, the statistical tuning led participants to make more confident belief judgments, which eliminated over-reaction.

UEcho, first proposed in Wang, Johnson, and Zhang (1998) as a model of belief evaluation in abduction, is further developed to model the experiment results.

A UEcho Model

UEcho is based on Echo, which is a connectionist implementation of the Theory of Explanatory Coherence (TEC), proposed by Thagard (1989, 1992) as a model of human abductive reason. Echo takes a coherence view of belief evaluation as its foundation. According to Echo, a belief should be accepted if it is coherent with other beliefs, and rejected if it is incoherent with other beliefs. By quantitatively defining (explanatory) coherence, an Echo system pursues highest coherence by considering all related beliefs in a holistic manner. When the system converges, the most believable hypothesis set will defeat any competitors and pop out.

Although Echo has gained much empirical support, they have serious limitations (e.g., Wang, Johnson, & Zhang 1998): (1) Echo does not handle sequential belief revision; (2) Echo does not learn from experience.; and (3) Echo does not distinguish confidence and probability. All these limitations cast doubt on Echo as a general model of human belief revision.

Wang, Johnson, and Zhang (1998) proposed UEcho as an extension of Echo to address the first two problems. It has been shown that UEcho is able to model order effects. UEcho is further extended here to embed the probability/confidence distinction. By doing so, we hope that UEcho, as a coherence-based model of belief evaluation, is able to model both the existence and the disappearance of order effects.

UEcho maintains that the activation of a node determines acceptability, thus representing the probability component of a belief. UEcho adopts three mechanisms to add a confidence dimension to the system.

All three mechanisms try to tune critical parameters based on previous experience. The first parameter is the parameter of skepticism θ . In Echo θ represents the decay rate in the activation updating – The higher θ is, the faster does the node activation decay. Confidence cures skepticism. Gradually tuning down θ , based on experience, is a natural way to represent confidence. The second mechanism has to do with the parameter α and β in the anchoring and adjustment model. As mentioned earlier, α and β represents one's sensitivity toward “negative” and “positive” evidence. Although the anchoring and adjustment model attributes the sensitivity to some unidentifiable factors of personality, the two parameters are functionally close related to confidence in the sense that as confidence goes up, the sensitivity to

new evidence goes down. Incorporating and gradually tuning α and β represents another aspect of confidence management. Finally, UEcho extends Echo's parameter of *data excitation*. In Echo, data excitation is used to represent that observed data nodes have independent support of their own. The hypothesis nodes in fact has an associate data excitation of always 0. By generalizing this parameter to hypothesis nodes, UEcho enables hypothesis nodes to learn and remember their activation values, thus to gradually gain self-support (or dis-support) of their own.

For a more detailed description of these tuning mechanisms, please see Wang (1998).

The exact same design and procedure was used to train a UEcho network, and the corresponding simulation results are shown in Figure 3.

The modeling results match the experiment results remarkably well. First, the gradual separation of the curves of the three ratio groups nicely reflects the statistical tuning toward the built-in environmental friendly-to-hostile ratios. From block1 to block4, the average belief judgment is 50.0, 50.0, 50.0, 50.0 for the 1:1 group, 58.8, 63.9, 67.2, 69.0 for the 3:1 group, and 41.4, 36.7, 33.6, 32.7 for the 1:3 group, respectively. Second, the order effect pattern change is evident. A recency effect is significant in block0 (23.3 for “+” vs 76.7 for “-”). The magnitude of the recency effect, measured as the difference between the final judgment in “+” and the final judgment in “-”, decreases significantly from block1 to block4. More specifically, they are 8.4, 3.6, 2.0, 0.8 for the 1:1 condition, 9.1, 3.5, 1.6, 1.3 for the 3:1 condition, and 11.1, 6.1, 3.6, 2.5 for the 1:3 condition, respectively. Finally, the areas inside the fish-tail shapes become systematically smaller with training as well, indicating the fluctuation in belief revision tends to be smaller as the training progresses.




Figure 3. The belief revision patterns in all three friendly-to-hostile ratio conditions, based on the UEcho simulation.

In summary, by embedding the probability/confidence distinction, UEcho is capable of capturing the changes of order effect patterns at different experience levels. The close match between the simulation results and the experi-

mental results in the decrement and disappearance of order effects with the increase of experience supports UEcho as a model of coherence-based and complex human belief revision.

Discussions and Conclusions

Human belief and human belief revision are ubiquitous in everyday life and scientific discovery. The order effect, a phenomenon in which the final belief is significantly affected by the temporal order of evidence is a robust empirical finding in human belief revision. The order effect is generally regarded as a manifestation of human biases and an indication of human irrationality. It is the attempt of this paper to study how the order effect occurs.

Previous research leads to the conclusion that human belief has a coherence foundation and consists of multiple components. Such a conclusion motivates and guides both the experimental study and the computational modeling work described in the paper. Both the experimental results and the UEcho modeling results show that order effects in belief revision exist at the early stage of training when the confidence level is low and they tend to diminish and disappear later when the confidence increases.

It is interesting to further speculate how the UEcho modeling results could tell us the possible rational basis of order effects. First of all, the fact that UEcho, constructed based on rational postulates and intended to prescribe what people should do, naturally shows order effects (when the confidence level is low) convincingly “debiases” order effects. Second, the existence of order effects has ecological implications. UEcho reveals that order effects appear when the relevant experience is scarce, and order effects disappear when the relevant experience becomes rich. When the relevant experience is rich, one has confident expectations, which eliminate the need to over-react. When the relevant experience is scarce, one has to sufficiently appreciate every single piece of information since its relevance cannot be easily and accurately determined in the first place. In this sense, both the existence and the disappearance of the order effect are rational.

It should be noted that this study involves only the recency effect. It would be of great importance to explore how it can be tuned or extended to model the primacy effect. Whether it can model the full range of order effects using the same mechanism is a strong test for UEcho as a general model of human belief revision.

Acknowledgements

This work is funded by Office of Naval Research Grant No. N00014-95-1-0241.

References

Ackermann, R. (1972). *Belief and knowledge*. Garden City, NY: Doubleday & Company.

Alchourron, C., Gardenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction functions and their associated revision functions. *Journal of Symbolic Logic*, 50, 510-530.

Almond, R. G. (1995). *Graphical belief modeling*. London: Chapman & Hall.

Anderson, N. H. (1981). *Foundations of information integration theory*. New York, NY: Academic Press.

Anderson, N. H., & Hubert, S. (1963). Effects of concomitant verbal recall on order effects in personality impression formation. *Journal of Verbal Learning and Verbal Behavior*, 2, 379-391.

Doyle, J. (1979). A truth maintenance system. *Artificial Intelligence*, 12, 231-272.

Gardenfors, P. (1990). The dynamics of belief systems: Foundations vs coherence theories. *Revue Internationale de Philosophie*, 172, 24-46.

Hogarth, R.M. & Einhorn, H.J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24, 1-55.

Josephson, J.R., & Josephson, S.G. (1994). *Abductive Inference: Computation, Philosophy, Technology*. Cambridge, NY: Cambridge University Press.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and Biases*. Cambridge, NY: Cambridge University Press.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 103(3), 582-591.

Miller, N., & Campbell, D. T. (1959). Recency and primacy in persuasion as a function of the timing of speeches and measurement. *Journal of Abnormal and Social Psychology*, 59, 1-9.

Ross, L., & Lepper, M. R. (1980). The perseverance of beliefs: Empirical and normative considerations. In R. A. Shweder (Ed.), *Fallible judgment in behavioral research: New directions for methodology of social and behavioral science, Vol. 4*. San Francisco, CA: Jossey-Bass.

Schlottmann A., & Anderson, N. H. (1995). Belief revision in children: Serial judgment in social cognition and decision-making domains. *Journal of Experimental Psychology: LMC*, 21(5), 1349-1364.

Shafer, G. (1976). *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press.

Thagard, P. (1989). Explanatory Coherence. *Behavioral and Brain Sciences*, 12(3), 435-502.

Thagard, P. (1992). *Conceptual revolutions*. Princeton, NJ: Princeton University Press.

Towne, D. (1995). CIC: Tactical Decision Making (Version 2.0). Unpublished manuscript. Behavioral Technology Laboratories, University of Southern California.

Walley, P. (1991). *Statistical reasoning with imprecise probabilities*. New York, NY: Chapman and Hall.

Wang, H. (1998). *Order effects in human belief revision*. Ph.D. Dissertation, The Ohio State University.

Wang, H., Johnson, T. R., & Zhang, J., (1998). UECHO: A model of uncertainty management in human abductive reasoning. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Zadeh, L. A. (1978). Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1(1), 3-28.

Zhang, J., Johnson, T.R., & Wang, H. (1997). The relation between order effects and frequency learning in tactical decision making. *Thinking and Reasoning*, 4(2), 123-145.