Modeling the Effects of Conditional Feedback on Learning

Navid Ghaffarzadegan, Thomas R. Stewart

Center for Policy Research, Department of Pubic Administration and Policy,

Rockefeller College of Public Affairs and Policy, University at Albany (SUNY)

Working Paper – Jan 15, 2009

This material is based in part upon work supported by the National Science Foundation under Grant Number SES 0451850. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Modeling the Effects of Conditional Feedback on Learning

Abstract

Most studies of decision threshold learning have assumed full feedback conditions, that is, regardless of the decision made, feedback is always provided. However, for many selection and detection problems the availability of feedback may be conditional on the decision made. For example, in passenger screening for airline security, security officers receive no feedback about the passengers who are not selected for searching. In this paper, we use simulation to investigate how such "conditional feedback" affects decisions. First, based on detection theory, a model of threshold learning based on error correction is proposed and used to examine conditional feedback situations. Simulation results show the model is able to replicate two results in threshold learning: conservative threshold placement in full feedback and threshold overestimation in conditional feedback. Finally, the model is used to replicate some other empirical findings. The results suggest that conditional feedback can be a barrier to learning by producing overconfidence about negative decisions.

Keywords: threshold learning, conditional feedback, detection theory

1. INTRODUCTION

Many management decisions involve elements of selection and detection. Decisions involved in hiring, promotion, and layoffs are only a few of the most obvious examples. In their simplest form, these decisions involve choosing one of two actions. Although more than two alternative actions may be available in many situations, this paper is limited to a discussion of binary choice problems. Decisions are assumed to be based on a continuous score or judgment. For example, a decision to subject an airline passenger to an interview and search is based on the security officers' judgment of degree of suspiciousness. If the judgment exceeds some threshold, then a search is conducted. Otherwise the passenger is allowed to pass unchallenged. Performance in such a task depends on both the ability to make accurate judgments and the ability to use an appropriate decision threshold. In this paper we address the problem of learning the appropriate decision threshold through experience.

Several mathematical models of threshold learning in binary choice problems have been proposed. Models range from error correction models that assume people learn from their errors (Kac 1962, Kac 1969), and Hill Climbing models in which individuals are assumed to search the neighborhood of the current solution in order to find a better answer (Busemeyer and Jae Myung 1992), to more complicated models assuming individuals are ideal learners who can remember all the right and wrong outcomes (Kubovy and Healy's 1977), or reinforcement learning models in which individuals assign a value for each member of a range of possible thresholds, modify those values after each trial, and then randomly choose one of the thresholds based on probabilities determined by the values (Erev et al. 1995, Erev 1998). These models are also applied to resource allocation decision making tasks (e.g., Rahmandad et al. 2008, Rahmandad 2008) and group decision making (e.g., Gunnthorsdottir and Rapoport 2006, Biele et al. 2008).

Although these studies pick different assumptions about how people learn, they share an important assumption about the nature of information feedback. They assume full feedback condition, that is, a decision maker makes a decision and always receives feedback about whether that decision was right or wrong. But full feedback may be the exception outside the laboratory.

Few studies have examined partial or conditional feedback. In conditional feedback, decision makers receive feedback for a specific kind of decisions - usually for positive decisions. For example, a human resources manager will know true job performance of a candidate if she decides to recruit the applicant. A police officer, who decides not to search a suspect, will not know whether or not the suspect possesses illegal substances. This is the same for the cases of admission decisions in universities, strategic decisions in companies, most medical decisions, etc. In all of these situations, and in many other real world conditions, there is dependency between one's decision and whether or not she receives clear feedback about the results (Einhorn and Hogarth 1978, Tindale 1989, Elwin et al. 2007, Dalgleish and Smillie 2006). The typical situation is that decision makers receive feedback about positive decisions (e.g. admitting a candidate, or deciding to search a suspect), but not about results of negative decisions.

Conditional feedback is argued to impose overconfidence. Based on a model, Einhorn and Hogarth (1978) claim that in conditional feedback (they call it partial feedback) people receive less disconfirming information that they would receive in a full feedback condition, so they become trapped in the illusion of validity. Fischer and Budescu (2005) make a similar argument based on a set of experiments. They study learning and development of confidence under different types of feedback and in different base rates. They argue that decisions made in the conditional feedback mode (screening tasks) tend to induce overconfidence, while under full feedback (discrimination tasks) people are able to learn fast and there is a high correspondence between their performance and confidence.

Learning performance is also different under conditional feedback and full feedback. Besides Tindale's (1989) study of 48 trials of decision making in a recruitment task with eight information cues and feedback that shows better individual performance under conditional feedback, other studies that we are aware of claim the opposite argument. Fischer and Budescu's (2005) results show that individuals perform better in full feedback conditions. Further, Elwin et al. (2007) investigate empirically the effects of conditional feedback on decision making. While observing that people underestimate the base rate (the ratio of signals to total observations), they argue that people assume their negative decisions, for which they do not receive feedback, are true. In addition, Stewart et al. (2007) also find that people overestimate their threshold (or underestimate their base rate) under conditional feedback.

Interestingly, the findings about overestimation of the threshold, even in low base rates, are inconsistent with the robust findings under full feedback condition about conservative threshold placement, that is, people pick less extreme thresholds in comparison to optimal thresholds (Green and Swets 1966, Kubovy and Healy 1977, Busemeyer and Jae Myung 1992). Therefore, it is important to develop a model of learning that can predict both behaviors.

The contribution of this study is to advance our understanding of imperfectness of decision making in a series of tasks with conditional feedback. While many scholars have emphasized the negative effects of misperception of delays (Sterman 1989a, Sterman 1989b, Rahmandad 2008), noise in feedback (Bereby-Meyer and Roth 2006), and feedback asymmetry (Denrell and March 2001) on learning, our work builds on few studies of conditional feedback and its effect on learning. This study does not reject other theories, but sheds more light from a new perspective on the problem of barriers to learning.

In the following, based on a brief review of detection theory (section 2), we build a simulation model of full feedback (section 3) and conditional feedback systems (section 4) and examine effects of different ways of coding negative decisions on learning optimal thresholds. Finally, we use the model to replicate data from a published empirical work (section 5).

2. DETECTION FRAMEWORK

Detection theory, also known as the signal detection theory (Macmillan and Creelman 1991, Green and Swets 1966; Swets 1991; Swets et al. 2000), divides events into two categories, the ones we want to detect (positive events) and all others (negative events). For example a police officer wants to detect guilty persons from innocent ones, and a human resources office wants to select capable people from a pool of capable and incapable candidates.

In detection theory, it is assumed that a decision maker makes decisions by comparing a continuous quantity (judgment, test score, observation, etc.) with a threshold. If the stimulus is above the threshold then it is assumed to be a positive event.

Due to environmental uncertainty it is not possible to totally differentiate positive and negative events. In other words, the distributions for positive and negative events overlap. That is to say, in the example of the police officer, although, on average, guilty persons appear more suspicious than innocents, the distributions of guilty and innocent persons overlap (Fig. 1) and whatever the police officer’s judgment is, there is always a possibility of making a wrong decision.

[Insert Fig.1 about here]

Figure 1 assumes normal distributions of suspiciousness for both innocent and guilty people, and that both distributions have the same variance. The distribution in the left depicts the distribution for innocent people, and the one in the right depicts the distribution for guilty people. The distance between the means of the two distributions is labeled d'. The vertical line represents the decision threshold. If, for a particular case, the judgment falls above the threshold then some action is taken, such as stopping or arresting the suspect. Otherwise a different action or no action is taken. Obviously, the portion of the distribution of innocent people that falls above the threshold represents mistakes, i. e., innocent people who are stopped or arrested. Similarly the portion of the distribution of guilty persons that falls below the threshold represents another kind of mistake—guilty people who are not stopped or arrested.

In any binary decision making situation, there are four possible decision outcomes. You can say "yes" and be right or wrong or you can say "no" and be right or wrong. These outcomes are often labeled true and false positives and true and false negatives. As Table.1 shows, a police officer can decide to stop a person (a positive decision) and the person maybe guilty (true positive) or innocent (false positive). The police officer can also decide not to stop the person (a negative decision). And again the person can be guilty (false negative) or innocent (false positive). Thus, there are two kinds of errors: false positives and false negatives.

[Insert Table.1 about here]

An important point is that different thresholds impose different error rates, and as the probability of one error decreases, the probability of the other error increases (see figure 1). For example, if the threshold is increased (moved to the right) false positives will be reduced, but false negatives will increase. Unless the distributions can be moved further apart (increasing d'), it is impossible to simultaneously decrease both errors by changing the threshold.

In this framework, the proportion of positive decisions is called selection rate (e.g. if 50 percent of people are selected for searching, selection rate is 0.5). On the other hand, the proportion of positive events (e.g., the proportion of suspects who are guilty) is called base rate (e.g. if 50 percent of people are guilty, base rate is 0.5).

In addition, it is important to mention that the threshold changes due to environmental pressures from different constituencies (Weaver and Richardson 2006, Hammond 1996), learning over time (Kubovy and Healy's 1977, Busemeyer and Jae Myung 1992, Erev et al. 1995, Erev 1998) or other factors. Based on a given set of values/costs for each cell in Table.1, there is an optimal threshold location. So, the question rises if individuals are able to learn to move their threshold toward the optimal level.

3. FULL FEEDBACK MODEL

Although not always the most effective way to learn (see, Balzer et al. 1989) we consider the role of outcome feedback in threshold learning. Given values for each of the decision outcomes, and normally distributed positive and negative events with known means and standard deviations, the optimal threshold can be calculated, but people do not always learn this threshold. We can assume that a person may require many trials to learn a threshold. In each trial, she will receive information about her performance and will try to correct/change her threshold, in order to increase the performance. For example, a human resources manager will find what are the minimum qualifications of an applicant, (e.g. education and experience) required for doing the job.

In many situations decision is based on judgment. For simplicity we do not consider the judgment process; our model creates a one-dimensional continuous observation (horizontal axis in Figure 1) that is the basis for decision, we consider two main processes in threshold learning: decision making, and adjusting (correcting) threshold. We assume that people use their current threshold as an anchor and adjust it to a new level using information about the payoff from their previous decision (Tversky and Kahneman 1974). Psychologically, it means that a decision maker has a threshold and tries to shift it toward the best threshold through experience. This assumption is consistent with many studies of decision science on anchoring and adjustment (e.g. Epley and Gilovich 2001). Actually, this assumption can be considered close to the formal assumption of individual learning in error correction models (Kac 1962, Kac 1969). Error correction models have been argued as "good first approximation[s] to decision rule[s]" while they do not fit to every effect (Kubovy and Healy's 1977: 427). In following, we model each of the processes for a full feedback condition.

3.1. Decision Making Process

From detection theory, a decision maker has a threshold and makes her decision by comparing the stimuli with the threshold. If the observation is greater than the threshold she judges it as a positive event, otherwise as a negative event. We assume the existence of a single threshold which can be formulated by an if-then-else decision rule:

$$d = 0 \quad \text{if } x < C$$
$$d = 1 \quad \text{if } x \geq C \qquad \text{(eq. 1)}$$

whereby *d* represents a decision, and is 1 for positive decisions and zero for negative decisions. *x* is the observation, and *C* is the decision maker's threshold. Let us show the true state of the world by *Q* which will be either 1, for positive events, or zero, for negative events. By comparison of *d* and *Q* we can find the payoff. The following formula does the same:

$$V_{Q,d} = (1-Q)*(1-d)*V_{tn} + Q*(1-d)*V_{fn} + (1-Q)*d*V_{fp} + Q*d*V_{tp} \qquad \text{(eq. 2)}$$

whereby payoff, $V_{Q,d}$, will be equal to V_{tn}, V_{fn}, V_{tp} and V_{fp}, called values, in true negative, false negative, true positive, and false positive decisions respectively.

For simplicity we assume symmetric values: $V_{tn}= V_{tp}=1$ and $V_{fn}= V_{fp}=-1$. The model is not restricted to these values, and other value structures will be explored in the future if the model proves useful.

3.2. Adjusting threshold

Different learning processes can be assumed in this stage. The most models assumed that people try to increase total payoff by changing the threshold after each trial based on the outcome.

We assume a process of learning from results: as a subject gets information about the value of the previous observation, she judges the payoff shortfall. Payoff shortfall is the difference between the maximum possible payoff for the event that occurs (Q) and the actual payoff obtained for the decision made. We can formulate the process as following:

$$\Delta V = V_{max,Q} - V_{Q,d} \qquad (eq.\ 3)$$

$$V_{max,Q} = V_{tn} + Q^*(V_{tp}-V_{tn}) \qquad (eq.\ 4)$$

whereby ΔV is Payoff shortfall, and $V_{max,Q}$ is the maximum possible payoff , i.e., the maximum value that a person can receive from making a decision about an event (Q), and as we assumed higher payoff values for correct decisions it can be calculated by a linear function of V_{tp} and V_{tn}.

Knowing that we have made a wrong decision ($\Delta V > 0$), the model assumes that the decision threshold will be amended toward the observation. Obviously, one observation can not change the whole assumption and the subject's mental model, but, in fact, it takes time for a person to change her threshold. Considering such a process, we can say:

$$\Delta C = k.(x-C) \qquad (\text{if } \Delta V > 0) \qquad (eq.\ 5)$$

whereby k is the correction ratio and we have $0 < k \leq 1$. The correction ratio can depend on many factors, such as the personal characteristics of the decision maker and her confidence, and the latter can change dynamically in the system.

So far, in addition to the threshold adjustment loop (threshold→change in threshold→ threshold) we have introduced one simple loop that formulates a full feedback system (threshold→decision→ payoff→payoff shortfall→change in threshold→threshold). This feedback forces the subject toward the optimal threshold.

We produce a set of random negative and positive events, consistent with detection theory (*negative events ~ N(0,1)* and *positive events ~ N(d',1)*), and choose randomly from them with a ratio that creates the desired base rate.

[Insert Fig.2 about here]

Now, we can examine simulation results of this simple full feedback system. We conduct 100 runs with the model for different random seeds. Fig.2 illustrates the average of those experiments. As the figure shows, the model converges near the optimal threshold which, in the base rate of 0.5, is equal to d'/2, i.e. 0.5. The initial threshold chosen was -1. In these runs d′ is assumed equal to one. The speed of approaching the optimal threshold depends on the correction ratio.[1]

Simulation runs for the base rates of 0.25 and 0.75 replicate the conservative cutoff-placement phenomenon and shows that based on the assumed model of learning, the model ends with a threshold which is less extreme than the optimal threshold. Fig. 3 shows the results from two sets of 100 runs in base rates of 0.25 and 0.75, and the systematic conservative bias in the results.

---

[1] Our sensitivity analysis shows that the model is not qualitatively sensitive to initial threshold and to random seeds, and it is able to find the optimal threshold.

[Insert Fig.3 about here]

Up to this stage the full feedback model fulfills our primary expectations, showing that the model is able to learn in the base rate of 0.5, and it is conservative in base rates of 0.25 and 0.75. Now, we will estimate the conditions under which it provides a good explanation and replication for conditional feedback situations.

4. CONDITIONAL FEEDBACK MODEL

We consider the case, discussed above, of conditional feedback in which no feedback is provided when the decision is negative. In this case, increasing the threshold will decrease the selection rate and the decision maker will receive feedback about decision outcomes less frequently. When the decision is negative, the decision-maker can only guess about the outcome and the resulting payoff. Incorrect judgments about payoffs can affect threshold setting. The rest of the paper will investigate this effect.

4.1. Constructivist Coding

The important issue in modeling conditional feedback is about how people judge (code) the results of negative decisions (Elwin et al. 2007). In the absence of feedback, would a human resources manager judge that all of her negative decisions about last year candidates were correct?

Constructivist coding is defined as a coding that represents what one believes is true (Elwin et al. 2007). We define *p*, proportion of negative decisions assumed to be correct in the absence of feedback, as a parameter to use for payoff estimation. So, when *p* is 1, the model assumes there is no wrong negative decision and when is equal to 0 the model assumes all of its negative decisions were wrong. Therefore, *p* represents subject's average confidence on her negative decisions - the decisions she has not received feedback.[2] Payoff estimation in conditional feedback can be calculated using eq.2:

$V' = V_{Q,d=1} = (1-Q)*V_{fp} + Q*V_{tp}$ if $d=1$ (eq. 6)

$V' = V_{p,d=0} = (1-p)*V_{tn} + p*V_{fn}$. If $d=0$ (eq.7)

whereby V' represents expected value. We run the model 100 times for the base rate of 0.5, d' of 1, and for $p \in \{0, 0.1, 0.2 \ldots 1\}$, that is at least 9 runs for each p. Fig. 4 illustrates the average of results for each p. As we see, the model is sensitive to the value of p, which means the way that people interpret their negative decisions can substantially influence their results. At two extremes, people who believe their negative decisions were always right or wrong end up with a considerable systematic bias. This raises the importance of investigating how people really judge their negative decisions' performance.

[Insert Fig.4 about here]

There are three points about why there is a possibility for different people to code absent feedback differently. First, different people have different personality traits; some are more conservative, presumably, coding more false for their negative decisions. Second, people's confidence can change and influence their coding, and it can result in different averages (p) for different people in different conditions. Third, a second loop learning process – the process of questioning assumptions and learning about the effects of coding negative decisions, if it exists, can lead to a more realistic coding of false negatives.[3] However, existence of second loop learning is an empirical question.

---

[2] In fact, p will vary from case to case. Some cases are so clearly negative that the decision maker is virtually sure that a negative decision is correct, even if there is no feedback. Other cases are going to be less clear, leading to a lower value of p. Here, we are assuming a constant p which represents the average of p in different conditions.

[3] In this paper, we do not attempt to model second loop learning, and leave it for further research; however, the existence of second loop learning in data has to result in p less than 1 and close the actual percentage of true negatives. So, although we are not including second loop learning in our model, but its existence can be tested by the current model.

Most empirical studies of conditional feedback suggest that there is a tendency to underestimate the optimal selection rate or, in another word, to overestimate the threshold (Elwin et al. 2007, Stewart et al. 2007). Elwin and his colleagues argue that, in conditional feedback situations, people tend to code negative decisions as always correct. We call these individuals, confident constructivists. For them, $p=1$.

Fig. 5 shows simulation results from 100 runs for the base rate of 0.5 for a confident constructivist. This figure compares simulation results from full feedback condition with conditional feedback. Other parameters for conditional feedback are the same as full feedback condition (section 2). In this figure, threshold moves higher than the optimal threshold.

[Insert Fig.5 about here]

But what do these simulation results really mean? Basically, in a full feedback situation, false positive decisions increase the threshold, and false negative ones decrease it. As in conditional feedback, the confident constructivist assumes all negative decisions are correct, there is only one adjustment force, and that is from false positive results. Therefore, forces are always toward increasing the threshold (decreasing selection rate), and it continues until no incorrect positive decisions can be made.

Fig. 6 shows the results from two sets of 100 runs in base rates of 0.25 and 0.75. Results are compared with full feedback conditions. Interestingly, Fig. 6a confirms the observation that in contrast to full feedback, in conditional feedback situations and under the confident constructivist assumption, threshold goes to the extreme for low base rates.

[Insert Fig.6 about here]

So far, we have shown how relaxing the assumption of full feedback, in the lack of second loop learning, can influence the final results, and the formerly observed conservative cutoff-placement phenomenon can disappear and lead to more extreme cutoff-placement in conditional feedback for low base rates. Particularly, considering suggestions of Elwin et al., (2007) and Stewart et al. (2007) we see how people can underestimate the optimal selection rate. But the question is what is the real value for *p*, or how do people really code their negative decisions? Don't they learn any thing about the performance of their negative decisions? Later on, we use data from Elwin et al. (2007) and narrow the possible values of *p* to find more about people's behavior.

5- REPLICATIONS OF AN EMPIRICAL INVESTIGATION

Elwin et al. (2007) conduct an experiment including binary and continuous decision making situations. Sixty four subjects performed a computerized task of predicting economic outcomes for companies varying on four continuous cues (e.g. number of staff) with values ranging from 0 to 10. Outcome was an additive function of the values of the four cues, with assignment of the cue weights of 4, 3, 2, and 1 to different concrete cue labels. The base rate of profitable companies was 0.5. In the binary set of experiments the subjects were supposed to select the companies for which they predict a positive profit.

The experiment had two major phases: First subjects had a series of training trials, and then they entered the test phase. In the training part, one group of subjects performed 120 trials of full feedback decision making, while another group performed 240 trials of conditional feedback. In the test phase for both groups, 60 judgments were made without feedback. They found that the subjects, who had the conditional feedback training, ended up with much lower selection rate in the test phase (0.33) in comparison with the other group (0.52). The authors propose a model of constructivist individuals that code all negative decisions (absent feedback) as correct in the training phase and their model fits the data.

Attempting to replicate their results with our model can help us to learn more about the dynamics of Elwin et al.'s argument and check whether or not the model can reproduce their results. Further, it might lead to new possible explanations for the data.

There are four main parameters in our model: base rate, level of expertise (*d'*), the correction proportion (*k*), and the level of confidence in coding absent feedback (*p*). In order to replicate the data, we run our model with base rate = .5 and nine conditions for different *d'* and *k*. In each of the 9 conditions we calculate the *p* that can replicate their selection rate.

Since there will always be a value of p that can reproduce a single data point, this replication is not validation of the model. We are investigating that, if the model is correct, what *p* will be implied from the data. Results are illustrated in Table-2. For each set of *d'* and *k*, 10 sets of simulations with different random seeds are conducted and the average of *p* is represented. Comparing *p* with the optimal *p* which results in unbiased threshold, named *p** (determined by simulation), shows the level of overconfidence.

[Insert Table-2 about here]

As we see, the value of *p* that most closely approximates the results is always relatively high. This suggests that people may tend to underestimate false negatives in conditional feedback. This is consistent with Fischer and Budescu's (2005) argument about individuals' overconfidence in screening tasks, under conditional feedback. Further, it shows even if a second loop learning exists, it is not effective enough as people are not able to find the optimal *p*.

6. DISCUSSION AND CONCLUSION

In this study, we proposed a mathematical model of threshold learning which, while producing conservative cutoff-placement phenomenon in full feedback conditions, is able to produce threshold overestimation phenomenon in conditional feedback. The simulation outcomes and the replication of data show that conditional feedback can result in bias and underestimation of the base rate. Basically, assuming people learn from their indirect decisions, in conditional feedback, all (or most) of negative decisions are treated as correct ones. Therefore, the dominant adjustment force comes from false positive results, not from false negative ones. Thus, forces are always toward increasing threshold (decreasing selection rate). Our simulation runs with different *d'* (level of expertise), and *k* (correction ratio) show that regardless of the values of these parameters, we will always face overconfidence and bias in conditional feedback situations. This implies that conditionality of feedback for, e.g., police officers, human resources management, university admission office, etc. can result in misperception of performance and overconfidence.

Our simple model of anchoring and adjustment behavior which is rooted on the error correction logic, without any second loop learning, fits the data from Elwin et al. (2007). Some may argue that in the existence of second loop learning, people may try new thresholds, correct their perception of false negative results, and find the optimal threshold. Although we do not have second loop learning in our model, our empirical investigation shows that people do not find the optimal threshold. The average *p* (subjective proportion of correct negative decisions) is always higher than the actual ratio (Table-2) indicating overconfidence. This simply shows that even if second loop learning exists, it works for a limited number of people, and the average person is not able to find the optimal *p*. All of these results show that conditionality of feedback can be considered as a barrier to learning as it makes it very difficult for people to learn the optimal threshold.

There are several possible ways for extending this study. First of all, another study should be conducted to compare different learning algorithms' performance in predicting the results from conditional feedback. While reinforcement learning (Erev et al. 1995) and hill climbing (Busemeyer and Jae Myung 1992) models have been better predictors of threshold learning in full feedback, there is a need for comparative study of these models under conditional feedback.

Second, discussing about how different *p* can be used to replicate the data, we find a wide range of possible *p* that can produce the data. This result comes from the fact that there is an interactive relation between the level of expertise (*d'*) and the optimal *p*. We may argue that, actually, none of *d'* or *p* are constant for an individual, but they may change dynamically through the process. Although this is more an empirical question, but intuitively we can accept that there can be some endogenous changes in these two variables. While experiencing, people learn about cue weights and it increases d'. Further, dynamics of confidence can lead to a change in *p*. Studying effects of these additional loops can be very interesting.

Third, personality traits are shown to influence learning in binary decision making and the level of confidence (Klayman et al. 1999). So, we can expect personality trait to affect how people interpret their negative decisions and learn in conditional feedback situations. In further studies, individual level data can be gathered, and the model can be calibrated for each individual. Different parameters can then be compared. Testing a hypothesized relationship between some of the Big Five personality characteristics (like openness) and the way that people code negative decisions (*p*) is another possible and interesting way to extend this study.

Overall, our main claim in this paper was that the proposed simple mathematical model can create seemingly conflicting behaviors: conservative threshold placement in full feedback and threshold overestimation in conditional feedback. In addition, we demonstrated how conditionality can be a barrier to learning and can lead to overconfidence.

REFERENCE

Balzer, W. K., Doherty, M. E., & O'Connor Jr, R. (1989). Effects of cognitive feedback on performance. ***Psychological Bulletin**, 106*(3): 410-433.

Bereby-Meyer Y. & Roth, A. E., 2006. The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation, ***American Economic Review**, 96*(4): 1029-1042.

Biele, G., Rieskamp, J., & Czienskowski , U. 2008. Explaining cooperation in groups: Testing models of reciprocity and learning, ***Organizational Behavior and Human Decision Processes.*** *106* (2): 89-105

Busemeyer, J. R. & Jae Myung, I. 1992. An adaptive approach to Human decision making: learning theory, decision theory, and human performance, ***Journal of Experimental Psychology: General**, 121*(2): 177-194.

Dalgleish, L. I., & Smillie, L. D. 2006. Decision-Making Under Partial Feedback: Perceptual Discrimination and Application to Personnel Selection Decisions.: University of Stirling.

Denrell, J. & March, J. G. 2001. Adaptation as information restriction: The hot stove effect. ***Organization Science** 12*(5): 523-538.

Einhorn, H. J., & Hogarth, R. M. (1978). Confidence in judgment: Persistence of the illusion of validity. ***Psychological Review**, 85*(5): 395-416.

Elwin, E., Juslin, P., Olsson, H. & Enkvist, T. 2007. Constructivist Coding – Learning From Selective Feedback, ***Psychological science** 18* (2): 105-110.

Epley, N. & Gilovich, T. 2001. Putting adjustment back in the anchoring and adjustment heuristic differential processing, ***Psychological science** 12* (5): 391-396.

Erev I. 1998. Signal detection by Human Observers, ***Psychological Review** 105 (2)*: 280-298.

Erev, I., Gopher, D., Itkin, R. & Greenshpan, Y. 1995. Toward a generalization of Signal Detection Theory to n-person games: The example of two person safety problem. ***Journal of Mathematical Psychology**, 39:* 360-376.

Fischer, I., & Budescu, D. V. 2005. When do those who know more also know more about how much they know? The development of confidence and performance in categorical decision tasks. ***Organizational Behavior and Human Decision Processes**, 98:* 39-53.

Gunnthorsdottir, A. & Rapoport, A., 2006, Embedding social dilemmas in intergroup competition reduces free-riding, ***Organizational Behavior and Human Decision Processes**, 101* (2):184-199.

Green, D. M., & Swets, J. 1966, ***Signal detection theory and psychophysics.*** New York: Wiley.

Hammond, K. R. 1996. ***Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice.*** New York: Oxford University Press.

Kac, M. 1962, A note on learning signal detection. ***IRE Transactions on Information Theory****,* IT-8: 126-128.

Kac, M. 1969, Some mathematical models in science. ***Science****, 166*, 695-699.

Klayman, J., Soll, J. B., González-Vallejo C. & Barlas. S. 1999. Overconfidence: It Depends on How, What, and Whom You Ask. ***Organizational Behavior and Human Decision Processes*** 79(3): 216-247

Kubovy, M. & Healy, A. F. 1977. The decision rule in probabilistic categorization: what it is and how it is learned, ***Journal of Experimental Psychology: General****, 106 (4)*: 427-446.

Macmillan, N. A., & Creelman, C. D. 1991. ***Detection theory: A user's guide.*** Cambridge: Cambridge University Press.

Rahmandad, H. 2008. Effect of delays on complexity of organizational learning. ***Management Science*** *54*(7): 1297-1312.

Rahmandad, H., Repenning, N. P. & Sterman, J. D. 2008. Effect of Feedback Delays on Learning, *Working paper*. Available from http://www.mit.edu/~hazhir/papers/Learning-Rahmandad.pdf

Sterman, J., 1989a. Misperceptions of Feedback in Dynamic Decision Making. ***Organizational Behavior and Human Decision Processes***. 43(3): 301-335.

Sterman, J., 1989b. Modeling Managerial Behavior: Misperceptions of Feedback in a Dynamic Decision Making Experiment. ***Management Science***. 35(3): 321-339.

Stewart, T., Mumpower, J. & Holzworth, J. 2007. More on learning to make judgment and decisions in an uncertain world, ***23rd Annual International Meeting of the Brunswik Society, Barcelona/Casablanca Rms, Westin Long Beach, Long Beach, CA***

Swets, J. A., 1991, The science of high stakes decision making in an uncertain world (Transcript of a Science and Public Policy Seminar). Washington, D.C.: Federation of Behavioral, Psychological and cognitive sciences.

Swets, J.A., Dawes, R.M. & Monahan, J. 2000. Psychological science can improve diagnostic decisions. ***Psychological Science in the Public Interest*** 1(1): 1-26

Tindale, R. S. (1989). Group vs individual information processing: The effects of outcome feedback on decision making. ***Organizational Behavior and Human Decision Processes**, 44*(3): 454-473.

Tversky, A. and Kahneman, D. 1974, Judgment under Uncertainty: Heuristics and Biases. ***Science**, 185*: 1124-1131.

Weaver, E.A. & Richardson, G.P. 2006. Threshold setting and the cycling of a decision threshold. ***System Dynamics Review**, 22*(1): 1-26.

Figures:

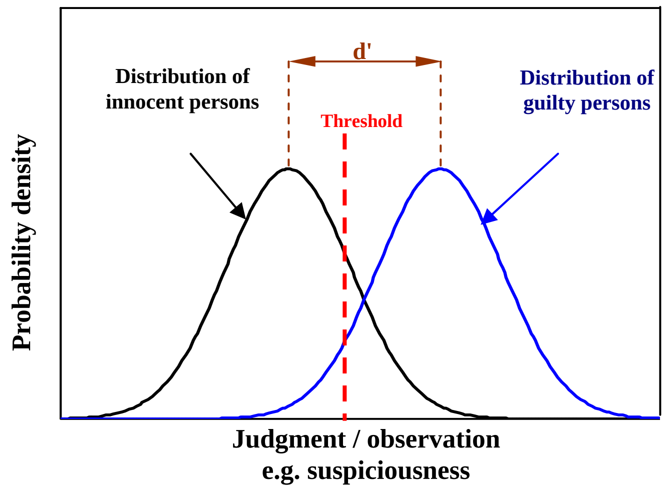

Fig.1: Distribution of innocent and guilty persons in detection theory

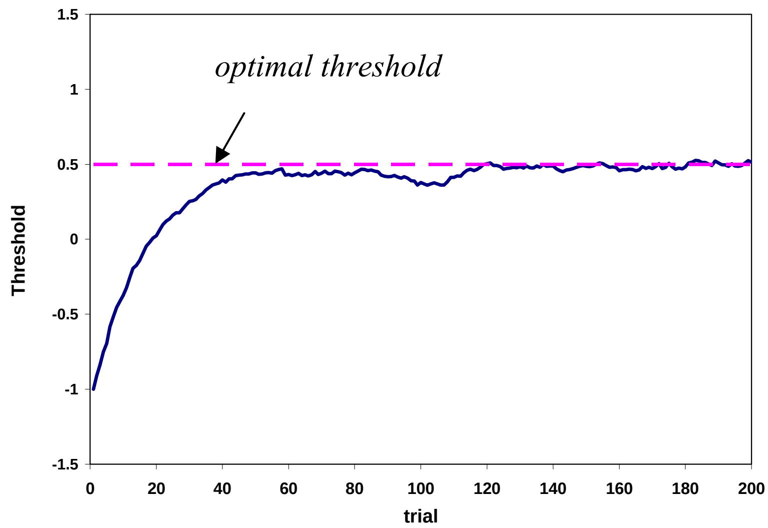

Fig. 2. Threshold dynamics (base rate = 0.5, k=0.05)

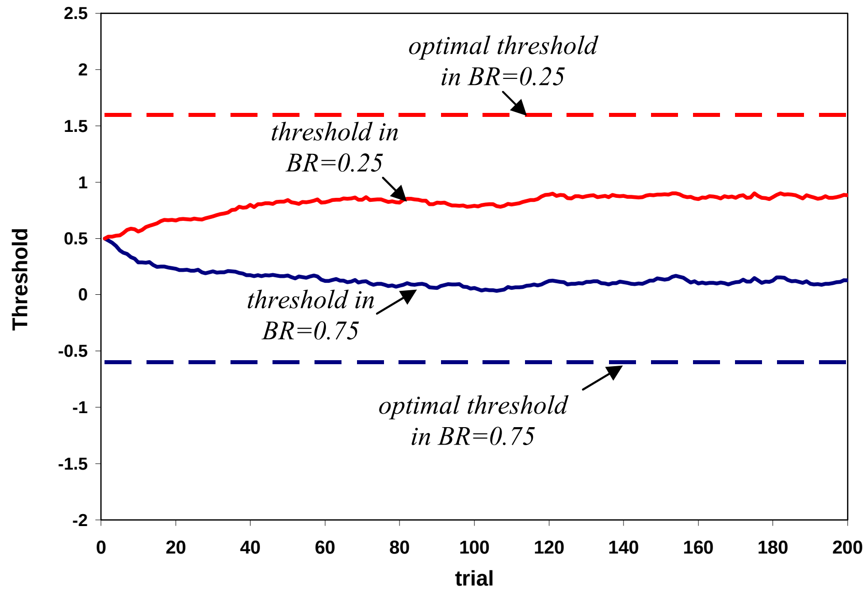

Fig. 3. Threshold dynamics for the base rate of 0.25 and 0.75 in full feedback, k=0.05

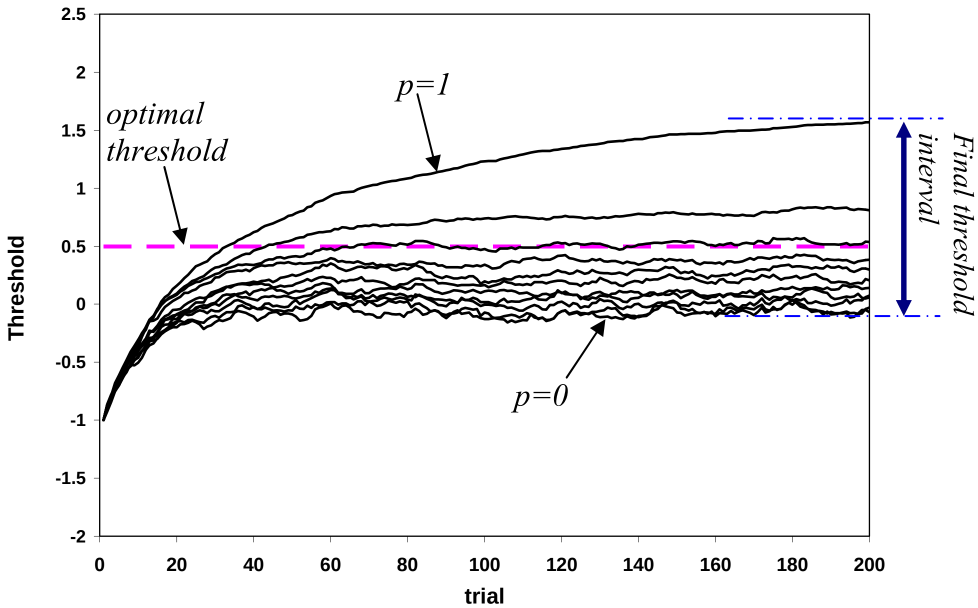

Fig. 4: Possible threshold dynamics for different *P*s (base rate = 0.5, k=0.05)

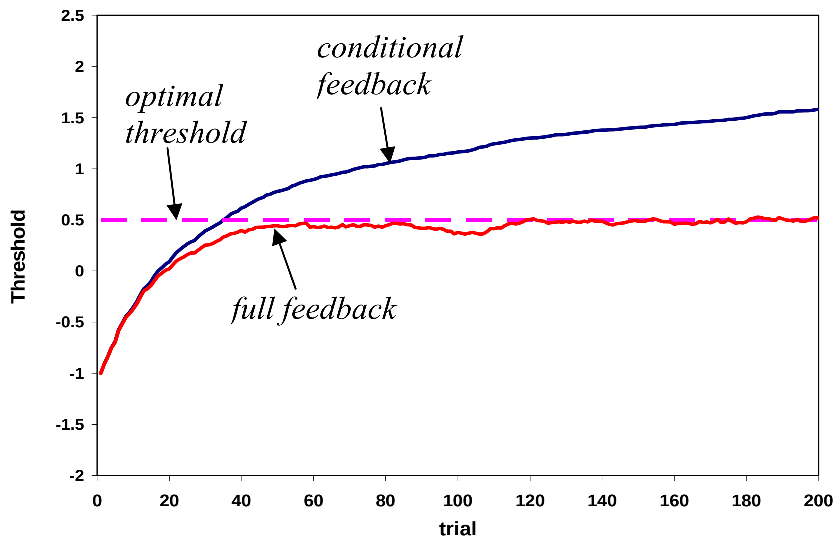

Fig. 5. Threshold dynamics for full feedback and conditional feedback (p = 1, base rate=0.5, k=0.05)

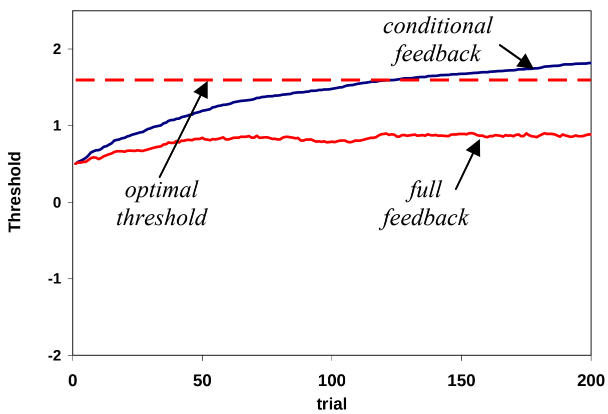

(a) base rate = 0.25

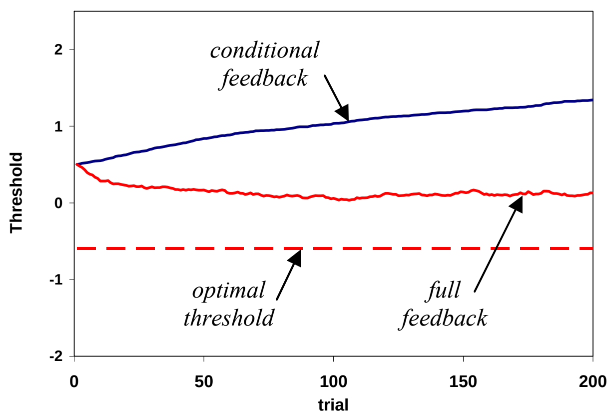

(b) base rate = 0.75

Fig. 6. Threshold dynamics for full feedback and conditional feedback (p = 1, k=0.05).

Tables:

State of the world		Decision: NO (not search)	Decision: YES (search)
	YES (guilty)	False negative	True positive
	NO (innocent)	True negative	False positive

Table 1 : Four possible outcomes

d'	k	average *p* to replicate data	*p** (optimal *p*)	overconfidence *p**- *p*
0.5	0.1	0.9	0.675	0.225
0.5	0.05	0.9	0.675	0.225
0.5	0.02	0.9	0.65	0.25
1.0	0.1	0.95	0.85	0.1
1.0	0.05	0.95	0.825	0.125
1.0	0.02	0.975	0.825	0.15
1.5	0.1	0.98	0.925	0.055
1.5	0.05	1	0.925	0.075
1.5	0.02	1	0.95	0.05

Table-2: Simulation results for different d' and k, to find 1) the range of *p* that can replicate Elwin et al.'s selection rate (.33) in the conditional feedback conditions, 2) the range of *p under which the model finds the optimal threshold, and 3) the difference between these two representing overconfidence.**