

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Order Effects and Frequency Learning in Belief Updating

Permalink

<https://escholarship.org/uc/item/9xn6v6vr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 18(0)

Authors

Zhang, Jiajie

Johnson, Todd R.

Wang, Hongbin

Publication Date

1996

Peer reviewed

Order Effects and Frequency Learning in Belief Updating

Jiajie Zhang¹, Todd R. Johnson², Hongbin Wang¹

Department of Psychology¹

Division of Medical Informatics, Department of Pathology²

Center for Cognitive Science^{1,2}

The Ohio State University

Columbus, Ohio 43210

zhang.52@osu.edu, johnson.25@osu.edu,

wang.190@osu.edu

Abstract

This paper examines order effects and frequency learning in belief updating. We present an experiment that tests for the existence of order effects for actual decisions during frequency learning and for belief evaluations after frequency learning in a realistic tactical decision making task. The experiment revealed that (a) subjects showed order effects for actual decisions during frequency learning—an effect not reported previously and (b) subjects still showed order effects for belief evaluations even after having correctly learned most of the frequency information. We also present a simulation for the frequency learning behavior and some preliminary results of a simulation for the order effect, and suggest networks for potential combinations of the order effect and frequency learning.

Introduction

A number of experiments indicate that the order in which evidence is presented can affect the strength of a person's belief in hypothesized causes. Different experiments have found that order of evidence can produce no effect, a recency effect, or a primacy effect. Hogarth and Einhorn (1992) proposed an anchoring and adjustment model of belief updating that predicts when these effects will occur based on features of the belief updating task. For example, step-by-step evaluation of beliefs for mixed positive and negative evidence items produces a recency effect: the final evaluation of belief is mainly determined by the last evidence item. However, step-by-step evaluation of beliefs for consistent evidence (all positive or all negative) produces no effect.

A different set of studies and models are concerned with the learning and use of frequency information. When conditional probability and base rates of occurrence are presented explicitly in terms of numeric values, they are very difficult to learn and utilize (see Kahneman, Slovic & Tversky, 1982). However, when they are presented in terms of real events and occurrences, they can often be learned implicitly and used correctly (e.g., Christensen-Szalanski, & Bushyhead, 1981; Medin & Edelson, 1988; for a review, see Hasher & Zacks, 1984). As a result of using real events to present frequency information, many of the well-known biases in human probabilistic reasoning (see Kahneman,

Slovic & Tversky, 1982) disappear. However, one important bias that has yet to be studied in this area is the order effect on belief updating described above.

The studies and models of the order effect are usually separated from those on frequency learning and use. As a result, models of belief updating based on frequency acquisition research cannot account for the complete spectrum of order effects. In fact, one of the most well-known examples, the Rescorla-Wagner model (Rescorla & Wagner, 1972), completely ignores the temporal sequence of information. Since many frequency acquisition tasks involve temporal sequences of information, it is important to consider the joint implications of these two research areas. This paper examines order effects and frequency learning in a common task environment. We present an experiment that tests for the existence of order effects during and after the acquisition of frequency information in a natural setting. The experiment revealed that (a) subjects showed order effects for actual decisions during frequency learning—an effect not reported previously and (b) subjects still showed order effects for belief evaluations even after having correctly learned most of the frequency information. We also present a simulation for the frequency learning behavior and some preliminary results of a simulation for the order effect, and suggest networks for potential combinations of the order effect and frequency learning.

Experiment

The experimental task was implemented on the CIC (Combat Information Center) simulator developed by Towne (1995) for the US Navy. Figure 1 shows a simplified radar display of the CIC simulator. It shows an unknown airplane heading toward the Naval ship which is at the center of the radar display. The captain of the ship can check whether the target is on or off a commercial air route by clicking the route button to display all available routes. He can also send a radio verbal warning to request the target to identify itself by clicking the warning button. The target may or may not respond to the warning. In this experiment, when it responds to the warning, it always identifies itself as a commercial airplane. The target can be either friendly or hostile. The task is to use the information about the air route and the

information about the identity (ID) obtained from the radio warning to identify whether the target is friendly or hostile. The constraint of the task is that the two evidence items (route and ID) can only be obtained sequentially, one at a time. The order can be either route followed by ID or ID followed by route.

The experiment tests three hypotheses. The first hypothesis is about frequency learning. In a given geopolitical environment, there are certain conditional probabilities about whether the target is friendly or hostile given the two evidence items. When subjects are trained on the task many times with a fixed base rate and conditional probabilities, they can implicitly and accurately acquire most of the frequency information. The second hypothesis is about the order effect for belief evaluations. Previous studies show that when frequency information is accurately and implicitly learned in actual events, certain biases such as the base rate fallacy can be eliminated. We test the hypothesis that even if most of the frequency information is acquired implicitly and accurately, the order effect for belief evaluations, a special type of bias, still exists. That is, when the two evidence items are presented one by one in different temporal orders, the final evaluations of hypotheses about the friendliness of the target are different. The third hypothesis is about the order effect for actual decisions. Previous studies showed order effects only for belief evaluations, not for actual decisions. We test the hypothesis that subjects show order effects for actual decisions during frequency learning.




Figure 1. A simplified radar display on the CIC simulator. See text for details.

Method

Subjects. The subjects were 40 undergraduate students in introductory psychology courses at The Ohio State University who participated in the experiment for course credit.

Design & Procedure. There were two evidence items: route and ID. Route indicates whether the target is on or off a commercial air route. ID indicates whether there is any response from the unknown target to a radio warning issued from the ship. The two evidence items were presented in two

different orders: *Route-ID order* in which route information was collected first, followed by ID information; *ID-Route order* in which ID information was collected first, followed by Route information. After having collected both evidence items, subjects made a forced-choice response indicating whether the unknown target was friendly or hostile. After each response, subjects were given feedback indicating whether the response was correct or incorrect. Each subject performed 50 trials. The conditional probabilities of hostility and friendliness for a given set of evident items are shown in Table 1. For half of the 40 subjects, the two evidence items were always presented in the Route-ID order for all 50 trials; for the other half, in the ID-Route order for all 50 trials. The 50 trials for each subject constituted the learning phase for the acquisition of frequency information.

After 50 trials, each subject was given a written questionnaire requesting belief evaluations about the hostility and friendliness of the unknown target after the presentation of a baseline fact and each of the two evidence items. In the questionnaire, Route was always negative (the plane was not on a commercial route) and ID was always positive (the plane indicated that it was a commercial plane). For half of the 20 subjects receiving each of the two training orders (Route-ID and ID-Route) of the 50 trials, the evaluation order of the evidence items was Route-ID; and for the other half, the evaluation order was ID-Route. For example, an evaluation order of Route-ID is shown in Figure 2. The written questionnaire constituted the evaluation phase for belief evaluations.

Thus, this experiment was a 2x2 between-subject design, with the two learning orders as one factor and the two evaluation orders as another factor.

1. You see a plane which is getting closer to your ship. On a scale from 0 to 100 (with 0 being total disbelief and 100 total belief) please rate your belief in the following hypotheses:
 1. How likely do you think the plane is hostile?
 2. How likely do you think the plane is friendly?
2. After consulting commercial air routes, you discover that the plane is not on a commercial air route. Given this new information, please answer the following questions. Again, express your answer on a scale of 0 to 100 with 0 being total disbelief and 100 being total belief.
 1. How likely do you think the plane is friendly?
 2. How likely do you think the plane is hostile?
3. When you asked the plane to identify itself, the plane identifies itself as a commercial airplane. Given this new information, please answer the following questions. Again, express your answer on a scale of 0 to 100 with 0 being total disbelief and 100 being total belief.
 1. How likely do you think the plane is friendly?
 2. How likely do you think the plane is hostile?

Figure 2. The questionnaire for the Route-ID evaluation order.

Table 1. Probability Distribution of Learning Trials

Route	ID	p(Friendly Route)	p(Friendly ID)	p(Friendly Route, ID)
+	+	0.80	0.80	0.94
+	-	0.80	0.20	0.50
-	+	0.20	0.80	0.50
-	-	0.20	0.20	0.06

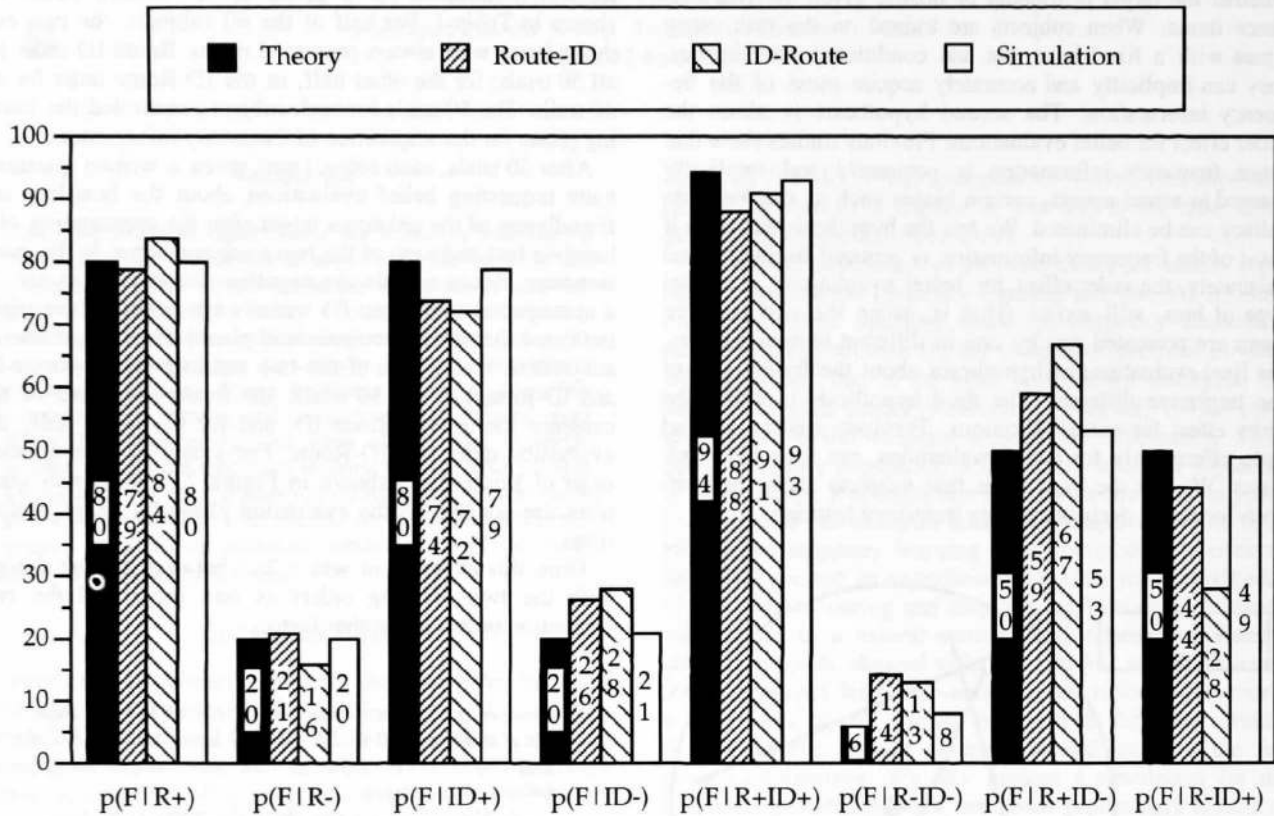


Figure 3. Conditional probabilities from theoretical calculation (Bayes rule), the experiment, and the simulation using the Rescorla-Wagner rule.

Results & Discussion

Frequency Learning. The responses of the 50 trials by each subject were transformed into conditional probabilities, which were then averaged across the 20 subjects for the Route-ID learning order and across the 20 subjects for the ID-Route learning order. The results are shown in Figure 3 under Route-ID and ID-Route.

The conditional probabilities from the experiment were compared with their corresponding theoretical values. For the Route-ID learning order, p(F|ID+), p(F|ID-), and p(F|R-ID-) were significantly different from the theoretical values (smallest $t(19) = 2.82, p < 0.05$); p(F|R+), p(F|R-), p(F|R+ID-), p(F|R-ID+), and p(F|R+ID+) were not signifi-

cantly different from the theoretical values (largest $t(19) = 2.0, p > 0.05$). For the ID-Route order, p(F|ID+), p(F|ID-), p(F|R-ID-), p(F|R+ID-), and p(F|R-ID+) were significantly different from their theoretical values (smallest $t(19) = 2.26, p < 0.05$); p(F|R+), p(F|R-), and p(F|R+ID+) were not significantly different from the theoretical values (largest $t(19) = 1.31, p > 0.20$). These results indicate two findings. First, subjects correctly learned some of the conditional probabilities but not all, and more for the Route-ID order than for the ID-Route order. Second, ID was a less objective evidence item than Route: p(F|ID+) was smaller than its theoretical value but p(F|R+) was not different from its theoretical value. This might be because Route was always objectively observed from the radar display whereas ID was obtained from possibly deceptive radio communication.

The conditional probabilities from the two different learning orders were also compared with each other. The only significant difference was $p(\text{FIR-ID+})$ ($F(1, 38) = 4.52, p < 0.05$). This is clearly a recency order effect: $p(\text{FIR-ID+})$ for the Route-ID order was larger than $p(\text{FIR-ID+})$ for the ID-Route order because the last evidence item in the Route-ID order was positive (ID+) whereas that in the ID-Route order was negative (R-). A similar order effect was also observed for $p(\text{FIR-ID-})$ although the difference between the two learning orders was not statistically significant. This result is an indication that subjects showed order effects for actual decisions during frequency learning.

Belief Evaluation. The results of belief evaluations after the learning phase are shown in Figure 4. For both learning orders, there was a clear order effect: when the two evidence items (positive ID, negative Route) were presented in different temporal orders, the final friendliness evaluations of the unknown target were different. An ANOVA for the final evaluations of friendliness was conducted for the two learning orders and two evaluation orders. The main effect for the two learning orders was significant ($F(1, 38) = 4.32, p < 0.05$), indicating that the learning order Route-ID produced a more hostile evaluation than the learning order ID-Route. The main effect for the two evaluation orders was also significant ($F(1, 38) = 13.41, p < 0.001$), indicating that the evaluation order Route-ID produced a more friendly evaluation than the evaluation order ID-Route. This order effect for evaluations was a recency effect, as predicted by Hogarth and Einhorn's model: the final evaluation of friendliness was determined by the last evidence item. For the Route-ID evaluation order, the last evidence ID was positive, producing a more friendly evaluation. In contrast, for the ID-Route evaluation order, the last evidence Route was negative, producing a more hostile evaluation. The interaction between learning order and evaluation was not significant ($F(1, 38) = 0.53, p = 0.47$).

Simulation: Frequency Learning

The experiment shows that subjects could learn most of the frequency information implicitly and accurately. This section describes a connectionist simulation of frequency learning based on the Rescorla-Wagner rule (Rescorla & Wagner, 1972).

The Rescorla-Wagner rule was initially proposed for classical conditioning in animal associative learning and later extended for human learning (e.g., Gluck & Bower, 1988). It is sensitive to frequencies of observations and conditional probabilities. Let w_{ij} denote the strength of association between observation o_i and hypothesis h_j . If h_j is a correct hypothesis for observation o_i , then the weight change is

$\Delta w_{ij} = \eta o_i (\max - \sum_{i \in o} w_{ij})$, where η is the learning rate, o_i reflects the reliability of the observation, \max is the maximum possible level of associative strength that can be associated with each hypothesis unit h_j and $\sum_{i \in o} w_{ij}$ is the total

associative strength for the hypothesis h_j that connects all observations present on that trial. If the hypothesis h_j is an incorrect hypothesis for observation o_i , then the associative strength between o_i and h_j decreases. The weight change is $\Delta w_{ij} = -\eta o_i \sum_{i \in o} w_{ij}$. After a series of trials, w_{ij} will reflect the conditional probability of the hypothesis h_j given the observation o_i .

Figure 5 shows a simple network used to simulate the experimental task described previously. The four observation units represent the four possible outcomes of the two evidence items: Route positive (R+), Route negative (R-), ID positive (ID+), and ID negative (ID-). The two hypothesis units represent the two possible outcomes of the evaluation: Friendly (F) and Hostile (H).

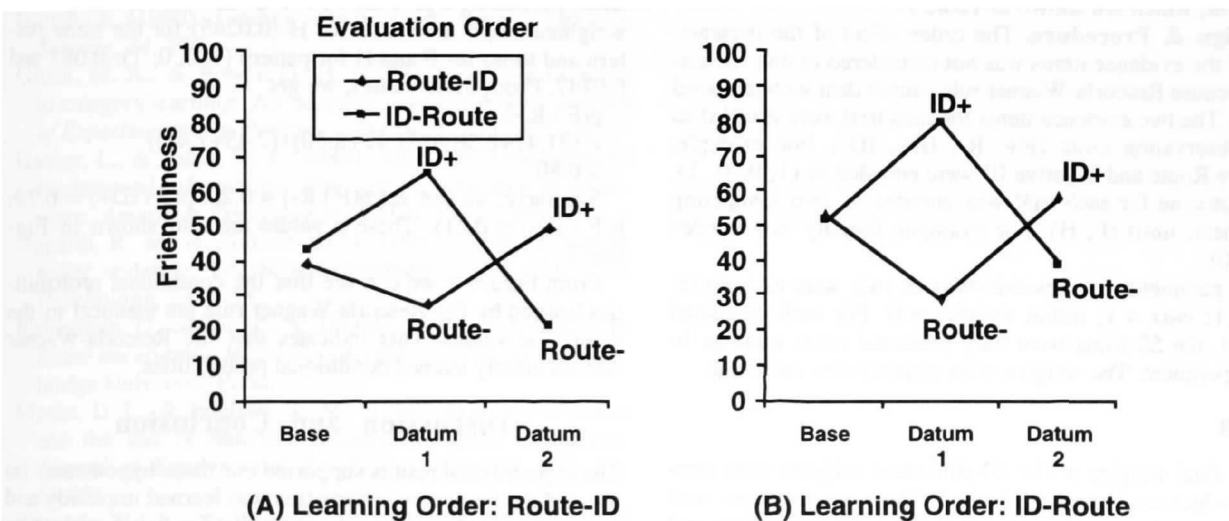


Figure 4. The evaluations of friendliness for the two different evaluation orders under the two different learning orders.

Table 2. The Estimation of Conditional Probabilities for One Evidence Item

input patterns	prob of patterns	output activation		p × output × 100	
R+ R- ID+ ID-	p	o(F)	o(H)	Friendly	Hostile
1 0 1 0	0.34	0.924	0.072	31.42	2.45
1 0 0 1	0.16	0.519	0.467	8.30	7.47
0 1 1 0	0.16	0.487	0.515	7.79	8.24
0 1 0 1	0.34	0.082	0.910	2.79	30.94




Figure 5. A simple network used to simulate the experimental task. The Rescorla-Wagner learning rule was used. The weights are the average of the weights learned by the network across 20 simulated subjects with 50 trials for each simulated subject.

Method

Subjects. 20 simulated subjects were trained on the network.

Materials. The 50 trials used in the experiment were the 50 trials for the simulated subjects. The presentations of the 50 trials were randomized for each simulated subject. The conditional probabilities of outcomes given conditions of the two evidence items were identical to those in the experiment, which are shown in Table 1.

Design & Procedure. The order effect of the presentation of the evidence items was not considered in this simulation because Rescorla-Wagner rule cannot deal with temporal orders. The two evidence items for each trial were encoded as four observation units (R+, R-, ID+, ID-). For example, positive Route and negative ID were encoded as (1, 0, 0, 1). The outcome for each trial was encoded as two competing hypothesis units (F, H). For example, friendly was encoded as (1, 0).

The parameters of Rescorla-Wagner rule were as follows: $\eta = 0.1$; $max = 1$; initial weights = 0. For each simulated subject, the 50 trials were only presented once, same as in the experiment. The weights were updated after each trial.

Result

The final weights of the 20 simulated subjects were averaged, which are shown in Figure 5. These weights were used to estimate the conditional probabilities of friendliness and hostility for different combinations of the two evidence items.

The calculation of $p(F | R, ID)$ is straightforward. For example, if we use $o(F)$ and $o(H)$ to indicate the activation values of the friendly and hostile units, then

$$\begin{aligned}
 p(F | R+, ID+) &= o(F)/(o(F)+o(H)) \\
 &= (w_{FR+}+w_{FID+})/((w_{FR+}+w_{FID+})+(w_{HR+}+w_{HID+})) \\
 &= (0.470+0.454)/((0.470+0.454)+(0.024+0.048)) \\
 &= 0.93
 \end{aligned}$$

Similarly, we can get $p(F | R+, ID-) = 0.53$; $p(F | R-, ID+) = 0.49$; $p(F | R-, ID-) = 0.08$. These p values are shown in Figure 3.

To estimate $p(F | R)$ and $p(F | ID)$, we need a different procedure because we cannot use the ratio of activation values or the ratio of weights directly for the following reason. Suppose we have an input vector (R+, R-, ID+, ID-). In an actual input pattern, one of the two Rs is 1 and the other is 0 and one of the two IDs is 1 and the other is 0, e.g., (1, 0, 1, 0). To estimate $p(F | R+)$, for example, we need an input pattern (1, 0, 0, 0), which is a pattern never presented in the training set. Thus, in order to estimate $p(F | R)$ and $p(F | ID)$, we need to consider the distribution of the 50 training patterns, as shown in Table 2. For example, to calculate $p(F | R+)$, we need to consider all the patterns in the 50 training patterns that contain the pattern (1, 0, 0, 0), which are (1, 0, 1, 0) and (1, 0, 0, 1). Given the input (1, 0, 1, 0), the output activation of F is 0.924, which is then weighted by the probability of this pattern (0.34) across the 50 training patterns, producing a product 0.3142. Similarly, we can get the weighted output activation for H (0.0245) for the same pattern and those for F and H for pattern (1, 0, 0, 1): 0.083 and 0.0747. From these values, we get

$$\begin{aligned}
 p(F | R+) &= (31.42+8.30)/((31.42+8.30)+(2.45+7.47)) \\
 &= 0.80
 \end{aligned}$$

Similarly, we can get $p(F | R-) = 0.20$; $p(F | ID+) = 0.79$; $p(F | ID-) = 0.21$. These p values are also shown in Figure 3.

From Figure 3 we can see that the conditional probabilities learned by the Rescorla-Wagner rule are identical to the theoretical values. This indicates that the Rescorla-Wagner rule accurately learned conditional probabilities.

Discussion and Conclusion

The experimental results supported our three hypotheses: (a) most of the frequency information was learned implicitly and accurately, (b) there was an order effect for belief evaluation after frequency learning, and (c) there was also an order effect for actual decisions during frequency learning. The first unique contribution of the present experiment is that it

shows that the order effect could not be eliminated even if most of the frequency information was learned implicitly and accurately. This is in contrast with previous studies which show that when frequency information is learned implicitly and accurately, certain biases such as the base rate fallacy can be eliminated. The second unique contribution of the present experiment is the demonstration of order effects for actual decisions, which have not been reported previously in the literature.

The simulation results on frequency learning with the Rescorla-Wagner rule were identical to the theoretical values and very close to the values learned by human subjects. It indicates that the Rescorla-Wagner rule can account for both theoretical and empirical data on frequency learning. In addition to the Rescorla-Wagner rule, we also tried a simple recurrent net (Elman, 1993) and a reinforcement net (Sutton, 1988) to simulate the order effect. The simulation results show that both types of networks could produce order effects for the order used for training but they could not be generalized to a new order not present in the training. Other types of networks are being examined for the order effect. The high-level objective of our simulation studies is to find a single architecture that can not only learn frequency information but also produce the full range of order effects.

Acknowledgments

This research was supported by Office of Naval Research Grant No. N00014-95-1-0241 and a summer fellowship from the Center for Cognitive Science at The Ohio State University.

References

- Christensen-Szalanski, J. J. J., & Bushyhead, J. B. (1981). Physicians' use of probabilistic information in a real clinical setting. *Journal of Experimental Psychology: Human Perception and Performance*, 7 (4), 928-935.
- Elman, J. (1990). Finding structures in time. *Cognitive Science*, 14, 179.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117 (3), 227-247.
- Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *American Psychologist*, 39(12), 1372-1388.
- Hogarth, R. M. & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24, 1-55.
- Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Medin, D. L., & Edelson, S. M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, 117 (1), 68-85.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical Conditioning II: Current Research and Theory*. New York: Appleton-Century-Crofts.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning*, 3, 9-44.
- Towne, D. (1995). *CIC: Tactical Decision Making (Version 2.0)*. Behavioral Technology Laboratories, University of Southern California.