JOURNAL OF MATHEMATICAL PSYCHOLOGY 39, 360–375 (1995)

Toward a Generalization of Signal Detection Theory to N-Person Games: The Example of Two-Person Safety Problem*

IDO EREV, DANIEL GOPHER, REVITAL ITKIN, AND YAAKOV GREENSHPAN

Technion—Israel Institute of Technology

Game theoretical analysis is used to extend the predictions of classical signal detection theory to situations in which the detection task is performed by two observers. The implications of the suggested generalization are demonstrated and examined in the context of a two-person safety game. In a controlled experiment, 28 dyads performed a letter classification task. In accordance with the predictions, naive subjects were less likely to detect critical signals when the responsibility of detection was shared, and the magnitude of the criteria shift depended on the subjects' sensitivity (d'). It is shown that a simple reinforcement-based learning model can approximate the observed results. This model is equally applicable to both the classical one-person and the two-person signal detections tasks.

Signal detection theory (SDT), as introduced by Green and Swets (1966), has been successfully used to analyze, prescribe, and predict behavior in almost all domains of psychological research. Among the distinct domains are psychophysics (e.g., Falmagne, 1986), vigilance (e.g., David & Parasuraman, 1982), attention (e.g., Bonnel, Stein, & Bertucci, 1992), subjective judgment (e.g., Ferrel & McGoey, 1980; Wallsten & Gonzalez-Vallejo, 1994), eyewitness testimony (e.g., Buckhout, 1974; Birnbaum, 1983), jury decisions (e.g., Kerr, 1993), and medical diagnosis (e.g., Swennen, Hessel, & Herman, 1977).

We contend that the exceptional success of SDT stems from the fact that it captures many of the important characteristics of human behavior in the face of uncertain or incomplete information. Yet, an important factor is neglected by classical SDT; the theory ignores the strategic (game theoretical) considerations that may affect signal detection when several observers interact. As a result, the theory's prescriptive and descriptive power decrease as strategic considerations become important. As demonstrated by Sperling and Dosher (1986), traditional SDT can be thought of as a corollary of Savage's (1954) subjective expected utility theory, which is limited to one decision maker.

As an example of the important of strategic interaction between observers in a typical signal detection task, consider a version of the example provided by Swets and Green (1978) concerning a driver who enters the highway and has to decide whether the opening in the traffic is wide enough. Strategic considerations are important in this example because the wideness of the opening can be changed by the behavior (strategy) of other drivers, which, in turn, may be sensitive to the strategy chosen by the entrant. Thus, the optimal strategy for the entrant cannot be calculated independently of the strategy of the other drivers. Indeed, a utility maximizer entrant is expected to be sensitive to the decision problem faced by the other drivers, as well as his/her own decision problem.

The present paper represents a first step in an attempt to generalize the prescriptive and descriptive predictions of SDT to situations in which the strategic aspects of the situation are important.[1] The major goal of this paper is to demonstrate that a generalization of SDT to n-person games can make this useful theory even more useful. No attempt is made here to provide a complete generalization as the task is too complex. Rather, the paper presents a formal description of a two-person signal detection problem, and then focuses on one example: a two-person safety problem—a situation in which accidents can be avoided if at least one of two individuals detects the risky situation and takes the necessary safety measures.

To appreciate the challenge of incorporating strategic considerations into SDT, note that the effect of interactions between observers cannot be added linearly to the effects of the factors considered by traditional SDT. Rather, interaction between observers alters the nature of the optimal behavior and its calculation altogether. Relations that hold in the traditional one-observer task may not hold when more than one observer is considered; whereas strategies

* This work was supported by a grant from the committee for Research and Prevention in Occupational Safety and Health in the Israeli Labor Ministry. Correspondence and reprint requests should be addressed to Ido Erev, Faculty of Industrial Engineering and Management, The Technion, Haifa, 3200 Israel. Fax: 972-4-235194. E-mail: erev@techunix.technion.AC.IL.

[1] Kalai (1993) provides a related game theoretical analysis of signal detection task. In his analysis there is one observer (a juror) who plays a game with the person who produces a signal (a prospective criminal).

0022-2496/95 $12.00

(decision criteria) can be studied independently of perceptual abilities (sensitivities) in classical SDT, decision criteria are likely to be affected by sensitivity parameters when more than one observer is involved. For example, in the highway scenario a driver may be more cautious when he or she has reason to believe that the other drivers are less likely to see the danger than he or she does, but less cautious given the belief that others see the possible danger better.

The paper proceeds as follows: The first section summarizes the predictions of traditional SDT. It is then demonstrated that when a second observer (we use the words *observer, detector, decision maker, operator,* and player interchangeably) is added, the detection task can be abstracted as a two-person game. The nature of the predicted behavior is described using game theoretical concepts. The game theoretical model is then used to explore the example of a two-person safety dilemma. It is shown that, in equilibrium, observers who share the responsibility of detecting warning signals are less likely to report signals. This predicted trend (criterion shift) is particularly strong for the less capable (less sensitive) observer. An experiment is then described that examines the game theoretical predictions. The results exhibit the general trends predicted by the game theoretical model. The psychological process by which behavior could converge toward the game theoretical equilibrium is discussed. It is demonstrated that a simple learning model that assumes that observers update their behavior in accordance with the law of effect (Thorndike, 1898) is sufficient to reproduce the results. Thus, the suggested generalization of SDT can be based on the same psychological process that underlies classical SDT. Theoretical and practical implications of the results are discussed.

TABLE 1a

The Four Outcomes (Utilities) of a One-Person Signal Detection Task

Player 1's decision	State of nature: S	State of nature: N
"S"	*U*(Hit)	*U*(FA)
"N"	*U*(Miss)	*U*(CR)

TABLE 1b

A Two-Person Signal Detection Game

Player 1's decision	Player 2's decision "S" — State of nature: S	Player 2's decision "S" — State of nature: N	Player 2's decision "N" — State of nature: S	Player 2's decision "N" — State of nature: N
"S"	HH_1, HH_2	FF_1, FF_2	HM_1, HM_2	FC_1, FC_2
"N"	MH_1, MH_2	CF_1, CF_2	MM_1, MM_2	CC_1, CC_2

Note. The bottom left entry in each cell is the notation of Player 1's utility, and the top right entry in each cell is the notation of Player 2's utility.

TABLE 1c

Player 1's Expected Utilities at Equilibrium of a Two-Person Signal Detection Game

Decision	State of nature: S	State of nature: N
"S"	$(Pm_2)(HM_1)+(1-Pm_2)(HH_1)$	$(Pf_2)(FF_1)+(1-Pf_2)(FC_1)$
"N"	$(Pm_2)(MM_1)+(1-Pm_2)(MH_1)$	$(Pf_2)(CF_1)+(1-Pf_2)(CC_1)$

THE MODEL

One-Person SDT

For the sake of completeness, at the risk of some redundancy, we will briefly review the basic postulates of SDT. Classical SDT considers situations in which one observer needs to detect signals in a noisy environment and/or perception system. Thus, the observer has two responses "S" (I saw a signal) and "N" (no signal), while there are two relevant states of the world ("signal and noise" and "only noise"). Because the perceived signal includes noise, the observer cannot be 100% accurate; rather, the four contingencies presented in Table 1a are possible. The observer can correctly detect a signal (hit, in common SDT terminology), correctly say "N" (correct rejection—CR, say "N" when a signal was presented (miss), or say "S" when there was no signal (false alarm—FA).

The simple analysis suggested by SDT was found to have useful applications on three distinct levels. For the engineer, SDT is a prescriptive model—a technique to solve engineering problems (e.g., Peterson, Birdsall, & Fox, 1954). For the data analyst, SDT is a data analysis model—a statistical tool that allows independent assessment of sensitivity (d') and a response criteria (β) (Macmillan & Creelman, 1991). For the behavioral decision theorist, SDT is a descriptive psychological model that can be used to predict behavior under certain assumptions.

According to the behavioral decision theory's interpretation of SDT, the decision to detect a signal (to say "S") is a function of the utilities of the four contingencies presented in Table 1a, the prior probabilities of the two states of nature, and the magnitude of the perceived signal (x). Specifically, a utility maximizer observer will consistently detect a signal if the expected utility from "S" is larger than the expected utility from "N," that is, if

$$U(\text{"S"}) > U(\text{"N"}),$$

where

$$U(\text{"S"}) = U(\text{Hit})\, P(\text{S} \mid x) + u(\text{FA})\, P(\text{N} \mid x)$$

and

$$U(\text{"N"}) = U(\text{Miss})\, P(\text{S} \mid x) + u(\text{CR})\, P(\text{N} \mid x).$$

When $U(\text{"S"}) = U(\text{"N"})$, the observer is assumed to be indifferent between the responses "S" and "N". It is easy to show -(see, e.g., Coombs, Dawes, & Tversky, 1970) that calculation of $P(\text{N} \mid x)$ and $P(\text{S} \mid x)$ using Bayes' rule implies that a utility maximizer (rational) observer is indifferent between the two responses if and only if the ratio $P(x \mid \text{S})/P(x \mid \text{N})$ equals a certain decision criteria, referred to as β^*, specifically, if

$$\frac{P(x \mid \text{S})}{P(x \mid \text{N})} = \beta^* = \frac{U(\text{CR}) - U(\text{FA})}{U(\text{Hit}) - U(\text{Miss})} \frac{P(\text{N})}{P(\text{S})}. \quad (1)$$

The rational observer is expected to consistently report a signal when the ratio $P(x \mid \text{S})/P(x \mid \text{N})$ exceeds the cutoff β^*.

With the addition of specific assumptions concerning the distributions of the magnitude of the perceived signal given the two states of nature, the signal detection task can be presented graphically. Figure 1 is a graphical presentation of a signal detection task under the assumption that the perceived signal is sampled from one of two symmetrical normal distributions. It is selected from the noise distribution when a signal is not presented and from the signal and noise distribution in the face of a signal. The distance between the two distributions in standard deviation units (d') was arbitrarily assumed to equal two ($d' = 2$). The predicted rational behavior is for the operator to set a cutoff β^* and say "S" if and only if the perceived signal (x) falls to the right of the cutoff. The rationale behind this prediction will become clearer when a numerical example is discussed below (the numbers in Fig. 1 are part of this example).

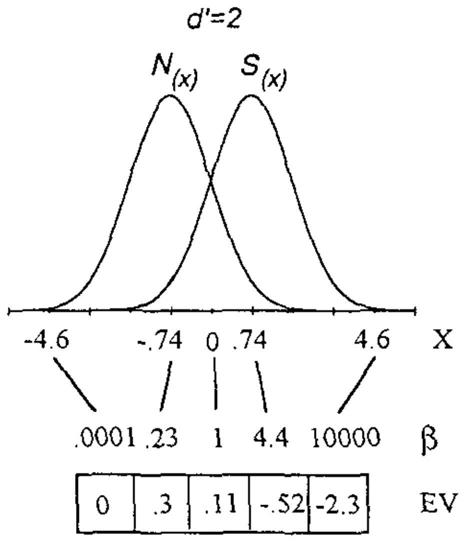

FIG. 1. Graphical presentation of a one-person signal detection task. X measures the distance from the signal to the intersection of the two distributions in standard deviation units. $N(x)$ describes the noise distribution, $S(x)$ is the signal and noise distribution. $\beta = S(x)/N(x)$ presents five possible cutoffs, and the EV box presents the expected outcomes when each of the five cutoffs is utilized.

Signal Detection as a Game—A Two-Person SDT

Consider now the case in which a second observer is added to the traditional signal detection task described above. To allow prediction of behavior in this setting precise assumptions concerning the way the two observers affect each other (structural assumptions), and the way the observers make their decisions (psychological assumptions) have to be added.

Previous examinations of signal detection by multiple observers (see reviews by Green & Swets, 1966, p. 248; Pete, Pattipati, & Kleinman, 1993 and Sorkin & Dai, 1994) assumed that multiple observers can be modeled as one decision maker that decides based on multiple observations. Whereas this assumption may hold in certain setting, coordinated decision process is not always possible. To address a wider set of situations, the coordination assumption is relaxed here. The present generalization retains classical SDT's main psychological assumption; it is assumed that observers try to maximize utility by setting optimal cutoff. Thus, each player makes one decision (chooses β), which implies one of two responses ("N" or "S") in each trial (game). In addition it is assumed that the observers are independent—given a known state of nature, knowledge of the magnitude of the signal perceived by Player i cannot be used to predict the magnitude of the signal perceived by Player j.[2] The decisions are assumed here to be made simultaneously. As demonstrated by Rapoport and his associates (Erev & Rapoport, 1990; Budescu, Rapoport, & Suleiman, 1992; Rapoport, Budescu, & Suleiman, 1993) the sequential case may differ from the simultaneous game. We plan to study the sequential protocol in future research.

[2] The independence assumption is likely to be violated in certain real-life situations. Theoretically, these situations can be addressed by the present model by dividing them into classes of subsituations in which conditional independence is likely to hold. Additional research is needed in order to evaluate the practical value of this solution.

Since the observers' behavior is limited to two responses, the interaction between the two observers can be summarized by the payoff matrix presented in Table 1b, which determines the observers' utility given the state of the world and the two responses. The bottom left and the top right entries in each of Table 1b's cells present Player 1's and Player 2's utilities, respectively. Note that the chosen notation retains a reference to classical SDT notation and that no symmetry constrains are assumed. For example, "HM_1," which reflects Player 1's utility when he or she hits while Player 2 misses, does not have to equal HM_2 or MH_2.

Figure 2 graphically presents the two-person signal detection task under the assumption of symmetrical normal distributions. In this graphical example, Player 2 is assumed to be more sensitive ($d'_2 = 3$) than Player ($d'_1 = 2$). Note that under specific assumptions concerning the signal distributions and the relevant utilities (Table 1b), it is possible to calculate the players' expected outcomes given each set of decision criteria. The set of possible expected outcomes presented in Fig. 2 is part of a numerical example, discussed below.

The calculation required in order to find the strategy that maximizes utility in the two-person case differs from the calculation required in the one-person case in three main respects. First, an assumption needs to be added that describes Player i's belief concerning his/her partner's strategy. In line with classical game theory, it is assumed here that Player i believes that his/her partner will choose the best response to his/her own strategy. That is, strategies will be in equilibrium—no player will be able to improve his/her payoff (utility) by unilaterally changing his strategy (decision rule). Formally, it is assumed that:

> Strategies will be in equilibrium—each Player i will choose the response rule (β_i^*) that maximizes his/her expected payoff given the other player's response rule.

A second difference between maximization in the one-person and the two-person signal detection tasks relates to the complexity of the required calculation. This quantitative difference is implied by the qualitative difference discussed above. In order to find the equilibrium strategy analytically, a set of two equations with two unknowns needs to be solved (assuming $0 < \beta^* < \infty$).

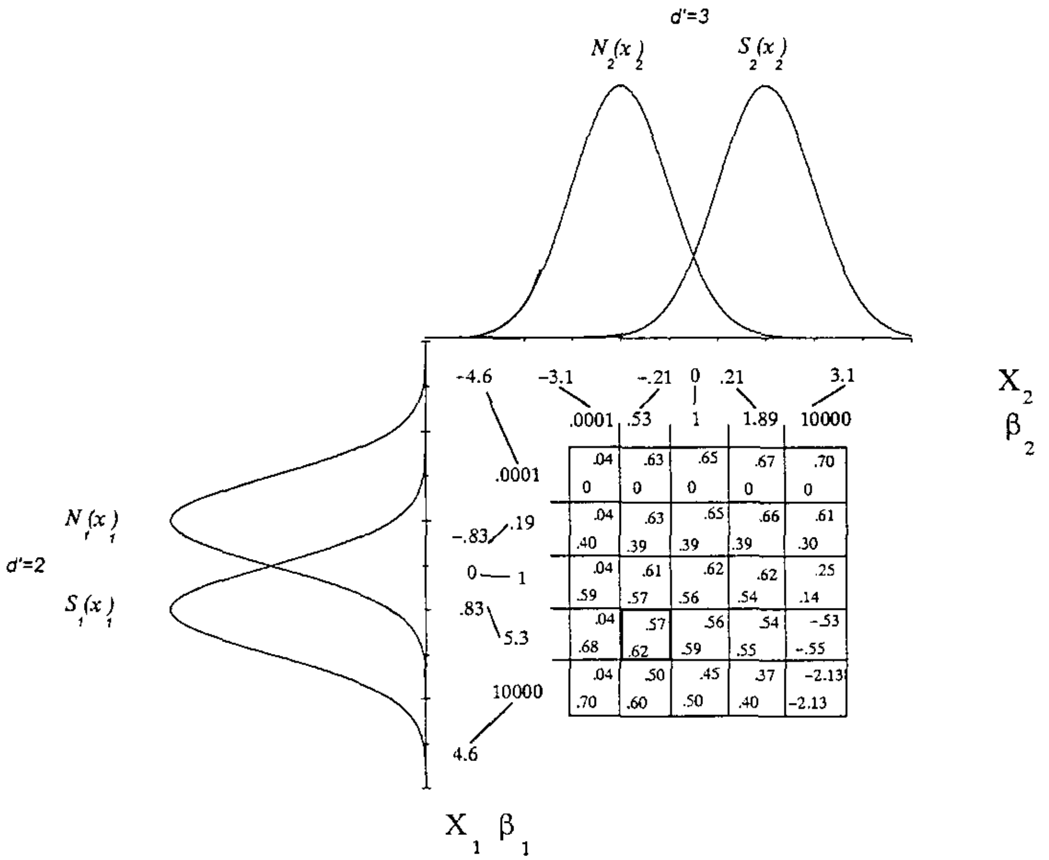

FIG. 2. Graphical presentation of a two-person signal detection task. Player 1 is the weaker player ($d'_1 = 2$) and Player 2 is the stronger ($d'_2 = 3$). The central matrix presents expected values of two players in the normal form of a simplified game in which each player considers five cutoffs ($\beta_1 = 0.0001$, 0.19, 1, 5.3, and 10,000; $\beta_2 = 0.0001$, 0.53, 1, 1.89, and 10,000). The simplified game has a unique equilibrium: $\beta_1 = 5.3$, $\beta_2 = 0.53$.

Let Pm_i and Pf_i be the probabilities of a miss and a false alarm, respectively, by Player i at equilibrium. Given Pm_2 and Pf_2, Player 1's maximization problem is identical to the one-person problem with modified expected utilities. Table 1c presents the modified problem. As can be seen in this table, the expected utilities are

$$\mathrm{EU(Hit')} = (Pm_2)(\mathrm{HM}_1) + (1 - Pm_2)(\mathrm{HH}_1)$$
$$\mathrm{EU(Miss')} = (Pm_2)(\mathrm{MM}_1) + (1 - Pm_2)(\mathrm{MH}_1)$$
$$\mathrm{EU(FA')} = (Pf_2)(\mathrm{FF}_1) + (1 - Pf_2)(\mathrm{FC}_1)$$
$$\mathrm{EU(CR')} = (Pf_2)(\mathrm{CF}_1) + (1 - Pf_2)(\mathrm{CC}_1).$$

Player 1's optimal response to Player 2's strategy (β_2^*) can be calculated directly from Eq. (1) by substituting Table 1a's utilities with Table 1c's expected utilities:

$$\beta_1^* = \frac{\begin{pmatrix}[(Pf_2)(\mathrm{CF}_1) + (1 - Pf_2)(\mathrm{CC}_1)] \\ -[(Pf_2)(\mathrm{FF}_1) + (1 - Pf_2)(\mathrm{FC}_1)]\end{pmatrix}}{\begin{pmatrix}[(Pm_2)(\mathrm{HM}_1) + (1 - Pm_2)(\mathrm{HH}_1)] \\ -[(Pm_2)(\mathrm{MM}_1) + (1 - Pm_2)(\mathrm{MH}_1)]\end{pmatrix}} \frac{P(\mathrm{N})}{P(\mathrm{S})}.$$

Namely,

$$B_1^* = \frac{\begin{pmatrix}\mathrm{CC}_1 - \mathrm{FC}_1 \\ +(\mathrm{CF}_1 - \mathrm{CC} - \mathrm{FF}_1 + \mathrm{FC}_1)(Pf_2)\end{pmatrix}}{\begin{pmatrix}\mathrm{HH}_1 - \mathrm{MH}_1 \\ +(\mathrm{HM}_1 - \mathrm{HH}_1 - \mathrm{MM}_1 + \mathrm{MH}_1)(Pm_2)\end{pmatrix}} \frac{P(\mathrm{N})}{P(\mathrm{S})}. \tag{2}$$

Let c_i be the magnitude of Player i's cutoff signal at equilibrium, the signal that satisfies the equality $\beta_i^* = P(c_i \mid \mathrm{S})/P(c_i \mid \mathrm{N})$, and let the functions $N_i(\cdot)$ and $S_i(\cdot)$ describe the height of the signal distributions perceived by Player i given noise and signal and noise, respectively. That is, $N_i(c_i)$ and $S_i(c_i)$ return the height of the signal distribution at point c_i. Note that

$$\beta_i^* = S_i(c_i)/N_i(c_i).$$

The probability that Player 2 will miss a true signal (Pm_2) at equilibrium is determined by c_2 and the distribution of the signals perceived by Player 2 given a signal (the function $S_2(\cdot)$); Pm_2 is the area to the left of c_2 below S_2:

$$Pm_2 = \int_{-\infty}^{C_2} [S_2(x)]\, dx.$$

In a similar fashion:

$$Pf_2 = \int_{c_2}^{\infty} [N_2(x)]\, dx.$$

Note that Eq. (2) describes the condition under which Player 1 responds optimally to Player 2's strategy. At equilibrium, both players maximize utility given their opponent's strategy. The equation that describes Player 2's best response is a mirror image of the equation that describes Player 1's best response. The two equations are presented below with the relevant probabilities and response rules expressed by the cutoff points:

$$\begin{aligned}\frac{S_1(c_1)}{N_1(c_1)} &= \frac{\begin{pmatrix}\mathrm{CC}_1 - \mathrm{FC}_1 + (\mathrm{CF}_1 - \mathrm{CC}_1 \\ -\mathrm{FF}_1 + \mathrm{FC}_1)\int_{C_2}^{\infty} [N_2(x)]\, dx\end{pmatrix}}{\begin{pmatrix}\mathrm{HH}_1 - \mathrm{MH}_1 + (\mathrm{HM}_1 - \mathrm{HH}_1 \\ -\mathrm{MM}_1 + \mathrm{MH}_1)\int_{-\infty}^{C_2} [S_2(x)]\, dx\end{pmatrix}} \frac{P(\mathrm{N})}{P(\mathrm{S})} \\ \frac{S_2(c_2)}{N_2(c_2)} &= \frac{\begin{pmatrix}\mathrm{CC}_2 - \mathrm{CF}_2 + (\mathrm{FC}_2 - \mathrm{CC}_2 \\ -\mathrm{FF}_2 + \mathrm{CF}_2)\int_{C_1}^{\infty} [N_1(x)]\, dx\end{pmatrix}}{\begin{pmatrix}\mathrm{HH}_2 - \mathrm{HM}_2 + (\mathrm{MH}_2 - \mathrm{HH}_2 \\ -\mathrm{MM}_2 + \mathrm{HM}_2)\int_{-\infty}^{C_1} [S_1(x)]\, dx\end{pmatrix}} \frac{P(\mathrm{N})}{P(\mathrm{S})}.\end{aligned} \tag{3}$$

That is, under the present assumptions a rational player is expected to solve (or behave as if he or she solves) the set of two equations presented above when selecting a strategy in a two-person signal detection task.

Examination of Eq. (3) reveals a third important difference in the calculation of the optimal strategy in one-person and two-person signal detection tasks. In order to compute the integrals in (3), Player i has to know (or be sensitive to) the exact signal distributions of both players. Such knowledge implies understanding of the distributions and the sensitivities (d'_1 and d'_2)—factors that play no role in the solution of the one-person signal detection task (Eq. (1)).

It is important to note that, whereas the logic that led to the game theoretical calculation was based on a description of Player i's beliefs, a top-down reasoning is not necessary. It has been shown that behavior can converge to equilibrium through a bottom-up learning process (e.g., Harley, 1981; Maynard Smith, 1984; Roth & Erev, 1995; Selten, 1991). As will be shown below, the present prediction can hold even if Player i is not aware of the fact that he or she has a partner, but simply learns from experience in an adaptive fashion. Indeed, the experiment described below focuses on the simplest case in which subjects do not know that they have partners.

We turn now to a comparison of examples of one-person and two-person signal detection tasks in order to demonstrate and evaluate the usefulness of the present analysis.

ONE-PERSON AND TWO-PERSON SAFETY PROBLEMS

One-person safety problem—definition. A one-person safety problem is defined here as a signal detection task in which the detection of a signal, the decision "S," amounts to taking safety measures, whereas the decision "N" implies neglect. Whenever the operator fails to detect a signal an accident occurs. That is, $U(\text{Miss})$ is the cost of an accident.

A numerical example of a safety problem. Consider an operator who works in a setting (safety problem) in which the costs of an accident is $U(\text{Miss}) = -10$, the benefit from correct rejection (e.g., a production of a unit when the system works well) is $U(\text{CR}) = 1$, and detection preserves the status quo, $U(\text{Hit}) = U(\text{FA}) = 0$. The prior probabilities are $P(\text{no signal}) = 0.70$, $P(\text{signal}) = 0.30$. These parameters are summarized in Table 2a.

If the observer is rational (behaves in accordance with Eq. (1)), he or she is predicted to set a decision rule (a cutoff) β^* and to detect a signal if and only if

$$\frac{P(x \mid \text{S})}{P(x \mid \text{N})} > \beta^* = \frac{(1-0)}{(0-(-10))}\frac{0.70}{0.30} = 0.233.$$

The table in Fig. 1 presents the expected outcome given five cutoffs. In line with the formal derivation, β^* maximizes the observer's expected outcome.

Two-person safety problem. In an attempt to abstract an important strategic aspect of real-world safety problems (see Battmann & Klump, 1993) while keeping the analysis simple, only minimal interaction between the two observers is assumed here. It is assumed that an accident is avoided if at least one of the two observers detects the risky situation and takes the necessary safety measures (chooses "S"). When an accident is avoided, both observers are assumed to receive the same payoff, $U(\text{Hit})$, regardless of their decision. The cost of the accident per observer is kept at $U(\text{Miss})$ as in the one-person problem. That is,

$$\text{HH}_1 = \text{HH}_2 = \text{HM}_1 = \text{HM}_2 = \text{MH}_1 = \text{MH}_2 = U(\text{Hit}),$$
$$\text{MM}_1 = \text{MM}_2 = U(\text{Miss}),$$
$$\text{FF}_1 = \text{FF}_2 = \text{FC}_1 = \text{CF}_2 = U(\text{FA}), \quad \text{and}$$
$$\text{CC}_1 = \text{CC}_2\text{CF}_1 = \text{CF}_2 = U(\text{CR}).$$

Table 2b presents the parameters of the two-person version of the one-person safety problem presented in Table 2a. Using these equalities and Eq. (2) we get

$$\beta_1^* = \frac{U(\text{CR}) - U(\text{FA})}{[U(\text{Hit}) - U(\text{Miss})]\, Pm_2}\frac{P(\text{N})}{P(\text{S})},$$

and using (1) and the players' symmetry,

$$\beta_i^* = \beta^*/Pm_j. \quad (4)$$

Namely, if $Pm_j < 1$ (the probability that player j will fail to detect a signal is below 1), β_i^* is expected to be larger than β^*. The first observation is therefore qualitative:

In equilibrium a utility maximizer player is less likely to take safety measures (say "S") in a two-person safety problem than in a one-person situation given the same parameters.

It is important, however, to note that this observation does not imply that an equilibrium point always exists. Indeed, the existence and the exact values of the equilibria are likely to depend on the specific parameters. We turn now to the derivation of the value of the equilibria under specific assumptions.

Note first that, given the present constraints on the values of the utilities, set (3) is simplified to

$$\int_{-\infty}^{C_2} [S_2(x)]\, dx = [N_1(c_1)/S_1(c_1)]\, \beta^*$$
$$\int_{-\infty}^{C_1} [S_1(x)]\, dx = [N_2(c_2)/S_2(c_2)]\, \beta^*. \quad (5)$$

When the functions $S_i(\cdot)$ and $N_i(\cdot)$ are assumed to be density functions of normal distributions, the solution (c_1 and c_2) cannot be derived analytically. It can only be approximated numerically. We approximated the numerical solution under the common assumption that $S_i(\cdot)$ and

TABLE 2a

A Numerical Example of a One-Person Safety Problem

		State of nature	
		S	N
	Prior probability	0.30	0.70
Decision	"S"	0	0
	"N"	−10	1

TABLE 2b

A Two-Person Variant of the Safety Problem

(Each cell: upper-right value; lower-left value.)

		Player 2's decision "S"		Player 2's decision "N"	
	State of nature	S	N	S	N
	Prior probability	0.30	0.70	0.30	0.70
Player 1's decision	"S"	0; 0	0; 0	0; 0	1; 0
	"N"	0; 0	0; 1	−10; −10	1; 1

$N_i(\cdot)$ are symmetrical normal distributions with variance $=1$ and with the means $-d'_i/2$ and $+d'_i/2$, respectively.

The solutions obtained for combinations of three levels of d'_i (1.5, 2.5, and 3.5) are presented in Table 3.

As can be seen in Table 3, all six distinct two-person safety problems have a unique equilibrium prediction in pure strategies. When the two players are equally sensitive ($d'_1 = d'_2$), the equilibrium point is symmetric ($\beta^*_1 = \beta^*_2$). In all other cases the more sensitive observer will be more cautious ($\beta^*_1 < \beta^*_2$ if and only if $d'_1 > d'_2$); the less capable operator is predicted to "feel free to let his friend do the job" (a similar prediction was made by Rapoport (1988) in the context of step-level public goods problems). Note also that in accordance with the qualitative observation, presented above, the equilibria cutoffs are larger than the predicted cutoff in the one-person case (0.23). The shift represents a decrease in the tendency to take safety measures.

These theoretical results imply that, under the present model, when the relevant parameters of the safety problem are known an exact and unique prediction of the observers' strategies is possible. According to the predicted trend when two operators work together they will tend to be less careful. This tendency will be particularly strong for the less capable worker. This shift can also be observed in the numerical example presented in Fig. 2. When a finite number of cutoffs are considered, the game can be reduced to a normal form game. A simplified game in normal form in which each player has five cutoffs is presented in Fig. 2. This game has a unique equilibrium in which the weaker player is less careful ($\beta^*_1 = 5.3$) than the stronger player ($\beta^*_2 = 0.53$).

In line with the present predictions social psychologists (see, e.g., Latane' & Darley, 1968, and a review by Dovidio, Piliavin, Gaertner, Schroeder, & Clark, 1991) have observed a "diffusion of responsibility" effect. For example, subjects in Latene' and Darley's (1968) experiment were more likely to detect smoke that was pumped into a "waiting room" when they were alone in the room; 75% of the subjects who were alone in the room reported this emergency, and only 12% of the subjects who were seated in a room with two confederates did the same. The present model provides a parsimony description of sufficient conditions to the occurrence of diffusion of responsibility. In addition, it provides new quantitative predictions concerning the effect of relative perceptual sensitivity on emergency detection.

The experiment described below was designed to test the descriptive power of the model's predictions in a minimalist signal detection setting. In the studied setting, observers do not know in advance how their payoff is determined. They are told only what the correct responses are and receive payoff feedback after each response. Thus, the experiment examines whether experience can lead observers to converge to the game theoretical predictions.

Since the speed at which bottom-up processes lead behavior to equilibrium is likely to depend on the players initial propensities,[3] two order conditions will be examined. Half the subjects will play a two-person SDT game first and then (without an explicit notice) will play a one-person game; the remaining subjects will play the one-person game first. Examination of the second game that each player plays allows evaluation of the sensitivity of the predictions' descriptive power to the observers' initial propensities.

TABLE 3

β Values at Equilibrium of Numerical Examples of Two-Person Safety Problems Given Different d' Values

	$d'_2 = 1.5$		$d'_2 = 2.5$		$d'_2 = 3.5$	
	β^*_1	β^*_2	β^*_1	β^*_2	β^*_1	β^*_2
$d'_1 = 1.5$	1	1	4.5	0.39	13.5	0.28
$d'_1 = 2.5$			1.6	1.6	7.1	0.72
$d'_1 = 3.5$					3.03	3.03

EXPERIMENT

Method

A combined within- and between-subject design was utilized to compare behavior in the one-person and the two-person safety problems described above (Tables 2a and 2b).

Subjects. Fifty-six male Technion students participated in the experiment. They were recruited by advertisements promising monetary reward for participation in a computerized experiment. The exact payoffs were contingent on performance and ranged from 12–40 Shekels ($4.5–$15). Subjects were divided into 28 pairs according to the order in which they responded to the advertisements.

Apparatus and stimuli. The experiment was programmed and run using a micro experimental laboratory system. This system was installed on a personal computer with a Super VGA 19" screen.

Figure 3 presents the experimental task and the time course of an experimental trial. Members of a dyad sat side by side in front of the screen, sharing the same keyboard on which two response keys were defined for each of them. Special carton dividers were used to hide the performance results window (the upper left and right corners of the screen) of one subject from the other. An experimental trial began with the appearance of a black fixation cross in the

[3] Roth and Erev (1995) have demonstrated that the speed of convergence to equilibrium is determined by an interaction of the game with the players' initial propensities. In certain games, behavior appears to converge to equilibrium in several rounds independently of the players' initial tendencies, whereas the learning path in other games is extremely sensitive to behavior in the initial rounds. Given certain initial propensities, convergence in these "sensitive" games may not be reached even after one million rounds.

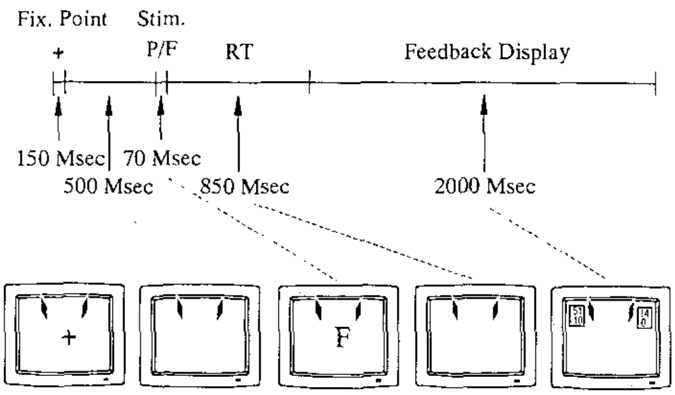

FIG. 3. The experimental setting and the time course of an experimental trial.

middle of the screen for 150 ms. After a 500-ms. interval, a single letter (5 mm in height) appeared at the same location, and was displayed for 70 ms. (because the subjects saw the stimuli from a distance of about 50 cm, the visual angle was about 0.57°). The letter "P" appeared in 70% of the trials, and the letter "F" appeared in the remaining 30%. To further increase the difficulty of letter discrimination, letters were presented in light blue on a grey screen.

Instruction and task. Subjects were told that the "P" stood for "perfect", meaning that the system was functioning as it should; while "F" represented "failure", implying a malfunctioning of the system. The correct response for each situation was defined explicitly. In the case of perfect ("P") detection, the subject on the left was instructed to press the ⟨1⟩ key of the digit raw while the subject sitting on the right was told to press the ⟨8⟩ key of the key pad. When "F" was detected, the left subject was told to press the ⟨2⟩ key and the right subject was told to press the ⟨9⟩ key. Response interval, for both operators, was limited to 850 ms.

Separate performance feedback was given to each player, in the upper corner of the screen on his respective side of the screen. Each operator could see only his own score. Feedback included the scores of the last trial and the sum total from the beginning of the game and was displayed for 2 s (2000 ms.). The initial score for each player before the first trial was set to 50 points. To reduce the possibility that a subject would ignore the feedback, a 4-s instruction to look at the result window appeared in the center of the screen after every ten trials. Subjects were given 2-min. break every 100 trials.

Subjects were given information only on the appropriate responses (to press ⟨1⟩ or ⟨8⟩ given "P", and ⟨2⟩ or ⟨9⟩ given "F"). They were not told that there were two payoff conditions (individual and team) with different payoff matrixes. Moreover, they were not told that there was a dependency between the dyad members; namely, that in the two-person condition, the payoff that one player received was determined by both responses. The payoff rules could be revealed only through the feedback scores. To ensure full understanding of the instructions, each dyad was given five training trials before the beginning of the first experimental block.

Experimental design. Two payoff conditions (one-person and two-person) were compared in a within-subject (within-pair) design. Each condition consisted of 200 trials blocked in two blocks of 100 trials. The order of the two conditions was counterbalanced across pairs (there were two, between-subject order conditions).

The two experimental conditions were manipulated by different payoff matrixes; Player i's payoff was calculated in accordance with Table 2a under one-person conditions and in accordance with Table 2b under two-person conditions.

Results and Discussion

Table 4 presents the main experimental results over subjects by experimental condition, order condition, and block (the Appendix presents a summary of the raw data by

TABLE 4

Summary of the Main Experimental Results

	One-person first		Two-person first	
	First experimental condition			
Condition:	One-person		Two-person	
Block:	1st	2nd	1st	2nd
P(Hit):	0.85 (0.1)	0.89 (0.11)	0.82 (0.13)	0.84 (0.12)
P(FA):	0.14 (0.1)	0.13 (0.11)	0.11 (0.09)	0.09 (0.09)
d':	2.43 (0.76)	2.8 (0.78)	2.55 (0.85)	2.72 (0.84)
log(β):	0.14 (0.50)	−0.08 (0.74)	0.35 (0.56)	0.37 (0.6)
β:	1.15	0.93	1.41	1.45
	Second experimental condition			
Condition:	Two-person		One-person	
Block:	3rd	4th	3rd	4th
P(Hit):	0.87 (0.12)	0.86 (0.16)	0.92 (0.09)	0.93 (0.07)
P(FA):	0.10 (0.09)	0.10 (0.1)	0.12 (0.14)	0.13 (0.16)
d':	2.83 (0.63)	2.78 (0.8)	3.02 (0.84)	2.94 (0.85)
log(β):	0.12 (0.99)	−0.02 (1.03)	−0.21 (0.78)	−0.26 (0.75)
β:	1.13	0.98	0.81	0.77

Note. Means and standard deviations (in parentheses). Geometrical means were computed for the β values.

subject). Recall that each subject pair participated in both experimental conditions. The top part of Table 4 presents the results of the first half of the experimental session (blocks 1 and 2), while the lower part of the table presents the results of the second half of the session.

Recall that, for the purpose of payoff calculation, slow responses were defined as incorrect responses. That is, when a subject did not respond within the time limit he received an "error" payoff. About 2% of the responses were too slow and fell within this category. We conducted the analysis twice: once by treating missing values as errors and once by excluding them from the analysis. Because the two analyses had similar outcomes, only the more common analysis that neglects missing values is presented here.

P(Hit) is the proportion of "F" responses given a presentation of an F. *P*(FA) is the proportion of "F" responses given a presentation of a P. Examination of Table 4 reveals that subjects had both higher hit rates and higher false alarm rates under one-person conditions than under two-person conditions.

The experimental hypotheses address the behavior of experienced subjects. Thus, only the results obtained in the last block played in each condition (blocks 2 and 4) were utilized for hypothesis testing. As subjects within a dyad cannot be considered to be independent, we used the dyad as the unit of analysis. The statistical tests reported below examine if the within-dyad trend is consistent with the model's prediction and if it is affected by the order condition.

In order to evaluate the working hypothesis which states that d' (players sensitivity) is not affected by the payoff manipulation, a d' score was calculated to each

TABLE 5

Summary of the Main Experimental Results by the Subjects' Relative Abilities

	One-person first				Two-person first			
	First experimental condition							
Condition:	One-person				Two-person			
Block:	1st		2nd		1st		2nd	
Relative ability:	S	W	S	W	S	W	S	W
P(Hit):	0.92	0.78	0.96	0.81	0.90	0.74	0.95	0.73
	(0.07)	(0.18)	(0.05)	(0.2)	(0.12)	(0.2)	(0.07)	(0.22)
P(FA):	0.08	0.19	0.06	0.20	0.05	0.16	0.07	0.11
	(0.08)	(0.17)	(0.07)	(0.19)	(0.06)	(0.15)	(0.09)	(0.11)
d':	2.98	1.88	3.41	2.19	3.15	1.96	3.28	2.15
	(0.75)	(1.01)	(0.49)	(1.19)	(0.89)	(1.14)	(0.71)	(1.18)
$\log(\beta)$:	0.12	0.18	−0.1	0.02	0.28	0.35	−0.11	0.46
	(0.55)	(0.55)	(0.88)	(0.85)	(0.49)	(0.84)	(0.74)	(0.73)
β:	1.13	1.19	0.91	1.02	1.32	1.43	0.9	1.59
	Second experimental condition							
Condition:	*Two-person*				*One-person*			
Block:	3rd		4th		3rd		4th	
Relative ability:	S	W	S	W	S	W	S	W
P(Hit):	0.94	0.80	0.95	0.77	0.97	0.87	0.97	0.89
	(0.08)	(0.22)	(0.1)	(0.29)	(0.04)	(0.14)	(0.04)	(0.1)
P(FA):	0.05	0.15	0.05	0.16	0.06	0.19	0.06	0.21
	(0.07)	(0.12)	(0.07)	(0.17)	(0.1)	(0.23)	(0.08)	(0.26)
d':	3.42	2.25	3.43	2.12	3.52	2.52	3.49	2.39
	(0.62)	(0.91)	(0.71)	(1.15)	(0.72)	(1.19)	(0.56)	(1.21)
$\log(\beta)$:	0.08	−0.01	0.06	−0.16	−0.12	−0.07	−0.16	−0.04
	(0.9)	(1.06)	(0.78)	(0.97)	(0.54)	(1.06)	(0.72)	(0.82)
β:	1.08	0.99	1.06	0.85	0.88	0.93	0.85	0.96

Note. The S columns present the statistics for the stronger player in each pair, and the W columns present the weaker player's statistics.

dyad.[4] The value of d' appears to increase with experience, but was not affected significantly by the order and the experimental condition. No significant change in the value of d' between the two conditions was observed ($t[27]=0.87$, repeated measures t-test, NS). Whereas this null result is consistent with the underlying assumption of the traditional SDT model (and with the present generalization), it is important to note here that it is possible that the rational players may be motivated to reduce their sensitivity in certain two-person signal detection games. Future research should focus on this possibility.

In order to test the first experimental hypothesis, which states that experience leads to higher β criteria under two-person conditions than under one-person conditions a $\log(\beta)$ score was calculated to each dyad under each condition. ($\text{Log}(\beta)$ was used rather than β, since β measures a ratio and for that reason has a skewed distribution.) In line with the model's prediction, following experience (in blocks 2 and 4) the average $\log(\beta)$ was lower (-0.17) under one-person conditions under two-person conditions (0.17) ($t[27]=1.89$, $p<0.05$ one tail repeated measures t-test). Although the "order by payoff condition" interaction was insignificant, it should be noted that support for the experimental hypothesis was particularly strong for dyads who performed the two-person task first. These dyads reduced their response rule from $\log(\beta)$ score of 0.37 in block 2 to -0.26 in block 4 ($t[13]=3.18$, $p<0.01$). A similar, but statistically insignificant, trend was observed for dyads who performed the one-person task first.

A between-subject analysis agrees with the within-subject analysis. A significant payoff condition effect was observed in the second block ($t(26)=1.75$, $p<0.05$ one tail test), but not in the fourth block.

In order to examine the effect of the observers' relative sensitivity on the response rule, a d' score was first calculated for each subject, and the subject with the highest d' in each pair was defined as the stronger subject. Then, Table 4's statistics were calculated for the 28 stronger and 28 weaker subjects in each block separately. Table 5 presents these results.

Inspection of Table 5 suggests that our second experimental hypothesis, like the first, holds when the two-person condition is played first. A significant "order by relative sensitivity" interaction was observed ($F[1,26]=5.21$, $p<0.05$). In the second block of the two-person condition the weaker players had a significantly higher ($t[13]=2.28$, $p<0.05$) $\log(\beta)$ score (0.46) than the stronger players (-0.11). The effect of the players' relative ability was insignificant in the forth block ($t[13]=0.89$, NS).

The fact that our experimental hypotheses were not supported in the fourth block suggests that the descriptive power of the equilibrium prediction is sensitive to the players' initial tendencies. No consistent trend was observed when the players played the two-person game after participation in the one-person game. The present data cannot be used to provide a clear account for this order effect. Among the possible explanations that we plan to examine in future research are: (1) Players who had high hit rate under one-person conditions could not have noticed that the payoff rule was changed. (2) Players with low d' who lost money under one-person conditions tried different strategies in an attempt to improve their positions, whatever strategy they chose in the beginning of the third block was successful, and for that reason they kept using it. Inspection of Appendix 1 suggests that the order effect may also be related to the higher variance in the distribution of missing values in the fourth block.

[4] Following Healy & Kubovy (1977), in all computations of SDT statistics the extreme standard scores of hit and falls alarm rates were taken as -2.00 or $+2.00$. That is, rates above 0.977 were replaced by 0.977, and rates below 0.023 were replaced by 0.023.

GENERAL DISCUSSION

The present results suggest that strategic considerations can, and should, be incorporated into SDT. It is shown that when a detection task involving two detectors is abstracted as a two-person game, the predicted behavior can be derived using game theoretical concepts. The effect of the social structure on the predicted (and prescribed) behavior is substantial. In equilibrium of the two-person safety problem, studied here, observers who share the responsibility to detect warning signals are less likely to report signals (relative to the one-person case). Moreover, the less capable observer is predicted to free ride on the effort of his/her coworker. Experimental results suggest that naive subjects exhibit the general trends predicted by the game theoretical model.

It should be reemphasized here that the safety game is only one example of a two-person signal detection game. The specific results obtained for this game suggest that the present approach can be used to predict behavior in two-person signal detection games, but do not imply that similar behavioral trends are to be expected in different games. The difference between one- and two-person tasks is expected to vary from game to game. For example, in another two-person signal detection task, the "consensus" game (Erev, Gopher, Meyer & Gilat, 1994), the present model correctly predicts overweighing of the base rates rather than reckless behavior. We believe the present extension of SDT provides a useful abstraction, and predications for many situations in which perceptual decisions are affected by the social context.

The Psychological Process

Recall that we speculated that the exceptional success of classical SDT stems from the fact that it captures many of the robust characteristics of human behavior in the face of uncertain or incomplete information. It seems that the psychological process implied by classical SDT holds (or at least provides a good approximation of the actual process) in many settings. Can the same be said for the present generalization? That is, can we be confident that the suggested game theoretical generalization of SDT provides a useful description of behavior in domains other than the two-person safety dilemma studied here? While a clear answer to this question requires additional research, some insight is provided by an examination of the psychological process that may give rise to the present pattern of behavior. The following discussion demonstrates that, although the calculations implied by the two models are quite distinct, the same basic process can lead behavior to converge to the predictions of both classical and generalized SDT.

A Reinforcement-Based Learning Rule

What learning rule might account for the present results? Several learning models were proposed to describe learning processes in one-person signal detection tasks (for a review and evaluations of the different models see Kubovy & Healy, 1977, and Busemeyer & Myung, 1992). Whereas these models can be extended to describe two-person games, we chose to take a more general approach and utilized the reinforcement-based learning (RBL) model that was suggested by Roth and Erev (1995). This model was found to provide a useful approximation of behavior in a wide set of experimental games (see Roth & Erev, 1995; Bornstein, Erev & Goren, 1994; Erev & Maital, in press; Ochs, 1995). The basic idea behind the model is the law of effect (Thorndike, 1898); it assumes that the probability that a certain strategy will be adopted increases when this strategy is positively reinforced. Similar models were suggested by Harley (1981) to describe animal learning processes and by Bush and Mosteller (1955).

The RBL model is adapted to the present game under the assumption that the set of strategies available to the players is a set of the possible cutoffs. The adapted model's basic assumptions are presented below.

Finite number of uniformly and symmetrically distributed cutoffs. According to the first assumption Player i considers a finite number of m cutoffs. The m cutoffs are uniformly distributed along the interval $[\text{cut}_{\min}, \text{cut}_{\max}]$, where $\text{cut}_{\min}$ and $\text{cut}_{\max}$ are the two extreme cutoffs. The distribution is assumed to be symmetrical around $c = 0$; that is, $\text{cut}_{\min} = -\text{cut}_{\max}$, and the distance between two adjacent cutoffs is $\Delta = 2(\text{cut}_{\max})/(m-1)$. Thus,

B1. Player i considers a finite set of m cutoffs. The location of cutoff j $(1 \leqslant j \leqslant m)$ is $\text{cut}_j = \text{cut}_{\min} + \Delta(j-1)$.

Initial propensities. The model assumes that the Player i starts the experiment with a certain tendency to choose each of the possible cutoffs.

B2. At time $t = 1$ (before any experience) Player i has an initial propensity $q_{ij}(1) > 0$ to choose his jth cutoff.

Reinforcement, generalization, and forgetting. The learning process is assumed to be a function of updating the propensities through reinforcement, generalization, and forgetting.

B3. It 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

$$q_{ij}(t+1) = \text{MAX}[\upsilon, (1-\varphi)\, q_{ij}(t) + G_k(j, R(\text{v})],$$

where $\upsilon > 0$ is a technical parameter that ensures that all the propensities are positive, ϕ is a forgetting parameter, $G_k(\cdot, \cdot)$ is a generalization function, and $R(\cdot)$ is a reinforcement function.

Roth and Erev noted that the shape of the reinforcement and the generalization functions should be determined based on experimental results. Most importantly, Hernstein's (1961) demonstration of linear relation between reinforcements and choice probabilities suggests that a linear reinforcement function can provide a good approximation. Thus, we choose the function

$$R(\text{v}) = \text{v} - \rho,$$

where ρ is a reference point parameter; outcomes that are larger than ρ are positive reinforcements whereas outcomes that are smaller than ρ are negative reinforcements.

Experimental investigation of generalization (e.g., Guttman & Kalish, 1956) suggests that strategies that are similar to the selected strategy will be affected by the reinforcement. Specifically, Guttman and Kalish (1956) observed a normal generalization distribution. To approximate a normal generalization distribution we assume that

$$G_k(j, R(\text{v})) = R(\text{v})(F_k\{[\text{cut}_j + \text{cut}_{j+1}]/2\} - F_k\{[\text{cut}_j + \text{cut}_{j-1}]/2\}],$$

where $F_k\{\cdot\}$ is a cumulative normal distribution with mean cut_k and standard deviation σ_g.

Relative propensities sum. The final assumption states the choice rule.

B4. The probability that player i sets strategy k at time t is determined by the relative propensities sum:

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

The model's parameters. In addition to the four general assumptions presented above, the derivation of the model's prediction requires specific assumptions to set the model's parameters. Since the goal of the present discussion is a mere demonstration that the RBL model can reproduce the present data, certain arbitrary parameter choices were made here.

The two strategy-set parameters were set to $m = 101$ and $\text{cut}_{\max} = 5$. These values were chosen to ensure that the interval $[\text{cut}_{\min}, \text{cut}_{\max}]$ is "wide enough" and the number of cutoffs is "large enough." Wide and large enough are defined here as values that can be increased without affecting the learning process in a meaningful way (sensitivity analysis supports this conjunction).

To eliminate the number of free initial propensities parameters Roth and Erev use the fact that the initial propensity to choose strategy k can be written as $q_{ik}(1) = P_k(1)\, S(1)$, where $P_k(1)$ is the covert (and assessable) choice probability and $s(1) = \sum_{j=1}^{m} q_j(t)$ is a strength parameter. Following Erev and Roth (1994) we estimate $S(1)$ as the expected absolute reinforcement sum from three random decisions and used experimental results (Lee & Janke, 1964) to assess the values of $P_k(1)$. Under the present model, the results of Lee and Janke imply that $P_k(1)$ can be approximated by a normal distribution with a mean at the center of the two distributions and standard deviation $\sigma_i = 1.5$.

Following Roth and Erev's convention the forgetting rate parameter was set to $\phi = 0.001$. The technical parameter in assumption B3 was introduced ensure that all the propensities are positive. This parameter's exact value does not affect the learning process as long as it is sufficiently small. It was set to $\upsilon = 0.000001$.

Since the subjects did not know the payoff rule at the beginning of the experiment, the reference point had to be acquired during the learning phase. In line with Erev and Roth, the present paper assumes that players start the experiment with a natural reference point ($\rho = 0$) and then adjust it based on their payoff. Specifically, the reference point in trial $t+1$ is assumed to be a weighted average of the reference point in trial t and the profit obtained in this trial (v):

$$\rho(t+1) = \begin{cases} (1-w)\,\rho(t) + (w)\,\mathrm{v} & \text{if } \mathrm{v} < \rho(t) \\ (1-w/2)\,\rho(t) + (w/2)\,\mathrm{v} & \text{if } \mathrm{v} \geqslant \rho(t). \end{cases}$$

Note that negative reinforcements are assumed to loom twice as much as positive reinforcements. This assumption is consistent with experimental results (see, e.g., Tversky & Kahneman, 1991). Overweighing of negative reinforcements is also necessary to ensure that the process will be consistent with the power law of practice. The weighting parameter was set to $w = 0.02$.

Finally, the generalization parameter was set to $\sigma_g = 0.025$. This small value was selected to ensure that with experience simulated subjects will converge to approximate the single cutoff hypothesis (Kubovy, Rapoport & Tversky, 1971).

Computer simulations. In order to compare the model's predictions to experimental results computer simulations were run. The simulations were run for 200 round of a direct replication of the first two blocks of the experiment. Half the simulations "played" the one-person condition, whereas the other half "played" the two-person game. At each round the following steps were taken:

(1) The state of the world (S with probability 0.3 and N otherwise) was randomly determined.

(2) The players' cutoffs (c_1 and c_2, which implies β_1 and β_2 values) were randomly determined in accordance with B4.

(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 response was S if and only if $x_i > c_i$).

(5) Profits were calculated using the relevant payoff table (Table 2a or 2b).

(6) Propensities were updated in accordance with B3.

(7) The reference point for the next trial was calculated.

Five hundred simulations (500 pairs of hypothetical players) were run, under each condition, in which the more capable player had $d' = 4$ and the less capable player had $d' = 2.5$.

Table 6 presents the simulations main results using the format of Table 5. As can be seen in this exhibit the simulated players showed the main trends observed in the experiment. Most importantly, in the second block they were "less careful" (had higher cutoffs) in the two-person problem than in the one-person problem. In this block the difference between the stronger and the weaker players moved toward the equilibrium prediction. Yet, the simulated players, like the experiment participants, did not converge to the exact quantitative predictions of the equilibrium analysis. The average inferred cutoffs in the second block were $\beta = 0.96$ for the more capable player and $\beta = 0.81$ for the less capable player in the one-person problem, and $\beta = 1.58$ and $\beta = 1.94$ for the stronger and the weaker player, respectively, in the two-person problem.

Evidence concerning the actual learning process. It is important to reemphasize that we do not contend to have evidence that our subjects updated their behavior in accordance with the RBL model. We intend to demonstrate only that the game theoretical generalization of SDT can be based on the same robust characteristics of behavior as classical SDT. Other learning rules can probably be used to provide an account of the present results.

Additional support for the suggestion that the same psychological process gave rise to the results under the two experimental conditions comes from informal discussions with subjects. Recall that our subjects were not told that they were to participate in experiments having two conditions. At the conclusion of the experiment, before any debriefing, the experimenter asked the subjects what they thought about the task. A few subjects said that one of the two halves (the two-person condition) was easier. One subject with a particularly low d' said that the computer was

TABLE 6

Summary of the Main Simulation Results by the Subjects' Relative Abilities

Condition:	One-person				Two-person			
Block:	1st		2nd		1st		2nd	
Relative ability:	S	W	S	W	S	W	S	W
P(Hit):	0.87 (0.06)	0.77 (0.07)	0.90 (0.06)	0.83 (0.08)	0.86 (0.08)	0.70 (0.11)	0.87 (0.07)	0.67 (0.13)
P(FA):	0.11 (0.04)	0.21 (0.07)	0.09 (0.04)	0.22 (0.07)	0.09 (0.04)	0.15 (0.06)	0.06 (0.04)	0.11 (0.05)
d':	2.47 (0.39)	1.56 (0.32)	2.73 (0.42)	1.79 (0.38)	2.55 (0.46)	1.61 (0.35)	2.77 (0.42)	1.76 (0.39)
log(β):	0.11 (0.49)	0.04 (0.26)	−0.04 (0.63)	−0.21 (0.46)	0.32 (0.59)	0.39 (0.37)	0.46 (0.69)	0.66 (0.49)
β:	1.11	1.04	0.96	0.81	1.38	1.48	1.58	1.94

Note. The difference between the estimated d' and the d' that were used to generate the data is a result of the probabilistic response rule. That is, the simulated subjects change their β from trial to trial.

more "forgiving" toward the end. Other subjects "noticed" that the letter "F" was more common in one of the two halves. A few subjects mentioned the possibility that their payoff was affected by their coworker's behavior, but not one of these subjects noticed that interdependence was important in only part of the experiment.

Thus, it seems safe to say that our subjects did not explicitly try to maximize distinct payoff rules under the two experimental conditions. In line of the logic behind the RBL model, the distinct outcomes (reinforcements) led them to select cutoffs that appear to be sensitive to the game theoretical considerations.

While a simple learning model was successful in approximating operators' behavior in the present experimental situation, it seems instructive to consider the psychological characteristics of the situation. Subjects were not informed of their experimental condition (one-person, two-person) nor of the exact payoff matrix. Hence, there was no top-down process that could have guided their strategic behavior. They were left to develop their own "model" of the world from direct experience with the task and from observing their performance scores. Under these conditions, the above depicted model seems to do well. We designed the experiment and introduced the constraints on information to specifically test the boundary conditions within which the initial influence of top-down processes and prior experience can be assumed to be limited. Our future experiments will be expanded to include systematic manipulations of advanced information and top-down strategies.

The Prescriptive Value

The support provided by the experimental results for the descriptive power of the suggested generalization of SDT validates the prescriptive power as well. That is, setting the equilibrium cutoff maximizes the player's utility only if the other players are likely to converge to the equilibrium. The observed convergence suggests that players should be motivated to behave in accordance with the prescription implied by the present generalization.

Note, however, that the fact that individual players should be motivated to follow the equilibrium strategy does not mean that this strategy is collectively optimal. As the prisoner dilemma game demonstrates, maximization at the individual player level does not necessarily imply efficient outcome for the team.[5]

CONCLUSIONS

As a first stage toward a generalization of SDT to n-person games, the present paper contains demonstrations rather than a systematic examination. It is demonstrated that: (1) strategic uncertainty often plays a role in signal detection; (2) signal detection tasks in which strategic uncertainty is important can be abstracted as games; (3) optimal behavior in two-person signal detection games can be derived using equilibrium analysis and is substantially distinct from optimal behavior in one-person tasks; (4) naive subjects appear to converge to the equilibrium predictions in a safety dilemma example of a two-person signal detection task; and (5) although the required calculations are quite distinct, the same basic process can lead behavior to converge to the predictions of both classical and generalized SDT. These demonstrations suggest that a generalization of SDT to n-person games can make this useful theory even more useful.

[5] Collectively optimal response rules in two-person signal detection task are considered and reviewed by Pete, Pattipati & Kleinman (1993) and Pados, Papantoni-Kazakos, Kazakos, & Koyiantis (1994).

APPENDIX: SUMMARY OF THE RAW DATA SUBJECT

Order condition: One-person first:

Pair	Player	H1	H2	H3	H4	F1	F2	F3	F4	MV1	MV2	MV3	MV4
1	L	0.94	0.92	0.97	1.00	0.03	0.03	0.14	0.06	1	2	8	42
	R	0.78	0.81	0.93	1.00	0.18	0.14	0.20	0.59	8	5	13	2
3	L	0.73	0.63	0.31	0.10	0.61	0.48	0.13	0.06	5	2	3	0
	R	0.91	0.88	1.00	1.00	0.06	0.02	0.01	0.01	2	2	2	1
5	L	1.00	1.00	0.96	1.00	0.04	0.03	0.00	0.00	0	0	0	0
	R	0.89	1.00	0.92	0.97	0.06	0.00	0.00	0.12	0	1	0	5
7	L	0.93	0.97	0.96	1.00	0.00	0.00	0.00	0.00	6	0	2	7
	R	0.84	0.97	0.63	0.93	0.03	0.02	0.00	0.00	0	1	0	0
9	L	0.91	1.00	1.00	0.93	0.13	0.18	0.07	0.00	2	0	1	0
	R	0.22	0.28	0.40	0.24	0.17	0.04	0.12	0.10	12	12	11	16
11	L	1.00	1.00	1.00	1.00	0.03	0.03	0.11	0.15	1	1	2	2
	R	0.90	0.93	1.00	1.00	0.01	0.01	0.00	0.00	2	0	1	1
13	L	0.84	0.80	0.94	0.97	0.18	0.36	0.18	0.19	6	0	3	0
	R	0.75	1.00	0.97	0.97	0.24	0.08	0.01	0.02	17	6	4	3
15	L	0.96	1.00	0.78	0.92	0.04	0.08	0.08	0.19	1	0	1	0
	R	0.92	0.97	0.85	0.62	0.18	0.39	0.38	0.45	2	1	0	0
17	L	0.93	1.00	0.90	0.82	0.04	0.00	0.03	0.04	5	3	3	5
	R	0.95	1.00	0.97	0.94	0.04	0.01	0.01	0.02	3	0	2	1
19	L	0.85	0.70	1.00	1.00	0.15	0.13	0.00	0.00	6	9	20	30
	R	0.91	1.00	1.00	1.00	0.06	0.00	0.05	0.00	4	3	1	1
21	L	0.76	0.92	0.81	0.78	0.05	0.05	0.03	0.07	1	0	1	4
	R	0.89	0.78	0.79	0.67	0.16	0.05	0.07	0.10	0	2	2	0
22	L	0.96	0.94	1.00	1.00	0.00	0.01	0.03	0.01	0	0	0	1
	R	0.74	0.93	0.90	0.86	0.45	0.55	0.28	0.23	1	1	1	1
24	L	0.81	0.88	0.81	0.68	0.26	0.08	0.03	0.04	3	1	0	1
	R	0.65	0.66	0.69	0.73	0.19	0.18	0.14	0.16	5	0	0	1
28	L	0.94	0.97	1.00	0.90	0.22	0.21	0.26	0.12	2	0	1	1
	R	0.88	0.91	1.00	1.00	0.26	0.36	0.36	0.21	1	0	1	1

Order condition: Two-person first:

Pair	Player	H1	H2	H3	H4	F1	F2	F3	F4	MV1	MV2	MV3	MV4
2	L	1.00	0.97	0.94	1.00	0.38	0.13	0.02	0.05	1	1	1	4
	R	0.92	0.97	1.00	1.00	0.00	0.04	0.06	0.05	0	0	0	0
4	L	0.64	0.94	0.96	0.93	0.04	0.03	0.00	0.07	27	6	3	9
	R	0.81	0.70	0.78	0.82	0.14	0.07	0.01	0.11	4	7	3	15
6	L	0.90	0.91	0.95	0.97	0.05	0.02	0.00	0.02	14	2	3	2
	R	0.97	1.00	0.98	0.94	0.03	0.03	0.12	0.00	2	1	0	3
8	L	1.00	0.97	1.00	1.00	0.00	0.02	0.00	0.03	0	0	2	1
	R	0.91	0.97	0.98	0.98	0.04	0.02	0.40	0.29	2	1	1	1
10	L	0.37	0.35	0.92	0.90	0.29	0.16	0.19	0.23	12	1	1	1
	R	1.00	1.00	1.00	1.00	0.01	0.00	0.01	0.00	0	4	2	1
12	L	0.43	0.52	0.58	0.69	0.42	0.39	0.29	0.35	19	13	19	19
	R	0.71	0.82	0.85	0.92	0.21	0.23	0.38	0.14	2	2	1	0
14	L	0.54	0.34	0.56	0.67	0.38	0.26	0.66	0.68	10	15	14	20
	R	0.97	0.97	0.97	0.85	0.03	0.08	0.03	0.03	2	0	2	0

16	L	0.86	0.50	0.80	0.88	0.07	0.04	0.07	0.14	0	2	1	4
	R	0.95	0.96	0.97	0.94	0.09	0.11	0.14	0.04	1	2	0	1
18	L	0.76	0.88	0.97	0.89	0.16	0.07	0.00	0.00	2	0	0	0
	R	1.00	0.97	1.00	1.00	0.02	0.01	0.03	0.01	1	0	1	0
20	L	0.62	0.79	0.97	0.89	0.02	0.00	0.01	0.03	2	1	0	1
	R	0.90	0.88	0.93	1.00	0.07	0.00	0.00	0.02	0	0	1	0
23	L	0.96	0.97	0.94	0.91	0.03	0.06	0.04	0.09	1	1	0	0
	R	1.00	1.00	1.00	0.97	0.00	0.00	0.00	0.02	0	1	0	0
25	L	0.78	0.79	0.96	1.00	0.10	0.06	0.00	0.09	14	3	1	1
	R	0.74	0.65	0.96	0.97	0.10	0.12	0.15	0.10	1	1	0	0
26	L	0.88	0.87	0.95	0.96	0.01	0.04	0.00	0.01	0	0	0	0
	R	0.96	1.00	1.00	1.00	0.03	0.07	0.00	0.00	0	1	2	1
27	L	0.56	0.80	0.96	0.93	0.16	0.17	0.64	0.86	6	4	3	6
	R	0.79	0.97	0.96	1.00	0.08	0.28	0.15	0.32	6	3	6	1

H1–H4 are the hit rates, F1–F4 are the false alarm rates, and MV1–MV4 are the number of missing values in the four blocks.

REFERENCES

Battman, W., & Klump, P. (1993). Behavioral economics and the compliance with safety regulations. *Safety Science*, **16**, 35–46.

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

Bonnel, A. M., Stein, J. F., & Bertucci, P. (1992). Does attention modulate the perception of luminance changes? *The Quarterly Journal of Experimental Psychology*, **44A**, 601–621.

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

Buckhout, R. (1974). Eyewitness testimony. *Scientific American*, **231** (6), 23–31.

Budescu, D. V., Rapoport, A., & Suleiman, R. (1992). Simultaneous vs. sequential requests in resource dilemmas with incomplete information. *Acta Psychologica*, **80**, 297–310.

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

Bush, R., & Mosteller, F. (1955). *Stochastic models for learning*. New York: Wiley. Coombs, C. H., Dawes, R. H., & Tversky, A. (1970). *Mathematical psychology: An elementary introduction*. Englewoord Cliffs, NJ: Prentice–Hall.

Davis, D. R., & Parasuraman, R. (1982). *The Psychology of Vigilance*. London: Academic Press.

Dovidio, J. F., Piliavin, J. A., Gaertner, S. L., Schroeder, D. A., & Clark, R. D., III. (1991). The arousal: cost-reward model and the process of intervention. A review of the evidence. In S. M. Clark (Ed). *Review of personality and social psychology: Prosocial behavior*. Newbury Park, CA Sage Publications.

Erev, I., Gopher, D., Meyer, J., & Gilat, S. (1994, November), *Signal detection in two person games: Implications to eye-witness behavior*. Paper presented at the 35th Annual meeting of the Psychonomics Society, St. Louis, MO.

Erev, I. & Maital, S. (in press). Melioration, adaptive learning and the effect of constant re-evaluation of strategies. In G. Antoniedes, F. van Raaij, & S. Maital (Eds.), *IAREP Conference Proceedings*. Dordrecht: Kluwer.

Erev, I., & Rapoport, A. (1990). Provision of step-level public goods: The sequential contribution mechanism. *Journal of Conflict Resolution*, **34**, 401–425.

Erev, I., & Roth, A. E. (1994, November). *On evolutionary game theory and its relation to judgment and decision making research*. Paper presented at the J/DM Society Meeting, St. Louis, MO.

Falmagne, J. C. (1986). Psychological measurement and theory. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance: Vol. 1. Cognitive processes and performance*. New York: Wiley.

Ferrel, W. R., & McGoey, P. J. (1980). A model of calibration for subjective probabilities. *Organizational Behavior and Human Performance*, **26**, 32–52.

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

Guttman, N., & Kalish, H. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology*, **51**, 79–88.

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

Healy, A. F., & Kubovy, M. (1977). A comparison of recognition memory to numerical decision: How prior probabilities affect cutoff location. *Memory and Cognition*, **5**, 3–9.

Herrnstein R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of Experimental Analysis of Behavior*, **4**, 267–272.

Kalai, E; (1993). A rational game theory framework for analysis of legal and criminal decision making. In R. Hastie (Ed.), *Inside the juror: The psychology of juror decision making*. Cambridge, UK: Cambridge Univ. Press.

Kerr, N. (1993). Stochastic models of juror decision making. In R. Hastie (Ed.), *Inside the juror: The psychology of juror decision making*. Cambridge, UK: Cambridge Univ. Press.

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**, 427–446.

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

Latane', B., & Darley, J. M. (1968). Group inhibition of bystander intervention in emergencies. *Journal of Personality and Social Psychology*, **10**, 215–221.

Lee, W., & Janke, M. (1964). Categorizing externally distributed stimulus samples for three continua. *Journal of experimental Psychology*, **68**, 376–382.

Macmillan, N. A., & Creelman, C. (1991). *Detection theory: A user guide*. Cambridge, UK: Cambridge Univ. Press.

Maynard Smith, J. (1984). Game theory and evolution of behavior. *Behavioral and Brain Sciences*, **7**, 95–125.

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

Pados, D. A., Papantoni-Kazakos, P., Kazakos, D., & Koyiantis, A. G. (1994). On-line threshold learning for Neyman–Pearson distributed detection. *IEEE Transactions on Systems, Man, and Cybernetics*, **24**, 1519–1531.

Pete, A., Pattipati, K. R., & Kleinman, D. L. (1993). Optimal team and individual decision rules in uncertain dichotomous situations. *Public Choice*, **75**, 202–230.

Peterson, W. W., Birdsall, T. G., & Fox, W. C. (1954). The theory of signal delectability, *Transactions of the IRE Professional Group of Information Theory*, **PGIT-4**, 171–212.

Rapoport, A. (1988). Provision of step-level public goods: Effects of inequality in resources. *Journal of Personality and Social Psychology*, **54**, 432–440.

Rapoport, A., Budescu, D. V., & Suleiman, R. (1993). Sequential requests from randomly distributed shared resources. *Journal of Mathematical Psychology*, **37**, 241–265.

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**, 164–212.

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

Selten, R. (1991). Evolution, learning, and economic behavior. *Games and Economic Behavior*, **3**, 3–24.

Sperling, G., & Dosher, B. A. (1986). Strategy and optimization in human information processing. In K. Boff, L. Kaufman, & J. Thomas (Eds.), *Handbook of perception and human performance: Vol. 1. Cognitive processes and performance*. 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.

Swennsen, R. G., Hessel, S. J., & Herman, P. G. (1977). Omission in radiology: Faulty search or stringent reporting criteria? *Radiology*, **123**, 563–567.

Swets, J. A., & Green, D. M. (1978). Applications of signal detection theory. In H. L. Pick, Jr., H. W. Leibowitz, J. E. Singer, A. Steinschneider, & H. W. Stevenson (Eds.), *Psychology: From research to practice* (pp. 311–331). New York: Plenum Press.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychology Monographs*, **2**.

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics*, **106**, 1039–1061.

Wallsten, T. S., & Gonzalez-Vallejo, C. (1994). Statement verification: A stochastic model of judgment and response. *Psychological Review*, **101**, 490–504.

Received: February 7, 1994