Journal of Experimental Psychology: Applied
1997, Vol. 3, No. 2, 83–104

1076-898X/97/$3.00

Beyond Bayes's Theorem: Effect of Base-Rate Information in Consensus Games

Sharon Gilat, Joachim Meyer, Ido Erev, and Daniel Gopher
Technion–Israel Institute of Technology

Many of the practical implications of behavioral decision-making research are based on the assumption that behavioral trends have to be compared with normative prescriptions. The present article demonstrates that in certain settings this approach is both inapplicable because there is no "normative" prescription and unnecessary because robust quantitative predictions can be made without reference to normative prescriptions. Experiment 1 demonstrates that a simple learning rule can be used to predict the base-rate effect in consensus games with multiple equilibria. Experiment 2 shows that information about the payoff rule affects participants' initial propensities but does not affect the learning process. Some implications of these results for the understanding of decision groups in social contexts, such as employment decisions in organizations, are pointed out.

Most attempts to derive practical implications from research on judgment and decision making are based on comparisons of observed descriptive tendencies to normative prescriptions of rational decision theory. Much of this research focuses on a debiasing technique that reduces the difference between observed decisions and the optimal prescription (e.g., see Von Winterfeldt & Edwards, 1986). Whereas this normative-based approach has led to useful results, its applicability is limited. The present article demonstrates this limitation and provides an evaluation of an alternative approach by focusing on one example: the effect of base-rate information in situations in which decision makers are motivated to reach a consensus. We chose this example because Bayes's theorem, which implies a specific utilization of base-rate information, is one of the best understood and most useful tools of rational decision theory.

Previous studies that explored the effect of base-rate information on subjective judgment and decision making focused on simple situations in which the optimal utilization of the base rates can be derived from Bayes's theorem. Generally, it has been found that people tend to use base-rate information less than optimally (e.g., Bar Hillel, 1980; Beyth-Marom & Fischhoff, 1983; Fischhoff & Bar-Hillel, 1984; Kahneman, Slovic, & Tversky, 1982; Kahneman & Tversky, 1973; Lyon & Slovic, 1976; Tversky & Kahneman, 1974, 1980). This finding is quite robust, although numerous variables were shown to affect the utilization of base-rate information (e.g., Ajzen, 1977; Gigerenzer, Hell, & Blank, 1988; Nisbett & Ross, 1980; Tversky & Kahneman, 1982). The robustness of the insufficient usage of base-rate information is particularly clear in binary decisions that can be considered as signal

Sharon Gilat, Joachim Meyer, Ido Erev, and Daniel Gopher, Research Center for Work Safety and Human Engineering, Technion–Israel Institute of Technology, Haifa, Israel.

This study was supported by the Committee for Research and Prevention in Occupational Safety and Health, Israel Ministry of Labor and Social Affairs. Parts of this study were presented at the 15th Biannual Conference on Subjective Probability, Utility and Decision Making (SPUDM-15), Jerusalem, Israel.

Correspondence concerning this article should be addressed to Ido Erev, Department of Industrial Engineering and Management, Technion–Israel Institute of Technology, Haifa 32000, Israel, or to Joachim Meyer, who is now at the Department of Industrial Engineering and Management, Ben Gurion University of the Negev, P.O. Box 653, Beer Sheva 84105, Israel. Electronic mail may be sent via Internet to joachim@bgumail.bgu.ac.il.

detection tasks. Experimental studies reveal that individuals tend to underweight the base rates, even after long practice periods (Healy & Kubovy, 1981; but see Birnbaum, 1983, for a critical assessment of Bayesian inference as the basis for assessing optimal behavior with base-rate information).

Whereas the comparison of base-rate utilization to the prescription of Bayes's theorem is instructive, it is important to recall that the situations under which Bayes's theorem describes rational behavior are quite limited. Bayes's theorem can be used to find the alternative that maximizes utility when decisions are made under uncertainty, but it is often insufficient for computing the best course of action in social interactions involving numerous actors who may be interdependent in the outcomes of their actions.

The case we present here are situations where actors not only attempt to be accurate in their decisions but also strive to reach agreement with another person. For instance, when deciding in what restaurant to reserve a table for dinner with a guest whose preferences are not known, restaurant goers will probably choose a place that they not only like themselves but also one they hope the guest will enjoy; they will make the decision by guessing from what they think most other people like. Base rates are likely to have a major effect on behavior in situations of this type. We refer to situations in which consensus with another person is reinforcing as *consensus games*. In the present article, we demonstrate that in these situations there often exists more than one optimal way to use a base rate and that the effect of base rate on choice behavior can be substantial.

The basic type of decisions that we model here are binary decisions where an actor has to decide to which of two mutually exclusive categories an event belongs. A large number, in fact, perhaps most decisions in social or organizational settings are of this type. For instance, a personnel manager may have to decide whether a candidate should be hired for a certain job. Other examples include the bank official who has to decide whether to grant a loan to an applicant, the equipment safety supervisor who has to decide whether the system is safe to operate, or the reviewer of a scientific journal who has to decide whether a submitted paper should be accepted for publication. Decisions of this type are naturally modeled in terms of signal detection theory (SDT), as presented by Green and Swets (1966/1988).

Classic One-Person Signal Detection Theory

It is convenient to distinguish between three submodels that are implicit in classical SDT: the model of the strategies, the payoff matrix, and the choice rule. The first model assumes that the observer considers cutoff strategies along the continuum of a single variable (e.g., the candidates' qualifications for the job) in which each instance (i.e., candidate) is summarized by a certain value. Assuming that the two categories of candidates (those who should be hired and those who should not) each have a distribution of values, one can specify the possibility of distinguishing between members of the two categories through the standardized distance between the means of the two distributions termed *sensitivity* or d' (see Figure 1). Usually one distribution is *Noise* (N), whereas the other is referred to as *Signal* (S) or *Signal + Noise*. The decision maker, or personnel manager, is assumed to consider cutoff strategies. Each strategy implies a cutoff point above which the decision maker will categorize an event as belonging to the S distribution and below which the event is categorized as belonging to the N distribution. This point is usually denoted as response criterion or β.

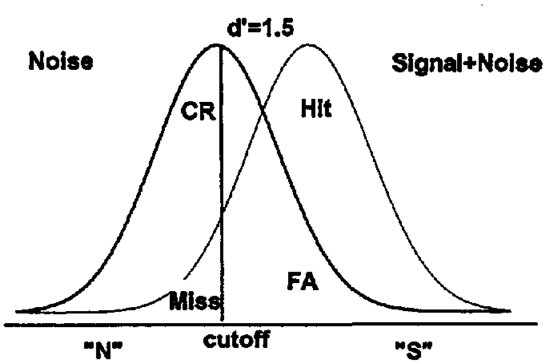

Figure 1. The Signal (S) and Noise (N) distributions in one-person signal detection. The distance between the means of the two distributions (d') is the sensitivity, and the cutoff point above which the decision maker indicates that this is a signal specifies the response criterion (β). CR = correct rejection; FA = false alarm.

The second model, the payoff matrix, is an abstraction of the possible outcomes. Because the perceived signal includes noise, the observer cannot be completely accurate; therefore, there are four possible outcomes (see Figure 2a): correct detection of a signal (*hit* in the common SDT terminology), correctly say N (*correct rejection* or CR), say N when a signal was presented (*miss*), or say S when there was no signal (*false alarm* or FA).

The third model, the choice rule, is often referred to as the *ideal observer model.* It implies that the detector is a utility maximizer who sets an optimal response criterion. For a single decision maker, the optimal response is computed by using Bayes's theorem. We should note that this is a strong interpretation of the ideal observer assumption. According to a weaker definition, an ideal observer uses a likelihood ratio response rule that may not be optimal.

a

		Player's Response	
		"Signal"	"Noise"
State of Nature	Signal	Hit	Miss
	Noise	False Alarm	Correct Rejection

b

	Player 1's Response	"Signal"		"Noise"	
	Player 2's Response	"Signal"	"Noise"	"Signal"	"Noise"
State of Nature	Signal	HH1 HH2	HM1 HM2	MH1 MH2	MM1 MM2
	Noise	FF1 FF2	FC1 FC2	CF1 CF2	CC1 CC2

Figure 2. Notations for one-person (a) and two-person (b) signal detection. For the two-person matrix (b), the upper left combination of letters indicates the utility for Player 1, and the lower right combination of letters indicates the utility for Player 2. The letters correspond to the four possible outcomes in signal detection theory, hit (H), miss (M), false alarm (F), and correct rejection (C), so that, for instance, the combination MH2 indicates the utility for Player 2 when Player 1's response was a miss and Player 2's response was a hit.

Two-Person SDT

Assume that the decision maker is not working in isolation; rather, the decision maker has to reach agreement with another person. For instance, the personnel manager mentioned above has to reach the same decision as another official in the company in order for this decision to be approved. It is natural to assume that not having a decision approved may be unpleasant. Under this assumption, the two actors are involved in a consensus game, with each trying to make decisions that are most likely to correspond to the other's decision. Analyzing situations like this requires an extension of SDT to two-person situations. Two complementary extensions have been suggested in the literature. Sorkin and Dai (1994) extended SDT to situations in which a team of observers who can communicate have to reach common decisions. Erev, Gopher, Itkin, and Greenshpan (1995) extended the theory to *noncooperative games,* situations in which each observer has to make an independent decision. The current article focuses on problems that are best modeled by Erev et al.'s extension.

Erev et al. (1995) noted that only the payoff matrix has to be modified in order to address two-person noncooperative signal detection games. Because there are only two players and two possible responses, the interaction between the players can be summarized in a payoff matrix (see Figure 2b). It presents the players' utility, given the state of nature and the two responses. The top left and bottom right entries in each cell present Player 1's and Player 2's utilities, respectively.

The models of the strategies and the decision rule do not have to be modified to consider two-person games. Yet the technical extension of the ideal observer model to the two-person case is not trivial. The calculation of the strategy that maximizes utility in the two-person game is different from the calculation in the case of a single detector. A utility maximizing player would have to behave differently when playing alone than when playing with other players. According to classical game theory, rational players choose their best response, which is based on the assumption that the other players will choose the best response to their own strategy. This logic leads to equilibrium predictions. Following Erev et al.

(1995), Appendix A presents the computation of the equilibria for two-person signal detection games when the two players are independent (see independent signals).

Formal Definition and Equilibria in Consensus Games

The notation of two-person games uses letter pairs to denote the payoffs for the various combinations. In each letter pair in the two-person condition, the left letter refers to Player 1 and the right letter refers to Player 2, where H = hit, C = correct rejection, F = false alarm, and M = miss. The numerals 1 and 2 indicate the payoff of either Player 1 or Player 2. For example, HH1 refers to the payoff for Player 1 when both Player 1 and 2 made a hit, and CF2 refers to the payoff for Player 2 when Player 1 made a correct rejection and Player 2 made a false alarm. In these notations, a two-person binary consensus game can formally be defined by two constraints on the possible payoffs: First, Player i's payoff, given a certain choice and a state of nature, is maximal when Player j makes the same decision; for example, for Player 1: HH1 $\geq$ HM1 and MM1 $\geq$ MH1 and CC1 $\geq$ CF1 and FF1 $\geq$ FC1; second, in at least one case, Player i benefits from consensus (for Player 1: HH1 $>$ HM1 or MM1 $>$ MH1 or CC1 $>$ CF1 or FF1 $>$ FC1. Note that in naturally occurring consensus games, the decision makers are likely to be in conflict between the attempt to find the best decision and the strive to reach consensus. Such a conflict can be modeled by the different payoffs.

In consensus games when the consensus is sufficiently important (for Player 1: MM1 $\geq$ HM1 and FF1 $\geq$ CF1), the extreme cutoffs ($\beta \rightarrow 0$ or $\ln \beta \rightarrow -\infty$ and $\beta \rightarrow \infty$ or $\ln \beta \rightarrow \infty$) are always equilibria. That is, when reaching a consensus is more important than being correct, there are equilibria in which players ignore the signal and always use one of the two responses; in this way, they can achieve 100% agreement rate. In addition, in some games additional equilibria points may exist. Moreover, in an extreme condition (considered below) where the two observers see the same signal in each trial (or a dependent signals case) every symmetrical pair of cutoffs is an equilibrium. Thus, generally there is no single rational cutoff in consensus games.

The case of dependent signals is probably extremely rare in actual social settings, and most situations that resemble consensus games should probably be considered independent games, that is, games in which the two players form their judgments, which are based on different stimuli. This is obvious when the two decision makers receive different information, as when two people interview an applicant independently, each obtaining somewhat different information from the interview. Yet even given that the two interviewers may have conducted the interview jointly, they will still be involved in an independent game. A candidate's evaluation is the result of the combination of the values in a number of dimensions (e.g., her or his past experience, communication skills, or demeanor), each of which is subject to some random error and may be weighted differently by different evaluators. Hence, the evaluators' scores on the combined dimension on which judgments are based (e.g., the applicant's suitability for the position) will most likely differ to some extent.

An Alternative Choice Model: A Reinforcement-Based Learning Rule

The observation that consensus games can have multiple equilibria does not imply that behavior in these situations is unpredictable. Rather, it implies that behavior is likely to be sensitive to the exact learning process. An understanding of the learning process is necessary in order to predict which equilibrium is an attraction point.

Roth and Erev (1995) suggested a simple reinforcement-based learning model that has been found to provide a useful approximation of behavior in a wide set of experimental games (see Bornstein, Erev, & Goren, 1994; Erev, Maital, & Or-Hof, in press; Erev & Roth, 1995; Ochs, 1995; Rapoport, Seale, Erev, & Sundali, in press; Roth & Erev, 1995). The model is based on the law of effect (Thorndike, 1898), which states that the probability that a certain strategy will be adopted increases when this strategy is positively reinforced. Similar models have been suggested by Harley (1981) to describe animal learning processes and by Bush and Mosteller (1955). Erev et al. (1995) adapted the model to detection games under the assumption that the set of strategies available to the players is a set of possible cutoff points (i.e., each value of β is considered a strategy). The adapted model's basic assumptions are

1. Each player considers a finite set of cutoff strategies.
2. The player has an initial propensity to select each of the strategies.
3. Reinforcements (payoffs relative to a reference point) affect the propensity to choose the selected strategy again in line with the law of effect with the addition of a generalization and a forgetting processes.
4. The choice probabilities are determined by relative propensities.

A quantification of these assumptions is required to allow a derivation of the model's predictions by a computer simulation. Erev et al.'s (1995) quantification is presented in Appendix B. This quantification provides a good approximation of Erev et al.'s results and reproduces the main experimental results obtained in traditional one-person signal detection studies (for a review, see Erev, 1995). In one-person tasks, the model typically (see exceptions in Erev) predicts a slow convergence toward the ideal observer's predictions. Thus, it reproduces the base-rate underutilization phenomenon in this case; the predicted cutoffs are less extreme than the cutoffs predicted under the assumption of efficient base-rate utilization.

The model's predictions for consensus games depend on the dependency between the two observers; it predicts an extreme base-rate effect when the two observers are independent, but a weak effect, as in the one-person case, when the two players see dependent signals.[1] In addition, the model predicts more variability in cutoff selection in the two-person task with independent observers. This variability is expected to lead to lower estimated d' scores in this condition. The model's quantitative predictions for the games studied here are presented in the *Results and Discussion* sections below. These sections also show that the model's qualitative predictions, which are stated above, are relatively insensitive to the exact choice of parameters.

The set of experiments described below were designed to examine how people perform in experimental consensus games. Specifically, the degree to which the equilibrium predictions and the predictions gained from the reinforcement-based learning model correspond to participants' behavior in these games was assessed. The experiments utilize an external signal detection paradigm (see Kubovy, Rapoport, & Tversky, 1971; Lee, 1963) in which participants are asked to categorize stimuli that are sampled from one of two normal distributions.

We decided to conduct an entire game rather than have participants play against the computer in order to avoid arbitrary decisions about the computer strategy. In games with multiple equilibria, like the game we describe below, minor changes in the opponent's strategy can lead to completely different behavior.

Experiment 1

Experiment 1 focused on the consensus game presented in Figure 3b, as it was played by two symmetrical players with administrated d' (the distance between the two experimentally determined normal distributions, referred to as D') of 1.5 when the prior probability (base rate) of a signal, which is denoted $p(\mathrm{S})$, is .6. Given these parameters and assuming that the functions $\mathrm{S}_i(.)$ and $\mathrm{N}_i(.)$ are symmetrical normal distributions, with variance 1 and with the means $-D'/2$ and $D'/2$, respectively, Appendix A implies that the game has three distinguishable sets of equilibrium points: Two low-equilibria cutoffs at $\beta \rightarrow 0$ ($\ln \beta \rightarrow -\infty$) and $\beta = .02$ ($\ln \beta = -4$), a moderately low-equilibrium cutoff at $\beta = .42$ ($\ln \beta = -.87$), and two high-equilibria cutoffs at $\beta = 163$ ($\ln \beta = 5.8$) and $\beta \rightarrow \infty$ ($\ln \beta \rightarrow \infty$).

In choice tasks, the base-rate effect can be

---

[1] Whereas the model's predictions are not trivial (we were surprised by them), they can be understood by examination of the reinforcement structure of a simplified situation in which Player i sets a cutoff at the center ($\ln \beta_i = 0$) and only j learns from experience. Given this cutoff, Player i will choose the frequent responses more than 50% of the time (assuming $d_i' > 0$). Consider the independent case first. From j's point of view, the simplified game is similar to a one-person game in which the frequent response is more likely to be reinforced because i chooses it in more than 50% of the trials. Thus, an extreme base-rate effect is predicted. In the dependent game, on the other hand, j maximizes reinforcement probability by selecting i's cutoff (the center). Some learning is still possible due to the generalization process.

a

		Player's Response	
		"Signal"	"Noise"
State of Nature	Signal	10	-10
	Noise	-10	10

b

	Player 1's Response	"Signal"		"Noise"	
	Player 2's Response	"Signal"	"Noise"	"Signal"	"Noise"
State of Nature	Signal	10 10	-10 -10	-10 -10	-10 -10
	Noise	-10 -10	-10 -10	-10 -10	10 10

Figure 3. Payoff matrices for the one-person (a) and the two-person (b) conditions in the experiments. For the two-person matrix (b), the upper left number is the payoff for Player 1, and the lower right number is the payoff for Player 2.

evaluated by the proportion of choices of the more frequent category. At the three distinct equilibria, the relevant proportions are expected to be 1.0, .72 and 0, respectively. Thus, rational decision theory cannot be used to predict the optimal base-rate effects.

As noted above, the rational behavior is even less clear when players are dependent (i.e., they see the same stimulus). In this case, every symmetrical cutoff set ($\beta_1 = \beta_2$) is an equilibrium because Figure 3 implies that Player *i* should try to make the same decision as Player *j*. In the dependent game, the probability for consensus reaches 1 when both players choose the same strategy. Thus, Player *i* cannot improve his or her gains by selecting a cutoff that differs from Player *j*'s cutoff. Clearly, in the dependent games, it is impossible to predict a specific optimal strategy that may be adopted by the players.

In comparison to the multiple equilibria that characterize the two-person games, there is typically one optimal cutoff in traditional one-person signal detection tasks. For example, in a one-person simplification of the present game (see Figure 3a), the optimal cutoff is $\beta = .67$ ($\ln \beta = -.4$). The expected proportion of the frequent category in this case is .63.

Method

Participants

Fifty Technion students served as paid participants in the experiment. They performed the experiment in unmixed male or female pairs and were randomly assigned to the experimental conditions. One pair had biased prior information regarding the experiment and was omitted from the experiment. The exact payoffs were contingent on performance and ranged from 18 to 26 shekels ($6.00–$8.50).

Apparatus

The experiment was programmed and conducted with Visual Basic 3 for Windows 3.1. This system was installed on a 486 PC with a Super VGA screen approximately 35.5 cm diagonal. In the display, there were two gray rectangles (6.6 cm wide × 13.0 cm high on a blue background as shown in Figure 4). Above each rectangle were two white fields (2.2 cm wide × 0.8 cm high) in which participants' scores were displayed. The lower field showed the last trial's score, and the upper field showed the cumulative score.

Participants in all experimental conditions sat side by side facing the screen at a distance of approximately 50.0 cm from the screen. They responded by pressing keys on the keyboard

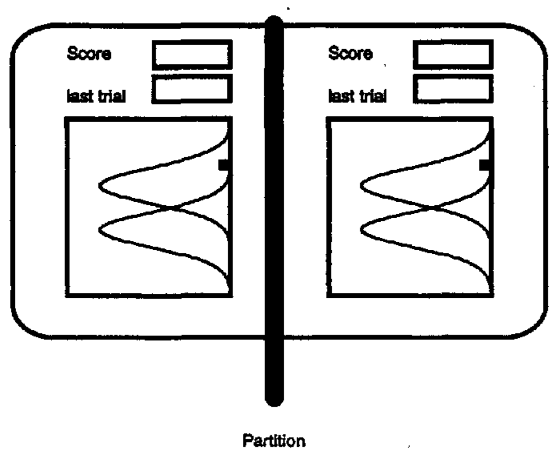

Figure 4. Schematic depiction of the experimental screen as seen by the participants. Each participant saw only one half of the screen, and the other half was hidden by the partition. The distributions are not shown during the experiment and are only added here for demonstration purposes.

positioned in front of them. A partition was placed between the participants in order to prevent each from seeing the displays or responses of the other participant.

In each trial, the participants saw for 1 s a white square, 0.2 cm wide and 0.2 cm high at 1 of 66 possible locations along the right border of the gray rectangle. After both participants in a pair responded, the feedback was shown, and after a 3-s interval the next stimulus was displayed. The stimuli were sampled from one of two normal distributions with equal variances, one of which was 1.5 standard deviations (2.8 cm) higher than the other. (The distributions are shown in Figure 4 but were not displayed to the participants.) The base rate was set so that 60% of the stimuli were drawn from one distribution and 40% from the other distribution. The location where the stimuli appeared was determined randomly for each participant in the pair in the independent condition in each trial. That is, the participants saw the stimulus at different locations, although the stimulus was always sampled from the same distribution, the higher or the lower, for both of them, keeping the state of nature the same for both participants. In the dependent condition, both participants saw the stimuli at exactly the same location in each trial, thus the state of nature was identical for both players. In the single-player condition, half of the pairs saw the stimuli at exactly the same location in each trial, as in the dependent condition, and half of the pairs saw the stimuli at different locations at each trial, as in the independent condition.

Participants responded to the stimuli by pressing one of two keys. A response that the stimulus was sampled from the higher distribution was given by pressing the ⟨*6*⟩ key by the right player and ⟨*A*⟩ by the left player. A response that the stimulus was sampled from the lower distribution was given by pressing the ⟨*3*⟩ and ⟨*Z*⟩ keys, respectively.

Procedure

Participants received written instructions stating that they would see the heights of men and women from a sample of students and that their task was to decide for each stimulus whether it was a man or a woman from that sample (i.e., whether it belonged to the higher or the lower distribution). The participants were also told that there might be more or fewer men than women in the sample and they should not assume that the number of men and the number of women would be equal. The payoff matrix was not explained. Instead, participants were told that correct answers were rewarded in a probabilistic fashion; that is, incorrect answers were always costly, but correct answers were also costly in some of the trials.

The experiment consisted of five blocks with 100 trials each. At the beginning of each block, the participants received 2,000 points. In each trial, they could gain or lose 10 points according to the payoff matrix, and the cumulative score for each block could range between 1,000 and 3,000 points. Blocks were separated by a short break. At the end of the experiment for each participant, one of the blocks was randomly selected, and the participant was paid according to his or her score in this block. The value of a point was .01 shekel. The experiment lasted approximately 1 hr.

Experimental Design

There were 24 pairs who participated in the experiment, 8 in the single-player condition, 8 in the independent condition, and 8 in the dependent condition. In the two-person conditions, participants were rewarded only when both were correct and agreed with each other. In the single-player condition, participants were rewarded according to the correctness of their own decisions. The single-player condition was the control condition for both two-person games. Therefore, 4 pairs in this condition saw dependent stimuli (i.e., the stimuli appeared at the same location in the display for both players), and the other 4 saw independent stimuli. In order to balance the base rate, for half of the pairs in each condition 60% of the stimuli were drawn from the higher distribution and 40% from the lower; for the other half, 40% were drawn from the higher distribution and 60% from the lower.

Results and Discussion

Descriptive Statistics

Table 1 presents the mean raw results for the experimental conditions and blocks. To derive the

Table 1
Means and Standard Deviations for Hit and False Alarm (FA) Rates and Cumulative Profits (P) in the Three Experimental Conditions in the Five Blocks

	Two-person										One-person				
	Independent game					Dependent game									
Measure	1	2	3	4	5	1	2	3	4	5	1	2	3	4	5
P(hit)															
M	0.72	0.83	0.84	0.87	0.91	0.73	0.78	0.77	0.8	0.81	0.77	0.79	0.79	0.79	0.81
SD	(0.12)	(0.11)	(0.11)	(0.07)	(0.08)	(0.08)	(0.07)	(0.07)	(0.04)	(0.08)	(0.05)	(0.02)	(0.03)	(0.06)	(0.05)
P(FA)															
M	0.35	0.53	0.51	0.57	0.6	0.31	0.29	0.27	0.32	0.32	0.31	0.31	0.36	0.32	0.35
SD	(0.13)	(0.19)	(0.16)	(0.17)	(0.18)	(0.08)	(0.10)	(0.07)	(0.13)	(0.07)	(0.1)	(0.09)	(0.12)	(0.07)	(0.09)
P															
M	1,985.0	2,005.0	2,082.5	2,057.5	2,130.0	2,227.5	2,327.5	2,377.5	2,397.5	2,407.5	2,473.8	2,502.5	2,458.8	2,496.3	2,486.3
SD	(194.1)	(206.1)	(136.7)	(183.8)	(97.4)	(161.8)	(112.1)	(88.44)	(89.08)	(127.4)	(91.95)	(83.79)	(74.34)	(57.8)	(85.01)

signal detection statistics, we arbitrarily designated the more frequent distribution as signal and the less frequent distribution as noise. According to this distinction, P(hit) and P(FA) are the proportion of frequent responses for stimuli that were sampled from the more and less frequent distribution, respectively. The profit (P) is the number of points accumulated during an experimental block beginning with initial 2,000 points. Table 1 presents the means and standard deviations of these averages for the participants.

In line with traditional signal detection research, the analysis reported below focuses on statistics derived from P(FA) and P(hit): ln β and estimated d'. In addition, we analyzed the proportion of frequent responses, P(frequent), that provides a direct assessment of the base-rate effect. Because players in a pair cannot be considered to be independent, we used pairs as the unit for statistical hypothesis testing. Hence, we first computed the variables separately for each participant and block, and then the mean for each pair was computed. These means were analyzed with two- way analyses of variance (ANOVAs) Condition $\times$ Block with repeated measures on the experimental block. We present these analyses following the derivation of the model's predictions.

The Learning Model's Predictions

Basic simulations. We conducted computer simulations in order to derive the model's predictions in a virtual replication of the experiment. There were 800 virtual pairs who participated in the simulation of each experimental condition (100 virtual replications of each of the eight pairs in each condition) with Erev et al.'s (1995) original parameters. Like the actual participants, each pair of virtual subjects participated in 500 trials. The following steps were taken in each trial.

1. The state of the world was randomly determined in accordance with the assumed prior probabilities.
2. The players' cutoffs (c_1 and c_2, which implies β_1 and β_2 values) were randomly determined in accordance with A4q (see Appendix B).
3. The perceived signals (x_1 and x_2) were selected from the assumed normal distribution, given the state of the world.
4. Players' decisions were determined (Player i's response was S if and only if $x_i > c_i$).

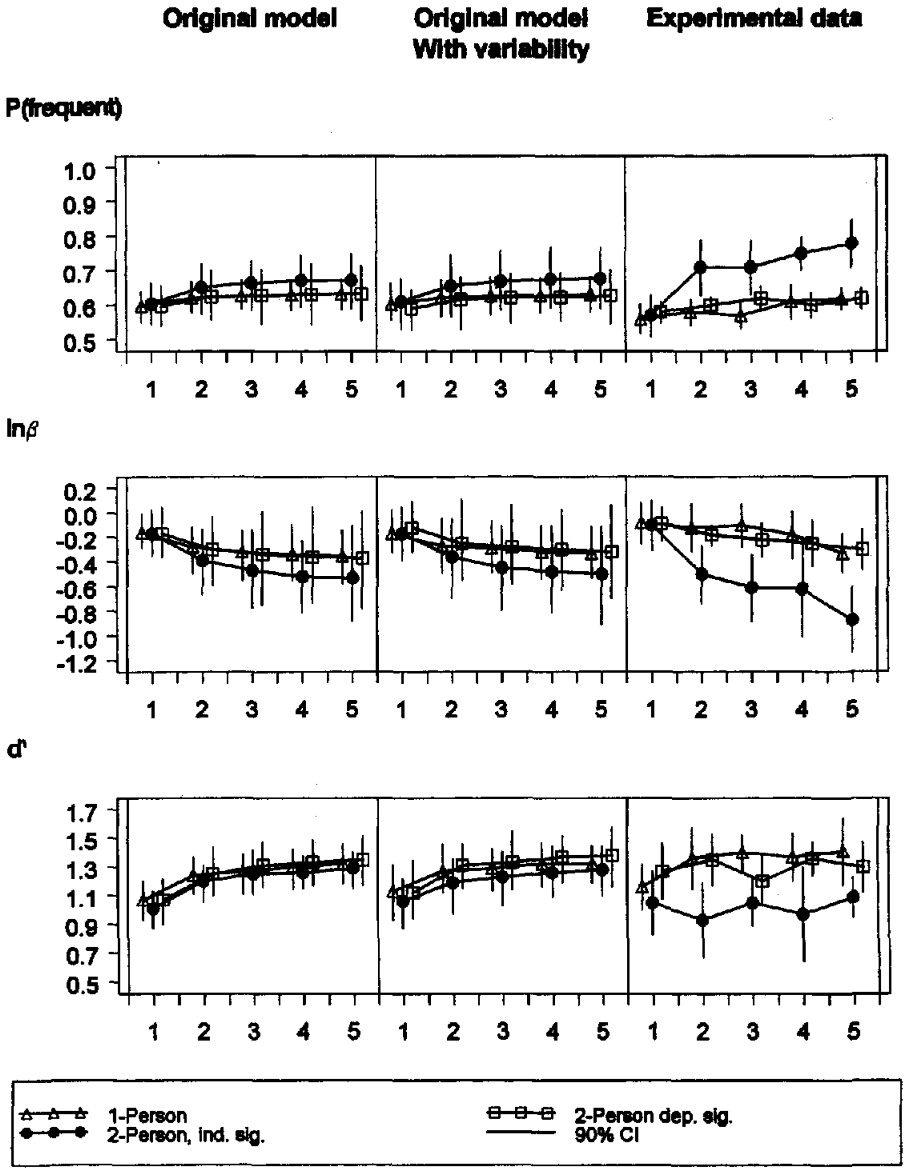

Figure 5. Results predicted from the learning model without variability (left) and with variability (middle) and the results obtained in Experiment 1 (right) for the three experimental conditions in the proportion of frequent responses, P(frequent), ln β and d'. ind. sig. = independent signal detection; dep. = dependent signal detection; CI = confidence interval.

5. Profits were calculated using the relevant payoff table.
6. Propensities were updated in accordance with A2q (see Appendix B).
7. The reference point for the next trial was calculated (see Appendix B).

To evaluate the model's predictive power, we calculated the same statistics that were calculated for the actual pairs for the virtual pairs. Figure 5 (left column) presents the three statistics that we chose to analyze. Each curve shows the average of 800 simulations summarized in five blocks of 100 rounds in one of the three experimental conditions. Three qualitative predictions can be observed: Teams with independent observers are expected to have more extreme ln β and P(frequent) and lower estimated d'. The first two effects are indications of faster learning in the direction of base-rate utilization in the two-person independent condition. The third effect is an indication of a larger within-subject cutoff placement variation in this condition.

The fact that the model is probabilistic means that it not only predicts a central tendency but it also makes a prediction of the complete distribution of possible outcomes. The 90% confidence intervals (CIs) predicted by the model are represented by the vertical lines in Figure 5. Recall that 100 virtual replications of the experiment were run and the CI were derived by setting the lower bound of the interval (for each statistic) by the 5th percentile's simulation and the upper bound by the 95th percentile. The CIs were also estimated under the assumption of normal distributions by estimating the virtual pairs' standard deviation, deriving the expected standard error in a sample of eight dyads, and by either adding or subtracting 1.645 time this value from the means. The two sets of estimated CI practically coincided. This finding implies that the observed means in the two-person condition do not reflect a prediction that each dyad will converge to one of the possible equilibria. Rather, there appears to be one attraction point in each condition.[2]

Sensitivity analyses. Two types of sensitivity analyses were conducted to evaluate the robustness of the model's predictions. The first set includes simulations in which one of the values of the nonzero original parameters (m, C_{max}, S(1), ϕ, υ, w^-, w^+, σ_i, σ_g) was either increased or decreased by 50%. There were 100 virtual dyads included in each of 18 additional simulations (nine parameters × two variations). Table 2 presents the average results over blocks of the original and the altered simulations. It shows that the three qualitative predictions stated above hold in all cases. This table also shows that the parameters affect the model's quantitative predictions.

A second set of sensitivity analysis simulations examined the effect of between-subjects variability in the learning parameters. In this analysis, the value of the seven nonzero learning parameters (all the parameters considered above with the exception of m and C_{max}, the cognitive strategies parameters) were randomly selected for each virtual subject. Following Erev and Roth (1995), the values were drawn from a uniform distribution with a mean at the value of Erev et al.'s (1995) original parameter. The values were selected from the interval, 0, 2 (original parameter). Thus, for example, the values for ϕ, .001 in the original model, were randomly selected from the interval (0, .002). Figure 5 (center column) presents the predicted learning curves under this variant of Erev et al.'s model. It shows that the between-subjects variability slightly improves the model's predictions.

Hypothesis Testing

Figure 5 (right column) presents the experimental statistics for which the model's prediction where derived. The data's 90% CI were calculated by either adding or subtracting 1.645 *SE* units from the data means.

Proportion of frequent responses. The proportion of frequent responses given by participants indicates the degree to which their responses correspond with the base rate. Figure 5 (upper right panel) presents the proportion of frequent responses observed in the three experimental conditions as a function of time. The main effect of the time was significant, $F(4, 21) = 9.8$, $p = .0001$, because of an increase of the proportion of frequent responses in all conditions. The overall proportion of frequent responses in the independent condition was significantly higher than in the other conditions, $F(2, 21) = 8.48$, $p = .002$. The last two blocks of the different conditions also revealed a significant difference, $F(2, 21) = 13.07$, $p < .01$, between the conditions that resulted from a higher proportion of frequent responses for the independent condition (.77) than the dependent (.62) and the single-player (.61) conditions. The linear trend between the first and last block was significantly stronger in the independent condition, $F(2, 21) = 4.1$, $p < .04$, than in the other two conditions. This pattern

[2] The overlap of the three confidence intervals (CIs) in each block implies that the power of the current test of the model is limited. Only over blocks does the probability of incorrect acceptance of the null hypothesis fall below .2. We have derived the CI prediction only after the completion of the first experiment, and we increased the design power by reducing d' and increasing the number of dyads in each condition in the second experiment.

Table 2
Sensitivity Analysis: Mean Predictions of One-Parameter Variants With Original Parameters

	Original game parameters								
	P(frequent)			ln β			d'		
	1-person	2-person independent	2-person dependent	1-person	2-person independent	2-person dependent	1-person	2-person independent	2-person dependent
Measure	.62	.65	.62	−.29	−.42	−.31	1.24	1.20	1.26
s(1)									
1.5	.61	.68	.61	−.37	−.54	−.26	1.26	1.24	1.32
4.5	.613	.633	.626	−.25	−.35	−.34	1.23	1.18	1.23
σ_g									
.0125	.64	.62	.66	−.31	−.43	−.40	1.24	1.19	1.27
.0375	.63	.69	.61	−.32	−.53	−.27	1.23	1.15	1.28
σ_i									
.75	.60	.61	.59	−.20	−.29	−.19	1.37	1.36	1.42
2.25	.66	.69	.65	−.40	−.53	−.38	1.13	1.09	1.15
w^-									
.01	.64	.67	.65	−.35	−.49	−.44	1.20	1.18	1.29
.03	.62	.67	.62	−.30	−.47	−.31	1.23	1.18	1.25
w^+/w^-									
.25	.624	.633	.626	−.31	−.34	−.31	1.23	1.18	1.24
.75	.65	.69	.66	−.39	−.54	−.46	1.21	1.19	1.26
υ									
.00005	.61	.65	.64	−.26	−.43	−.36	1.25	1.23	1.28
.00015	.62	.66	.62	−.29	−.43	−.28	1.26	1.21	1.22
ϕ									
.0005	.63	.67	.63	−.30	−.50	−.33	1.21	1.18	1.24
.0015	.63	.66	.60	−.34	−.45	−.20	1.24	1.21	1.27
m									
51	.59	.61	.58	−.18	−.28	−.12	1.38	1.35	1.39
151	.62	.68	.62	−.31	−.49	−.29	1.23	1.15	1.27
C_{max}									
2.5	.63	.68	.62	−.3	−.51	−.31	1.23	1.15	1.27
7.5	.63	.67	.61	−.35	−.52	−.27	1.24	1.19	1.29

Note. The original values of the parameters were $S(1) = 3$, $\sigma_g = .025$, $\sigma_i = 1.5$, $p(1) = 0$, $w^- = .02$, $w^+/w^- = .5$, $\upsilon = .0001$, $\phi = .001$, $m = 101$, $C_{max} = 5$. In each of the variations, the value of one of the nonzero parameter was either increased or decreased by 50%.

of results corresponds with the predictions of the learning model.

Response criterion ln β. The ln β statistic is an estimate of the players' response criterion, that is, the point above which all stimuli will be perceived as drawn from the signal distribution, and below it, from the noise distribution. The variable is computed as follows: $\ln \beta = cd'$, where c is $c = -0.5[Z(\text{hit}) + Z(\text{FA})]$. Figure 5 presents ln β values in the three experimental conditions as a function of time. The main effect of time was again significant, $F(4, 21) = 9.15$, $p = .0001$, as was the difference between the independent condition and the dependent and single-player conditions, $F(2, 21) = 6.33$, $p < .008$. The values of ln β were significantly lower in the independent condition. For the last two blocks of the different conditions, a significant difference between the conditions, $F(2, 21) = 7.51$, $p < .01$, was found. The value of ln β was lower for the independent condition (−.75) than for the dependent (−.26) and single-player (−.28) conditions. The linear trend between the first and last block was again stronger in the independent condition, $F(2, 21) = 3.94$, $p < .04$. These results also correspond with our predictions.

The sensitivity d'. Figure 5 (bottom row) shows the average estimated d' scores in the three conditions. The observed condition effect is significant, $F(2, 21) = 7.3, p < .004$. In line with the model's predictions, lower d' values were observed in the independent condition. Recall that under the model this effect is a result of an increase in cutoff variability; our virtual subjects had a fixed perceptual ability. Indeed, the participants in the experiment exhibited more violations of the assumption of static cutoff in the independent conditions.[3]

Cumulative profits. Table 1 presents the cumulative profits in Israeli shekels in the three experimental conditions and blocks. The main effect of time was significant, $F(4, 21) = 3.35$, $p < .02$, because of an increase of the cumulative profits in all three conditions. The difference between the independent condition and the dependent and single-player conditions, $F(2, 21) = 68.92$, $p < .0001$, was significant. The cumulative profits were significantly lower in the independent condition.

Quantitative Assessment of the Model's Predictions

The experimental results support the robust qualitative predictions made by the model. As noted above, these predictions are relatively insensitive to the exact values of the model's parameters. The current section examines the quantitative predictions made by the model with Erev et al.'s (1995) original parameters (with and without between-subjects variability). To derive quantitative fitness scores, we created a data set from each row in Figure 5. Each data set includes 15 observations (5 blocks $\times$ 3 conditions) and its variables are Figure 5's statistics.

Table 3 presents two types of fitness scores derived from these data sets: the proportion of observations outside the 90% CI, and the correlations between the data and the models. This table and Figure 5 reveal that the experimental statistics fall outside the original model's CI in 10 of the 45 cases (3 statistics $\times$ 3 conditions $\times$ 5 blocks). The central tendencies predicted by the original model fall outside the data's CI in 12 of the 45 cases.

Between-subjects variability improves the model's calibration. Only in 5 of the 45 cases did the observed statistics fall outside the model's CI when variability was assumed. This value is as close as possible to the predicted 10% under the assumption of perfect calibration. The model fell outside the data's CI in 8 of the 45 cases.

Table 3
Model Fitness Scores for Experiment 1

Model	P(frequent)	ln β	d'	Total
	Original parameters			
Exp out	3/15	1/15	6/15	10/45
Model out	4/15	3/15	5/15	12/45
r(model, exp)	.93*	.89*	.37	
	With variability			
Exp out	2/15	0/15	3/15	5/45
Model out	4/15	2/15	2/15	8/45
r(model, exp)	.94*	.91*	.53*	

Note. Frequency of experimental observations (exp) that fell outside the model's 90% confidence interval (CI, exp. out), the frequency of average predictions that fell outside the data's 90% CI, model out, and Pearson's correlations, r(model, exp). P(frequent) = proportion of frequent responses.
$^*p < .05$.

Positive correlations between the model's mean predictions and the mean results were observed in all 6 (3 statistics $\times$ 2 versions of the model) cases. Only one of the six correlations, d' with no variability, did not reach significance. The correlations for the response criteria and P(frequent) are above .89.

These results suggest that in addition to the accurate qualitative prediction, the current model with Erev et al.'s (1995) parameters provides a useful quantitative approximation of the data. In fact, with variability, the model's quantitative predictions cannot be rejected, even if the 45 statistics are assumed to be independent. The probability that of a sample of 45 independent observations 8 or more will fall outside the 90% CI is above .05. Of course, this does not mean that the model is accurate. Rather, the quantitative finding is important: The model provides a

[3] To assess cutoff variability, we assessed the proportion of *static cutoff violations* (SCV), which are defined as choices that are inconsistent with the static cutoff that best describes the participants' decisions. Over blocks, the proportion of SCV was 14.85 in the independent condition compared with 9.26 in the dependent condition and 10.21 in the single-player condition.

90% CI that is not too wide (over the three statistics, the geometric average is only 25% wider than the CI estimated from the data) and is well calibrated.[4]

Experiment 2

In Experiment 1, we demonstrated that participants adjust their response criteria in accordance with the predictions of the reinforcement-based learning model. This occurred even though participants were unaware of the game's payoff structure and their interdependence. One purpose of Experiment 2 was to understand the effect of prior knowledge regarding the exact payoff structure on participants' behavior. The question of how prior knowledge of being engaged in a consensus game affects behavior is of major importance for applying the results to actual social situations where people usually are aware that they are participating in a consensus game. Participants' prior knowledge regarding the payoff structure has obviously no effect on the optimal cutoffs. However, the knowledge may affect the actual learning process. In particular, as noted by Roth and Erev (1995), instructions are likely to affect the player's initial propensities. The sensitivity analysis (see Table 2) suggests that whereas such an effect is not likely to affect the qualitative trends, it can affect the magnitude of the differences.

A second purpose of Experiment 2 was to examine the learning process for games in which players have a lower sensitivity. Therefore, we used the same consensus game as in Experiment 1, which is presented in Figure 3, but the two symmetrical players' D' (administrated d') was 1.0 instead of 1.5. The prior probability of a signal was .6, as in Experiment 1. Given these parameters, Appendix A implies that the game has only two distinguishable equilibria: a low equilibrium at $\beta \rightarrow 0$ ($\ln \beta \rightarrow -\infty$) and two experimentally indistinguishable high equilibria at $\beta = 5.1$ ($\ln \beta = 1.63$) and $\beta \rightarrow \infty$ ($\ln \beta \rightarrow \infty$). The proportion of choices of the more frequent category at the two distinguishable sets of equilibria are expected to be 1 and 0, respectively. In the single-player conditions, the optimal cutoff is $\beta = .67$ ($\ln \beta = -.4$). The expected proportion of the frequent category in this case is .67.

Method

Participants

There were 96 Technion students who served as paid participants in the experiment. They performed in unmixed pairs of men and women who were randomly assigned to one of four experimental conditions. The exact payoffs were contingent on performance and ranged from 16 to 24 shekels ($5–$8). None of the participants in Experiment 2 had taken part in Experiment 1.

Apparatus

The experiment was conducted with the same computer system that was used in Experiment 1. Again, the base rates were 60% for one distribution and 40% for the other distribution. One major difference was the use of $D' = 1.0$ instead of 1.5, as in Experiment 1. The screen distance between the means of the two distributions remained the same (2.8 cm); thus the D' was reduced to 1.0 by increasing the variance of the distributions. The result was a greater difficulty to discriminate among the distributions.

Procedure

The general procedure was the same as in Experiment 1, but the experiments differed in the instructions the participants received. In the without-information conditions, participants received the same instructions as in Experiment 1. In the with-information conditions, participants received basically the same instructions, but they were also informed about the exact payoff structure and the game's rules according to the condition to which they were assigned (two person or one-person). Participants in the two-player condition were told that whenever they were correct and in agreement with each other, they would receive 10 points; in all other situations (when

---

[4] Note, however, that there appears to be an increase with time in the number of observations outside the model's confidence intervals (CIs). This observation suggests that the model does not capture all the factors that affect the experimental participants. For example, our virtual subjects are not getting tired even after 500 rounds.

they disagreed or when they agreed on an incorrect judgment) they were told they would lose 10 points. In the single-person condition with information, players were told that they would receive 10 points when they were correct and lose 10 points when they were wrong.

Experimental Design

There were 48 pairs who participated in the experiment: 12 in the single-player condition without information, 12 in the single-player condition with information, 12 in the independent condition without information, and 12 in the independent condition with information. In the two-person conditions, players were rewarded only when both were correct and agreed with each other. In the single-player conditions, players were rewarded according to the correctness of their own decisions. In order to balance the base rate, for half of the pairs in each condition, we drew 60% of the stimuli from the higher distribution and 40% from the lower. For the other half, 40% were drawn from the higher distribution and 60% from the lower.

Results and Discussion

Descriptive Statistics

Table 4 presents mean raw results. As in Experiment 1, we analyzed the three statistics derived from P(hit) and P(FA): $\ln \beta$, d', and P(frequent). The values of these statistics are presented in Figure 6 (see Columns 3 and 4). They are discussed following the derivation of the predictions of the learning model.

The Learning Model's Predictions

Again, computer simulations were conducted in order to derive the model's predictions with the parameters of the experiment. As in Experiment 1, 100 virtual replications of each of the conditions (1,200 virtual pairs) were conducted in which each simulated subject participated in 500 rounds summarized in five blocks of 100 rounds. The model's predictions with and without variability are presented in Figure 4 (see the lefthand columns). Comparison of these columns with the left-hand columns of Figure 5 shows that the decrease in D' (from 1.5 to 1.0) did not affect the trends predicted by the model. This assertion was reinforced in an additional sensitivity analysis that examined the effect of each of the parameters, which is in line with Table 4's analysis. As in Experiment 1, neither addition nor subtraction of 50% from any of the parameters affected the qualitative predictions.

Hypothesis Testing

Proportion of frequent responses. Figure 6 (top panel, columns 3 and 4) presents the proportion of frequent responses in the four experimental conditions as a function of time. The main effect of time was significant, $F(4, 44) = 8.65$, $p = .0001$, because of an increase of the proportion of frequent responses in all conditions, except for the single-player condition with information. The proportion of frequent responses over all five blocks in the two-person conditions was not significantly different from the single-player conditions, $F(1, 44) = 1.62$, *ns*. However, the proportion of frequent responses in the last two blocks combined was significantly higher in the two-person conditions, $F(1, 44) = 5.43$, $p < .03$, than in the single-player conditions. The linear trend between the first and last block was stronger in the two-person conditions, $F(1, 44) = 7.96$, $p < .001$. This pattern of results corresponds with the predictions of the learning model (see Figure 4).

In the overall analysis, neither the main effect of information nor any of the interactions involving information approached significance. In the separate analysis of the last two blocks, there was a significant main effect of the information, $F(1, 44) = 4.44$, $p < .05$. With information, the proportion of frequent responses was lower than without information.

Response criterion ln β. Figure 6 (see middle panel, columns 3 and 4) presents $\ln \beta$ values in the four experimental conditions as a function of time. The main effect of time was again significant, $F(4, 44) = 7.31$, $p = .0001$. The values of $\ln \beta$ over all five blocks in the two-person conditions were not significantly different from those in the single-player conditions, $F(1, 44) = .49$, *ns*. For the analysis of the combined $\ln \beta$ values in the last two blocks, there was a marginally significant difference between the single- and

Table 4
Means and Standard Deviations for the Proportion of Hit [P(hit)] and False Alarm [P(FA)] and Cumulative Profits in the Four Experimental Conditions in the Five Blocks

	With information					Without information				
Measure	1	2	3	4	5	1	2	3	4	5
					Two-person					
P(hit)										
M	0.66	0.70	0.71	0.73	0.77	0.70	0.71	0.69	0.75	0.81
SD	(0.11)	(0.12)	(0.13)	(0.14)	(0.12)	(0.08)	(0.01)	(0.12)	(0.14)	(0.12)
P(FA)										
M	0.37	0.44	0.47	0.47	0.49	0.45	0.49	0.51	0.55	0.62
SD	(0.14)	(0.15)	(0.19)	(0.22)	(0.19)	(0.15)	(0.16)	(0.14)	(0.18)	(0.15)
P										
M	1,863.33	1,880.00	1,878.33	1,886.67	1,968.33	1,845	1,886.67	1,816.67	1,903.33	1,926.67
SD	(132.34)	(134.57)	(148.56)	(153.35)	(134.15)	(84.48)	(143.04)	(171.38)	(160.64)	(168.27)
					One-person					
P(hit)										
M	0.69	0.69	0.70	0.70	0.70	0.70	0.75	0.74	0.75	0.76
SD	(0.06)	(0.06)	(0.05)	(0.06)	(0.06)	(0.08)	(0.08)	(0.08)	(0.08)	(0.06)
P(FA)										
M	0.39	0.42	0.39	0.36	0.39	0.40	0.43	0.44	0.44	0.46
SD	(0.09)	(0.06)	(0.06)	(0.08)	(0.08)	(0.09)	(0.18)	(0.10)	(0.08)	(0.08)
P										
M	2,330.00	2,300.00	2,315.00	2,345.00	2,318.33	2,311.67	2.353.33	2,336.67	2,336.67	2,350.00
SD	(83.65)	(65.29)	(56.32)	(58.58)	(44.34)	(77.87)	(63.65)	(76.17)	(70.06)	(84.04)

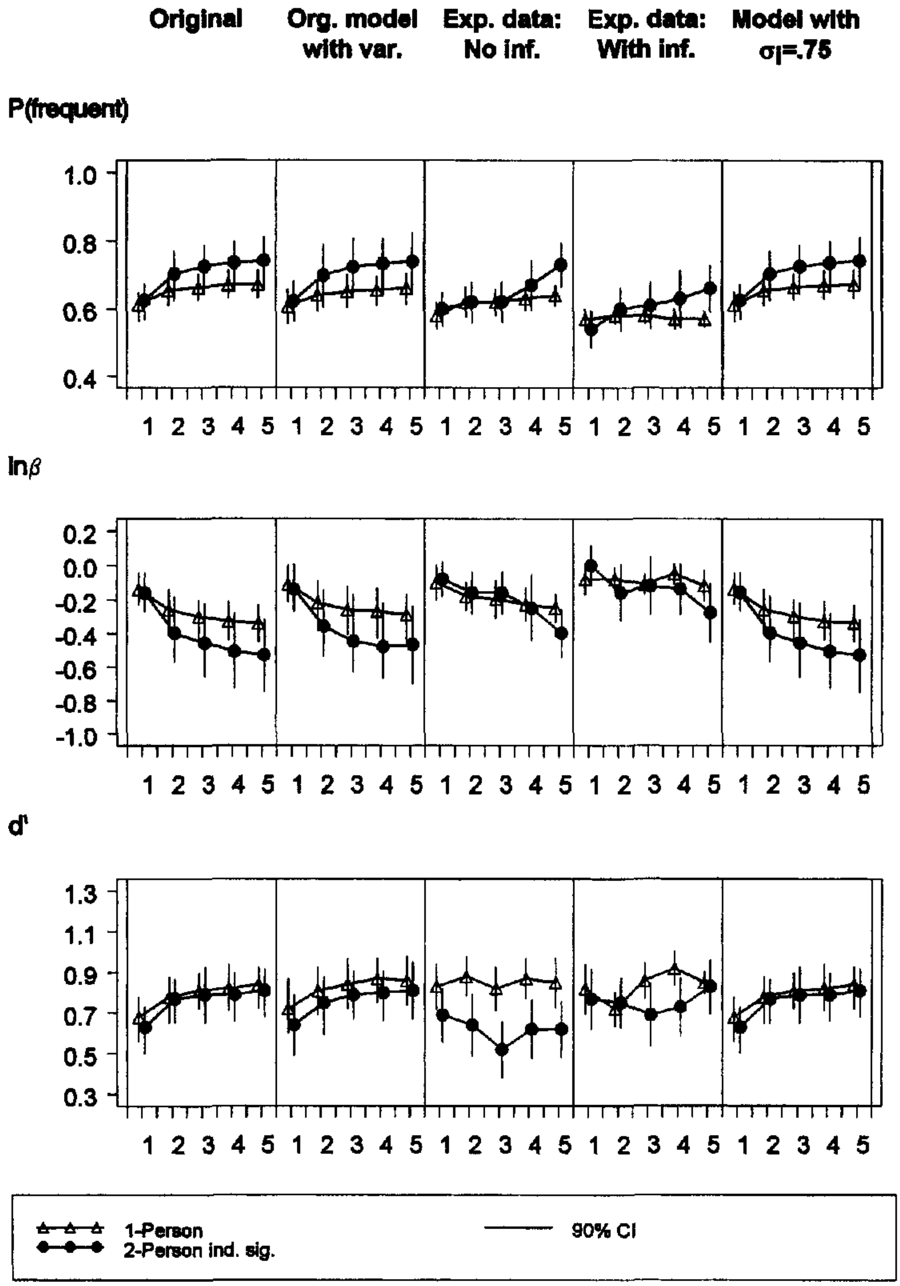

Figure 6. Results of the simulations and the empirical data for the proportion of frequent responses (P(frequent)), lnβ and d' for Experiment 2. From left to right are shown predictions of the learning model without variability (var.), predictions of the learning model with variability; experimental results for the condition in which participants had no information (inf.) about the game; experimental results when participants received information about the game; and predictions of the learning model with variability and a smaller variance of the distribution of the initial propensities. ind. sig. = independent signal detection; CI = confidence interval.

the two-person conditions, $F(1, 44) = 3.01, p < .09$. The linear trend between the first and last block for the two-person conditions was significantly stronger, $F(1, 44) = 5.62, p < .03$, than for the single-player conditions. These trends were predicted by the learning model.

Information had no main effect, and none of the interactions involving information approached significance in the overall analysis. For the last two blocks, there was a significant main effect of the information, $F(1, 44) = 4.72, p < .04$, which resulted from higher values of ln β in the conditions with information compared with the conditions without information.

The sensitivity d'. Replicating the findings of Experiment 1, the overall analysis revealed a significant main effect of the condition on the values of d', $F(1, 44) = 12.68, p < .001$. As in Experiment 1, the sensitivities were lower in the two-person conditions compared with the single-player conditions. Neither the main effect nor the interactions of information had a significant effect on the values of d'. Again, comparing the results of the experiment with the predictions of the learning model shows that the results correspond with the predictions (see Figure 6).

Cumulative profits. The main effect of time was significant, $F(4, 44) = 2.80, p < .03$, because of a slight increase of the cumulative profits in all three conditions. The difference between the two-person conditions and the single-player condition was significant, $F(1, 44) = 446.71, p < .0001$. The cumulative profits were significantly lower in the two-person conditions compared with the single-player conditions. The information had no main effect, and none of the interactions involving information approached significance.

Quantitative Assessment of the Model's Predictions

Table 5 presents Table 3's fitness scores for Experiment 2. It shows that whereas Erev et al.'s (1995) parameters with variability provide a good fit for the conditions without information, they have to be modified in order to account for the slower learning process observed with information. For example, the frequency of experimental statistics that fall outside the model with variability CI are 5 of 30 in the no-information condition, but 14 of 30 in the condition with payoff-matrix information.

Table 4 results are in line with Roth and Erev's (1995) assertion that manipulation of the initial propensities parameter can account for the observed information effect. Slower learning is predicted when the players have tighter distributions of initial propensities (smaller σ_i). Note that such an effect is consistent with the intuition that with information the players start the experiment with a clearer notion of their preferred strategies.

Figure 6 (right side) presents the prediction of a variant of the model with variability and $\sigma_i =$

Table 5
Model Fitness Scores for Experiment 2 in Table 5's Format

	Without information				With information			
Model and fitness score	P(freq)	ln β	d'	Total	P(freq)	ln β	d'	Total
Original parameters								
Exp. out	4/10	3/10	5/10	12/30	8/10	9/10	2/10	19/30
Model out	5/10	5/10	5/10	15/30	10/10	9/10	2/10	22/30
r(model, exp)	.75*	.78*	.08		.89*	.78*	.26	
Exp. out	1/10	1/10	3/10	5/30	8/10	6/10	0/10	14/30
With variability								
Model out	2/10	3/10	3/10	8/30	10/10	9/10	1/10	20/30
r(model, exp)	.77*	.68*	.32		.90*	.74*	.42	
With variability and $\sigma_i = .75$								
Exp. out	0/10	0/10	4/10	4/30	2/10	1/10	3/10	6/30
Model out	0/10	1/10	5/10	6/30	3/10	5/10	4/10	12/30
r(model, exp)	.76*	.66*	.39		.89*	.77*	.43	

Note. P(freq) = proportion of frequent responses; exp. = experimental observation.
$^*p < .05$.

.75, which is half the original value. This column and Table 7 (bottom columns) show that with the smaller σ_i the virtual subjects behave similarly to the players who know the payoff rule. Most important, the lower σ_i, like the information in the experiment, slowed the learning process or led to higher $\ln \beta$. In addition, it increased the estimated d'. The information in the experiment had a similar effect that did not reach significance.

In line with the results of Experiment 1, Table 5 shows positive correlations between the model's predictions and the observed statistics. Only the correlations for d' did not reach significance.[5]

Finally, again in line with the results of Experiment 1, Table 5 shows that between-subject parameter variability tend to improve the model's fitness scores. A similar conclusion was reached by Erev and Roth (1995).

General Discussion

The present research demonstrates that the effect of base-rate information on decision making cannot always be described by considering the optimal response rule (by means of Bayes' theorem) and the experimentally observed deviation from it. In natural situations in which interdependence between decision makers is likely, Bayes's theorem cannot be used to compute the rational decision. Moreover, in these situations, multiple equilibria may exist. For example, under certain conditions, all consensus games have at least two equilibria.

However, the fact that rational decision theory does not allow unique predictions of behavior does not imply that decisions are random. Rather, the present article demonstrates that behavior in two-person consensus games can be predicted from an understanding of the reinforcement structure and the learning process. In line with the predictions derived from Erev et al.'s (1995) learning model, an extreme base-rate effect was observed in consensus games that involve two decision makers who evaluate independent information.

The attempt to go beyond Bayes's theorem, which distinguishes the present research from most previous base-rate studies, is only one indication of a deeper difference between the approach taken here and the traditional judgment and decision-making approach. The present research was designed within the framework of cognitive game theory (Erev & Roth, 1995; Roth & Erev, 1995). Cognitive game theory relies on previous results in judgment and decision-making research but replaces some of the basic implicit assumptions of the traditional framework.

Most importantly, traditional experimental research implicitly accepts the framework of rational (high) game or decision theory. Under this framework, it is possible to distinguish between the economic and the psychological determinants of behaviors. The economic factors are modeled by a game, and the psychological factors are modeled by expected utility theory with certain additions. According to Savage's (1954) influential version of expected utility theory, the psychology of decision making can be summarized by two processes: probability assessments and choice among gambles. Moreover, the probability assessment process is expected to be consistent with Bayes's theorem. Whereas previous experimental studies demonstrate that the specific assumptions of rational decision theory are often not descriptive, the general framework has been kept. The utilization of Bayes's theorem as a benchmark for the study of base-rate effects is an example of the implicit acceptance of the rational framework.

Under cognitive game theory, psychological processes play a major role in decision making, and psychology is not limited to the choice among gambles and the assessment of subjective probabilities. Rather, in a cognitive game theoretical analysis, three factors have to be considered in order to predict choice behavior:

First, the set of strategies considered by the players has to be abstracted. The identification of the relevant strategies and the initial propensities is usually based on experimental results. In the

[5] The relative inaccuracy of the predicted d' scores appears to characterize the two experiments. Parameter fitting might improve the model's predictions even for the no-information condition. Yet, given the high predicted and the observed variance (large CIs), we believe that a larger data set is needed to perform this task.

present research, following results from signal detection research, we assumed that players utilize cutoff strategies.

A second abstraction involves the payoff rule. The specific payoff structure that applies to a situation has to be defined. Payoff here may depend on the state of the world and the actions of other people (e.g., the consensus with them in our study). The combination of strategies and associated payoffs gives rise to the abstraction of the "cognitive game," that is, the expected payoff for each strategy set. In the context of two-person signal detection games, it is the normal form presentation of the game. (For an explicit presentation of the game, see Erev et al., 1995.)

Finally, the way by which players select among the different strategies in the cognitive game has to be modeled. Traditional approaches assume that people try to select the strategy that maximizes expected utility. In contrast, the present approach distinguishes between long-term (possibly equilibrium) predictions and intermediate (learning-curve) predictions. In the present setting, the learning model predicts convergence toward one of the game's equilibria, but in different settings there may not be such a convergence.

Our results demonstrate the potential of cognitive game theory for the analysis of consensus games as created in an experiment. However, this approach also has much wider implications. We are all constantly engaged in consensus games where the outcome of our actions depend at least in part on consensus with other people. For instance, most decisions in organizational settings have to be approved by a number of people, and participants are motivated to reach consensus with others who are involved in the decision process. This situation can be modeled as a consensus game; from the results of our study, we may expect a strong weighting of base-rate information. For example, assume that there are two applicants, one belonging to a group whose members are usually hired when they apply for a job and another person who belongs to a group whose members have been less likely to be hired. In this case, we expect to find a strong tendency to chose the person belonging to the more readily accepted group. Hence, considering hiring decisions as consensus games leads to the prediction of discrimination against certain groups in the population.

Another set of situations for which the current model applies involves detection of malfunctioning in a system by a team of workers. The model predicts that an incentive to reach consensus will lead teams to ignore (miss) rare signals.

In addition to the identification of a consensus effect that can bias decisions, two specific nontrivial implications are directly derived from the current experimental results. First, the problem is expected to increase as decision makers gain experience. Thus, practice will not have a positive effect in this setting, and the severity of the problem cannot be assessed by examining the behavior of inexperienced decision makers.

A second direct implication is related to the effect of explicit information. Our results suggest that trying to hide the problematic game (e.g., by not informing the decision makers that they are evaluated based on between-judge agreements) might increase the bias.

Other implications can be derived from the quantitative model supported in the current research. This model can be used to derive quantitative predictions of the effect of the incentive to reach consensus in different settings. Given an abstraction of the incentive structure and the player's sensitivity, the model predicts the expected decision criteria as a function of experience. Most important, the model can be used to identify the conditions in which the bias is likely to be particularly strong and to suggest possible debiasing techniques. For example, the model suggests (in simulations that have not been reported and studied here) that two major factors are likely to increase the bias: a decrease in the observers sensitivity and an increase in the ratio of consensus-to-accuracy payoffs.

The examples considered above are only a small set of social situations that can be analyzed in terms of consensus games specifically or cognitive game theory generally. Specifying the strategies, payoff structure, and learning model that apply in a social interaction allows the researcher to predict intermediate term outcomes in this interaction through applying a minimum of basic principles and may be a valuable tool for the understanding of decision making under various circumstances.

References

Ajzen, I. (1977). Intuitive theories of events and the effects of base rate information on prediction. *Journal of Personality and Social Psychology, 35,* 303–314.

Bar-Hillel, M. (1980). The base rate fallacy in probability judgments. *Acta Psychologica, 44,* 211–233.

Beyth-Marom, R., & Fischhoff, B. (1983). Diagnosticity and pseudodiagnosticity. *Journal of Personality and Social Psychology, 45,* 1185–1195.

Birnbaum, M. H. (1983). Base rates in Bayesian inference: Signal detection analysis of the cab problem. *American Journal of Psychology, 96,* 85–94.

Bornstein, G., Erev, I., & Goren, H. (1994). Learning processes and reciprocality in intergroup conflicts. *Journal of Conflict Resolution, 38,* 690–707.

Bush, R., & Mosteller, F. (1955). *Stochastic models for learning.* New York: Wiley.

Erev, I. (1995, August). *A cognitive game theoretical model of binary categorization decisions under uncertainty.* Paper presented at the Conference on Subjective Probability Utility and Decision Making (SPUDM 15), Jerusalem, Israel.

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,* 360–375.

Erev, I., Maital, S., & Or-Hof, O. (in press). Melioration, adaptive learning and the effect of constant re-evaluations of strategies. In G. Antoniedes, F. van Raaij, & S. Maital, (Eds.), *Advances in economic psychology.* Sussex, England: Wiley.

Erev, I., & Roth, A. (1995, June). *Low versus high game theory, part II: Experimental investigation of learning and adaptive behavior.* Paper presented at Games 1995: International Conference on Game Theory and Applications, Jerusalem, Israel.

Fischhoff, B., & Bar-Hillel, M. (1984). Diagnosticity and the base-rate effect. *Memory & Cognition, 12,* 402–410.

Gigerenzer, G., Hell, W., & Blank, H. (1988). Presentation and content: The use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 513–525.

Gilat, S. (1995). *Hahashpaa shel tagmul avur haskama bemis'chakei gilui otot.* [The effect of the incentive to reach consensus in detection games]. Unpublished masters thesis, Technion, Israel Institute of Technology, Haifa, Israel.

Green, D. M., & Swets, J. A. (1988). *Signal detection theory and psychophysics.* Los Altos, CA: Peninsula Publishers. (Original work published in 1966)

Harley, C. B. (1981). Learning the evolutionary stable strategy. *Journal of Theoretical Biology, 89,* 611–633.

Healy, A. F., & Kubovy, M. (1981). Probability matching and the formation of conservative decision rules in a numerical analog of signal detection. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 344–354.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under certainty: Heuristics and biases.* Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review, 80,* 237–251.

Kubovy, M., Rapoport, A., & Tversky, A. (1971). Deterministic vs. probablistic strategies in detection. *Perception & Psychophysics, 9,* 427–429.

Lee, W. (1963). Choosing among confusably distributed stimuli with specified likelihood ratios. *Perceptual & Motor Skills, 16,* 445–467.

Lyon, D., & Slovic, P. (1976). Dominance of accuracy information and neglect of base rates in probability estimation. *Acta Psychologica, 40,* 287–298.

Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and short-comings in social judgment.* Englewood Cliffs, NJ: Prentice-Hall.

Ochs, J. (1995). Games with unique mixed strategy equilibrium: An experimental study. *Games and Economic Behavior, 10,* 202–217.

Rapoport, A., Seale, D., Erev, I., & Sundali, J. A. (in press). Coordination success in market entry games: Tests of equilibrium and adaptive learning models. *Management Science.*

Roth, A., & Erev, I. (1995). Learning in extensive form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior, 8,* 3–24.

Savage, L. J. (1954). *The foundation of statistics.* New York: Wiley.

Sorkin, R. D., & Dai, H. (1994). Signal detection analysis of the ideal group. *Organizational Behavior and Human Decision Processes, 60,* 1–13.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monographs, 8(Suppl. 2),* 1–109.

Tversky, A., & Kahneman, D. (1974, September 27). Judgment under uncertainty: Heuristics and biases. *Science, 185,* 1124–1131.

Tversky, A., & Kahneman, D. (1980). Causal schemata in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp. 49–72). Hillsdale, NJ: Erlbaum.

Tversky, A., & Kahneman, D. (1982). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases.* (pp. 153–160). Cambridge, England: Cambridge University Press.

Von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research.* Cambridge, England: Cambridge University Press.

Appendix A

Two-Person Signal Detection Model

The two-person signal detection model is an extension of classic signal detection theory (SDT) for the case of two independent detectors. Like traditional SDT, the model assumes that in each trial, Player i observes a stimulus (x_i) that was randomly selected from one of two distributions: $N_i(.)$ or $S_i(.)$. Player i is assumed to set a cutoff c_i, which implies $\beta_i = S_i(c_i)/N_i(c_i)$ value, and to consistently respond "S" if x_i is larger than c_i. Erev et al. (1995) showed that when the two players are independent, the equilibrium cutoffs (for $\infty > \beta > 0$) satisfy the following set of equations:

$$\frac{S_1(c_1)}{N_1(c_1)} = \frac{CC_1 - FC_1 + (CF_1 - CC_1 - FF_1 + FC_1)\int_{c_2}^{\infty} N_2(x)\,dx}{HH_1 - MH_1 + (HM_1 - HH_1 - MM_1 + MH_1)\int_{-\infty}^{c_2} S_2(x)\,dx}$$

and

$$\frac{S_2(c_2)}{N_2(c_2)} = \frac{CC_2 - FC_2 + (CF_2 - CC_2 - FF_2 + FC_2)\int_{c_1}^{\infty} N_1(x)\,dx}{HH_2 - MH_2 + (HM_2 - HH_2 - MM_2 + MH_2)\int_{-\infty}^{c_1} S_1(x)\,dx},$$

where C = correct rejection, F = false alarm, H = hit, and M = miss. In addition, the game can also have extreme equilibria ($\beta \rightarrow 0$, $\beta \rightarrow \infty$). The equilibrium calculation can be greatly simplified when symmetrical consensus games of the type studied here (where HH1 = CC1 > HM1 = MH1 = MM1 = FF1 = CF1 = FC1) are considered. The following hold for these games: (a) All equilibria are symmetrical (see Gilat, 1995), and (b) At all nonextreme equilibria ($0 < \beta < \infty$), Player i is indifferent between the two responses ("N" and "S"), thus,

$$P_i(N|(c_j = x_i))P_i(\text{“}N_j\text{”}|(x_i, c_j = x_i)) = P_i(S|(c_j = x_i))P_i(\text{“}S_i\text{”}|(x_i, c_j = x_i)),$$

where $P_i(Z|(c_j = x_i))$ is the probability that i assigns to the event that the state is Z (S or N) given a signal x_i that equals to her or his cutoff, and $P_i(Z_j|(x_i, c_j = x_i))$ is the probability that i assigns to the event that j will respond Z given i's signal and the assumption that j uses the same cutoff (when j's signal is not known to i). because

$$P_i(S|(c_j = x_i)) = 1 - P_i(N|(c_j = x_i)),$$

and

$$P_i(\text{“}S_j\text{”}|(x_i, c_j = x_i)) = 1 - P_i(\text{“}N_j\text{”}|(x_i, c_j = x_i)).$$

Simple algebra implies that

$$P_i(N|(c_j = x_i)) = P_i(\text{“}S_j\text{”}|(x_i, c_j = x_i)).$$

Appendix B

A Reinforcement-Based Learning Rule

Erev et al.'s (1995) adaptation of Roth and Erev's (1995) learning rule can be summarized by the following four assumptions (A).

A Finite Number of Uniformly and Symmetrically Distributed Cutoffs

A1q: Player i considers a finite set of m cutoffs. The location of cutoff j ($1 \le j \le m$) is $c_j = c_{\min} + \Delta(j-1)$. Erev et al. defined $\Delta = 2(c_{\max})/(m-1)$ and set the two strategy-set parameters to $m = 101$ and $c_{\max} = 5$.

Initial Propensities

A2q.: At Time $t = 1$ (before any experience), Player i has an initial propensity to choose the jth cutoff.

Two initial propensities parameters were set: S(1) and σ_i. To set these parameters, Erev et al. (1995) defined $q_{ij}(1) = p_k(1)S(1)$, where $S(1) = \Sigma_j^{m=1} q_j(t)$ and $p_k(1)$ is the probability that cutoff k will be chosen at the first round. They assumed that S(1) = 3 * (the average absolute profit in the game, given randomly selected cutoffs) and determined $p_k(1)$ by the area above cutoff k under a normal distribution with a mean at the center of the two distributions and SD $\sigma_i = 1.5$.

Reinforcement, Generalization, and Forgetting

The learning process is the result of the updating of the propensities through reinforcement, generalization, and forgetting.

A3q: If cutoff k was chosen by Player i at Time t and the received payoff was v, then the propensity to set cutoff j is updated by setting

$$q_j(t+1) = \max\,[\upsilon, (1-\varphi)q_j(t) + G_k(j, R, (v))],$$

where υ is a technical parameter that ensures that all propensities are positive, φ is a forgetting parameter, $G_k(.,.)$ is a generalization function, and $R(.)$ is a reinforcement function.

Erev et al. (1995) set $\upsilon = .0001$ and $\varphi = .001$. The reinforcement function was set to $R_t(v) = v - \rho(t)$, where ρ is a reference point that is determined by the following contingent weighted averge adjustment rule

$$\rho(t+1) = \begin{cases} \rho(t)(1-w^+) + v(t)(w^+), & \text{if } v(t) \ge \rho(t); \\ \rho(t)(1-w^-) + v(t)(w^-), & \text{if } v(t) < \rho(t), \end{cases}$$

where $w^+ = .01$ and $w^- = .02$ are the weights by which positive and negative reinforcements affect the reference point.

The generalization function was set to

$$G_k(j, R(v)) = R(v)(F\{[c_j + c_{j+1}]/2\} - F\{[c_j + c_{j-1}]/2\}),$$

where $F\{.\}$ is a cumulative normal distribution with mean c_k and $SD = \sigma_g$. Erev et al. set $\sigma_g = .25$.

The Relative Propensities Sum

The final assumption states the choice rule. **A4q:** The probability that the observer sets strategy k at time t is determined by the relative propensities sum

$$P_{ik}(t) = q_{ik}(t) \Big/ \left[\sum_{j}^{m=1} q_{ij}(t)\right].$$

Received September 11, 1995
Revision received October 15, 1996
Accepted October 22, 1996 ■