

Modeling the effects of payoff on response bias in a perceptual discrimination task: Bound-change, drift-rate-change, or two-stage-processing hypothesis

ADELE DIEDERICH

International University Bremen, Bremen, Germany

and

JEROME R. BUSEMEYER

Indiana University, Bloomington, Indiana

Three hypotheses—the *bound-change hypothesis*, *drift-rate-change hypothesis*, and *two-stage-processing hypothesis*—are proposed to account for data from a perceptual discrimination task in which three different response deadlines were involved and three different payoffs were presented prior to each individual trial. The aim of the present research was to show (1) how the three different hypotheses incorporate response biases into a sequential sampling decision process, (2) how payoffs and deadlines affect choice probabilities, and (3) the hypotheses' predictions of response times and choice probabilities. The two-stage-processing hypothesis gave the best account, especially for the choice probabilities, whereas the drift-rate-change hypothesis had problems predicting choice probabilities as a function of deadlines.

One of the hallmarks of signal detection theory is its ability to separate response bias from discriminability on the basis of the observed hit and false alarm rates (Green & Swets, 1966; Swets, Tanner, & Birdsall, 1961). According to signal detection theory, the response bias parameter is assumed to be influenced by the prior probability of the signal, as well as the payoffs for each type of error. Numerous empirical tests of these assumptions have been conducted over the past 40 years (see Maddox, 2002, for a review of this research). A major limitation of the classic signal detection model is that it provides a static (fixed sample) description of the decision process, and so it is unable to simultaneously account for choice probabilities and response times (RTs). Sequential sampling models (e.g., Laming, 1968; Link & Heath, 1975; Ratcliff, 1978) provide a dynamic extension of signal detection theory that elegantly accounts for the systematic relations between choice probabilities and RTs. Like the classic signal detection model, sequential sampling models also make a distinction between discriminability and response bias. Although the effects of prior probability have been examined (see, e.g., Green, Smith, & von Gierke, 1983), the effects of payoffs on response bias have rarely been

studied within the sequential sampling framework (Kornbrot, 1988). The purpose of the present research was to test three different hypotheses about the effects of payoffs on the sequential sampling decision process.

Sequential Sampling Models

Sequential sampling models seek to account for both RT and accuracy data. They assume that the stimuli (or choice alternatives) can be mapped onto a hypothetical numerical dimension representing the instantaneous level of activation. Furthermore, they assume some random fluctuation of this value over time in the course of information accumulation.

For a binary choice, say between A and B, they assume that the decision process begins with an initial state of evidence. This initial state may favor either Option A or Option B or may be neutral with respect to A or B. For simplicity, assume that a positive initial state favors Option A and a negative initial state favors Option B. Upon presentation of the stimulus, the decision maker sequentially samples information from the stimulus display over time. The small increments of information sampled at any moment in time are such that they favor either Option A or Option B. Again, assume that a positive increment indicates an increment favoring Option A and a negative increment is an increment favoring Option B. The evidence is accumulated from one moment in time to the next by summing the current state with the new increment. This process continues until the magnitude of the cumulative evidence exceeds a threshold criterion. The process

This research was supported by Deutsche Forschungsgemeinschaft Grant Di 506/6-2 to the first author. We are grateful to Hans Colonius for helpful comments. Correspondence should be addressed to A. Diederich, School of Humanities and Social Sciences, International University Bremen, P.O. Box 750 561, D-28725 Bremen, Germany (e-mail: a.diederich@iu-bremen.de).

stops and Option A is chosen as soon as the accumulated evidence reaches a criterion value for choosing A, or it stops and Option B is chosen as soon as the accumulated evidence reaches a criterion value for choosing B. The probability of choosing A over B is determined by the accumulation process reaching the threshold for A before reaching the threshold for B. The criterion is assumed to be set by the decision maker prior to the decision task.

The initial starting position of the process reflects prior beliefs before the stimulus is presented. Therefore, it is assumed to be influenced by the prior probabilities of each stimulus. The crucial question for the present work was how to represent the effects of payoffs. Three alternatives are considered next.

Bound-Change Hypothesis

Edwards (1965; see also Rapoport & Burkheimer, 1971) derived the optimal stopping rule for the sequential sampling model—that is, the stopping rule that maximizes expected payoff. According to the optimal rule, the payoffs bias the threshold bounds. On one hand, when larger losses are produced by incorrectly choosing Option A, more evidence is required to stop and choose Option A relative to Option B—that is, the criterion for choosing Option A is larger than the criterion for choosing Option B. On the other hand, when larger losses are produced by incorrectly choosing Option B, more evidence is required to stop and choose Option B relative to Option A—that is, the criterion for choosing B is larger than the criterion for choosing A. Changing the threshold bounds in response to payoff biases is therefore an optimal decision rule for sequential sampling models. Thus, the first hypothesis assumes that the payoffs affect the distance between the starting position and the boundary. This assumption was also incorporated into the model by Link and Heath (1975) and Ratcliff (1978).

Drift-Rate-Change Hypothesis

Ashby (1983) and Ratcliff (1981) proposed an alternative account for the effects of bias, in which the bias affects the way the increment in evidence is interpreted. To express this idea, define $u(t)$ as the increment that occurs with each sample under a neutral payoff condition that elicits no bias. Then, the increment under a biased payoff condition can be expressed as $v(t) = u(t) + \text{bias}$, where bias is positive if the payoffs favor Option A and negative if the payoffs favor Option B. According to this hypothesis, the payoffs affect the accumulation rate of the sequential sampling process.

Two-Stage-Processing Hypothesis

Diederich (1997) proposed a third possibility in which the decision maker views the task as a multiattribute decision problem. The two primary attributes to consider are the payoffs and the stimulus evidence. According to this hypothesis, the decision maker begins the sequential sampling process by evaluating one of the attributes (e.g., the payoffs) and then later switches to processing the second attribute (e.g., the evidence). There are two accumulation

processes: One operates for a period of time during the processing of the first attribute (payoffs), and the second operates later when processing the second attribute (stimulus evidence).

Comparison of Approaches

If no decision is ever reached during the first stage of these two-stage processes, then the two-stage-processing hypothesis and the bound-change hypothesis are very similar. Assuming no decision in the first stage, the second stage is mathematically equivalent to a one-stage hypothesis with the distance between the start position and threshold bounds determined by the previous processing of the payoffs. However, if a decision is frequently reached in the first stage of the two-stage process, then the hypotheses make quantitatively different predictions.

The drift-rate-change hypothesis differs qualitatively from the other two hypotheses in several ways. First, it predicts that the effect of payoffs persists over time, whereas the other two hypotheses predict an early transient effect of payoffs that disappears over time. Second, if the bias is stronger than the mean of the evidence under neutral conditions, then the drift-rate-change hypothesis predicts that errors will actually grow over processing time under a biased payoff condition. The other two hypotheses always predict that errors will decline as the processing time is extended. Figure 1 demonstrates these predictions for the bound-change hypothesis (left panels) and the drift-rate-change hypothesis (right panels). The curves show choice probabilities for choosing one alternative as a function of time limits (arbitrary units) for various bias conditions. The no-bias predictions (solid lines) are the same in all four figures. The upper panels refer to choice probabilities when positive payoffs for choosing that alternative are involved, whereas the lower ones refer to those with negative payoffs.

These hypotheses were tested in an experiment that varied both the bias produced by the payoffs and the processing time manipulated by deadline time limits. A *same-different* line judgment task was employed with discriminability held fixed at a moderate level. Three levels of bias (larger losses for incorrectly choosing the first option, symmetric payoffs, and large losses for incorrectly choosing the second option) were combined with three levels of deadline time limits (200, 350, 500 msec) to form a $2 \times 3 \times 3$ factorial design. We compared the ability of the three hypotheses to account for both the choice probabilities and mean RTs under these 18 conditions.

METHOD

Materials and Experimental Procedure

Stimuli were presented on a 21-in. monitor with a resolution of 640×480 pixels. They consisted of two white lines of 2 pixels in width and of at least 100 pixels in length, presented horizontally to the left and right from the midpoint of the black computer screen, 160 pixels (about 100 mm) apart. The *same* stimuli were 100 pixels (62.5 mm) long, whereas the *different* stimuli differed by 6 or 8 pixels. Prior to the experiment, the actual difference between stimuli was determined, for each participant individually, to satisfy a d'

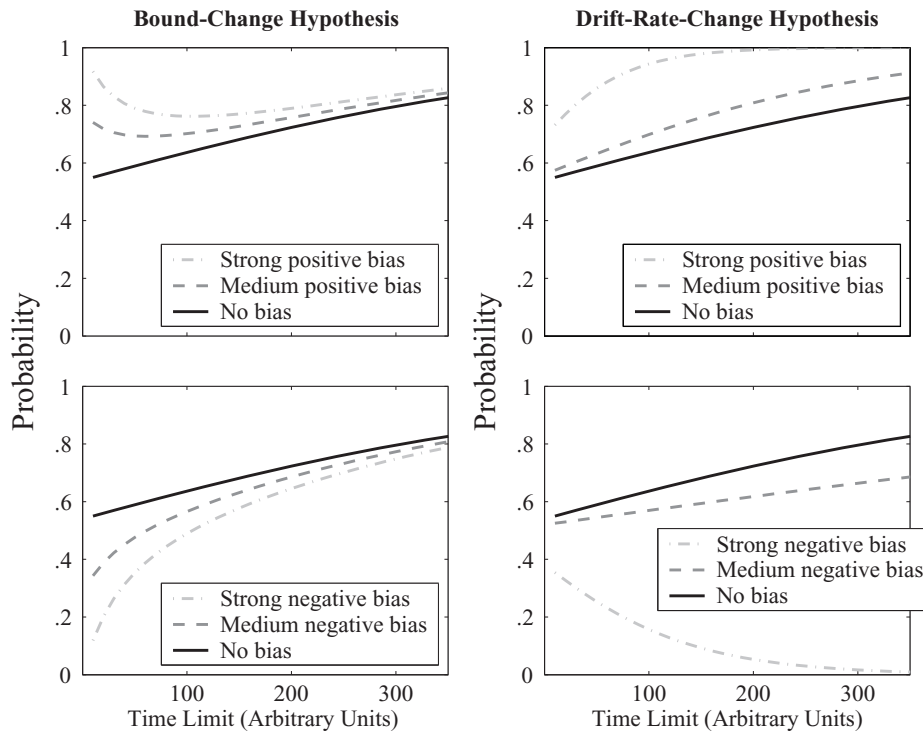


Figure 1. Predictions of the bound-change hypothesis (left panels) and the drift-rate-change hypothesis (right panels). The curves show choice probabilities for choosing one alternative as a function of time limits (arbitrary units) for various bias conditions. The no-bias predictions (solid lines) are the same in all four panels. Choice probabilities with bias induced by positive payoffs are in the top panels, with bias induced by negative payoffs in the bottom panels.

of .5 for a 500-msec stimulus. The display of the payoff matrix was the following. On the left side of the screen appeared two white lines of the same length on top of each other. Below these lines appeared the payoff for a correct response in blue including the plus sign, and below that appeared the payoff for an incorrect response in red including the minus sign. On the right side of the screen appeared the respective arrangement for different lines (i.e., two lines of different length appeared on top of each other including the respective payoffs). The specific payoff values are found in Table 1. The whole display was about 100 × 85 mm. Each matrix was explained to the participant.

The participant sat in a darkened room in front of the screen. To avoid head movements, the head rested on a chinrest. The chinrest was mounted on a table 500 mm from the screen such that the participant’s eyes were directed toward the midpoint of the screen. Eye movements were recorded, but we will not report them here. The response device was a box with two buttons—one to indicate a “same” response (left button), the other to indicate a “different” re-

sponse (right button). The box lay freely movable on the table. Most participants used their left and right thumbs to press the buttons.

Each trial started with the presentation of one of the three payoff matrices for 1,500 msec. In addition, the time limit for making a choice was displayed on the screen. Then, a fixation cross appeared for 500 msec followed by the pair of lines presented for 200, 350, or 500 msec. The participant was required to make a “same” or “different” response within the given time limit. Feedback was given after each trial, with “correct,” “wrong,” or “too slow” displayed on the screen. Violating the deadline limit was punished by subtracting the sum of the two possible losses from the points earned so far. After a pause of 1,500 msec, the next trial started in the same way. At the end of a block of 120 trials, the points earned during it were displayed on the screen. Note that payoff matrices varied, but time limits were fixed within a block of trials and were randomized across blocks of trials.¹ A time course of one trial is presented in Figure 2.

Table 1
Three Different Payoff Conditions

Response	Payoff Condition					
	SAME		DIFFERENT		NEUTRAL	
	Same Stimulus	Different Stimulus	Same Stimulus	Different Stimulus	Same Stimulus	Different Stimulus
“Same”	5	-1	1	-5	5	-5
“Different”	-5	1	-1	5	-5	5

Note—The payoff matrices are assumed to induce a bias for responding “same” (payoff matrix SAME) or “different” (payoff matrix DIFFERENT) or to induce no response bias (payoff matrix NEUTRAL).

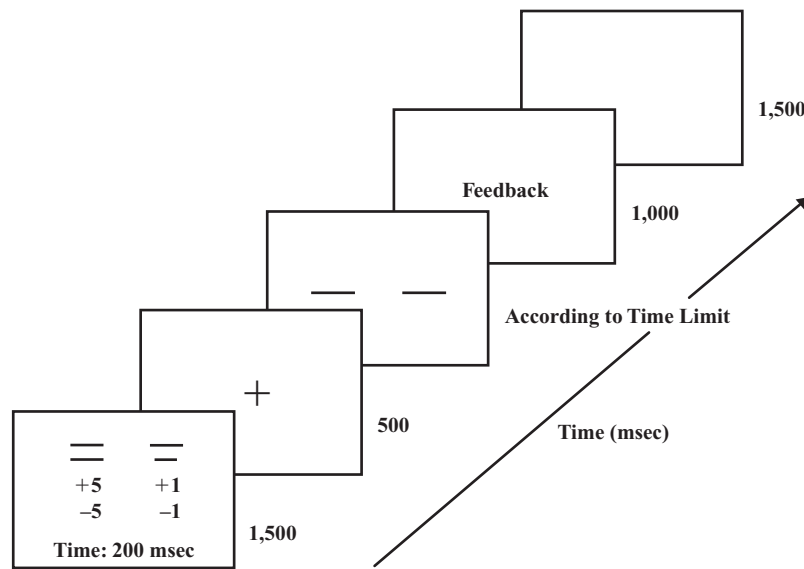


Figure 2. Time course of one trial. The lines were presented for 200, 350, or 500 msec. The feedback was “correct,” “wrong,” or “too slow.” The payoff matrix was randomized within one block of trials; the time limit was fixed within one block of trials. However, the participant was reminded of the time limit before each trial.

Two blocks of 120 trials each were presented in one experimental session (about 1 h). Each of the 18 stimulus combinations (two line length, three payoff matrices, three time limits) was presented 60 times, for a total of 1,080 experimental trials per participant.

Participants

Seven women and 7 men participated in this experiment. They had normal or corrected-to-normal vision. The participants received €5 per experimental session, plus the amount of money they won during each session. Including training sessions, each participant was involved in the experiment for about 6–8 h.

RESULTS

For all participants, choice probabilities differed depending on the experimental condition. Interestingly, however, the patterns of individual choice probabilities seem to fall within two distinct categories: (1) choice probabilities varied as a function of the imposed time limit and the payoff matrix (Group 1), or (2) choice probabilities varied as a function of the imposed time limit but not of the payoff matrix (Group 2). A cluster analysis confirmed two groups. Thus, we combined (averaged) the data from the participants within the two groups, and the following data analysis and hypothesis fits refer to these two groups. Group 1 consisted of 8 participants (4 women and 4 men). Group 2 consisted of 6 participants (3 women and 3 men). Figure 3 shows, for both groups, the relative frequencies of incorrect responses for *different* and *same* stimuli as a function of time limits and payoffs.

As mentioned earlier, the payoff matrices were assumed to induce a bias for saying “same” (SAME), and “different” (DIFFERENT) and to induce no response bias (NEUTRAL). This could be observed for Group 1. For payoff matrix SAME, the participants responded predominantly

with “same,” and, for payoff matrix DIFFERENT, they responded predominantly with “different,” regardless of the presented stimulus. For payoff matrix NEUTRAL, no such effect could be observed. Across time limits, the incorrect responses decreased with increasing time limit for payoff matrix SAME and increased with increasing time limit for payoff matrix DIFFERENT for *different* stimuli. The opposite was true for *same* stimuli: The incorrect responses increased with increasing time limit for payoff matrix SAME and decreased with increasing time limit for payoff matrix DIFFERENT. The results for Group 2 were different since choice probabilities varied as a function of time limit but not of payoff matrix in a systematic way. For both *different* and *same* stimuli, the incorrect responses decreased with increasing time limit for both payoff matrix SAME and payoff matrix DIFFERENT. Statistical tests for probability differences are presented below.

First, the null hypothesis of equal probabilities for all time limits was tested, separately for each payoff condition and stimulus type (*same/different*).

The ν -test statistic

$$\nu = n \cdot \left(\sum_{j=1}^m \sum_{k=1}^r \frac{h_{jk}^2}{h_j \cdot h_k} - 1 \right)$$

for this test is χ^2 distributed with $(m - 1)(r - 1) df$, where n is the number of observations, $m = 3$ is the number of different time limits, r is the number of choices within each time limit (i.e., $r = 2$ for binary choices), and h denotes the respective frequencies within the given time limits and choice categories. The χ^2 s are shown in Table 2. Adjusting the level of significance, $\alpha = .05$, accordingly—that is, with six tests, $\chi^2_{2,983} = 8.15$ —the probabilities differed

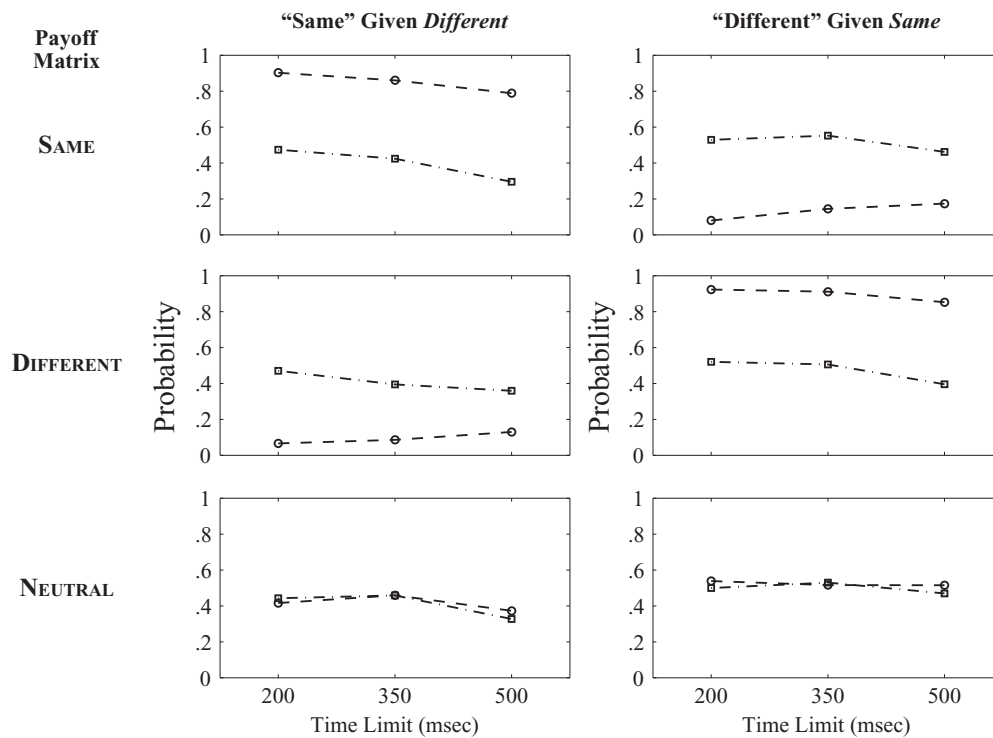


Figure 3. Relative frequencies of incorrect responses as a function of time limits conditioned on the stimulus presented. Each row represents a different payoff. The payoff matrices were assumed to induce a bias for responding “same” (SAME) or “different” (DIFFERENT) or to induce no response bias (NEUTRAL). For Group 1 (dashed line), the responses were biased in the assumed direction: For payoff matrix SAME, the participants responded predominantly with “same,” and for payoff matrix DIFFERENT, the participants responded predominantly with “different,” regardless of the presented stimulus. For payoff matrix NEUTRAL, no such effect could be observed. Moreover, for *different* stimuli, the incorrect responses decreased with increasing time limit for payoff matrix SAME and increased with increasing time limit for payoff matrix DIFFERENT. The opposite was true for *same* stimuli: The incorrect responses increased with increasing time limit for payoff matrix SAME and decreased with increasing time limit for payoff matrix DIFFERENT. For Group 2 (dash-dotted line), choice probabilities varied as a function of time limit but not of payoff matrix in a systematic way. For both *different* and *same* stimuli, the incorrect responses decreased with increasing time limit for payoff matrix SAME and payoff matrix DIFFERENT.

statistically for Group 1 when payoff matrices SAME or DIFFERENT were presented, but they were not different for the NEUTRAL payoff matrix. For Group 2, probabilities were different for response/stimulus pairs “same”/ *different* and “different”/ *same* but also when presenting the NEUTRAL payoff matrix in the *different* stimulus condition.

Table 2
v-Test Statistic Values Across Time Limits

Stimulus	Payoff Matrix	Group 1	Group 2
<i>Different</i>	SAME	23.05	21.72
	DIFFERENT	10.96	8.11
	NEUTRAL	6.79	13.43
<i>Same</i>	SAME	17.52	5.58
	DIFFERENT	13.91	12.28
	NEUTRAL	0.55	2.39

Note—The level of significance, $\alpha = .05$, was adjusted for six tests. If $v > \chi^2_{2,983} = 8.15$, the hypothesis that the $m = 3$ random variables have the same distribution is rejected.

Second, the null hypothesis of equal probabilities for all payoff conditions and stimulus types (*same-different*) was tested, separately for each time limit condition. Thus, $m = 6$, and $r = 2$, for binary choices. The χ^2 s are shown in Table 3. Adjusting the level of significance, $\alpha = .05$, accordingly—that is, with three tests, $\chi^2_{5,992} = 15.48$ —probabilities were statistically different except for the 200-msec condition for Group 2.

In perceptual matching tasks, correct “same” responses are sometimes faster than correct “different” responses, which has been called the *fast-same* effect. Furthermore, incorrect “different” responses are sometimes more frequent than incorrect “same” responses, called the *false-different* effect (Bamber, 1969; Krueger, 1978; for a discussion, see Van Zandt, Colonius, & Proctor, 2000). Both effects could not be observed in the present data. Rather, the probability patterns of both groups could be called *false-same* effect; however, no systematic pattern for RTs could be observed across all experimental conditions. In the following, we will demonstrate whether and how the approaches can account for this complex pattern of data.

Table 3
v-Test Statistic Values Within Time Limits

Time Limit (msec)	Group 1	Group 2
200	1,119	6.15
350	1,094	24.18
500	852	34.28

Note—The level of significance, $\alpha = .05$, was adjusted for three tests. If $v > \chi^2_{5,.992} = 15.48$, the hypothesis that the $m = 6$ random variables have the same distribution is rejected.

Model and Hypotheses

As mentioned above, the three hypotheses investigated here belong to the class of sequential sampling models holding that information (in sensory or cognitive systems) is accumulated sequentially over time until a preset criterion is reached and a response is initiated. The rate of accumulation is determined by how well the lines can be discriminated: The easier the task, the higher the rate, and, therefore, the faster the process approaches the criterion. Assume a criterion for saying “same” (denoted θ_s) and one for saying “different” (denoted θ_d). These decision criteria can be adjusted by the decision maker: Increasing the boundary allows for more evidence to be accumulated and may lead to a more accurate response but, at the same time, may increase the time to initiate a response.

Under deadline conditions, the decision maker may employ an additional strategy. As Ratcliff and Rouder (2000) pointed out, the decision maker may use internal deadlines as well as decision bounds, and, therefore, the accumulation process may be terminated either by a decision bound or by an internal deadline. Both criteria to initiate a response are obviously quite different. The former operates on the state space, whereas the latter is based on the time set. Under short experimenter-set deadlines, the decision maker is inclined to use an internal deadline, and, with increasing time allotted to the task, the decision maker uses the criterion bound. Note that these ideas are related to the deadline model (e.g., Swenson, 1972; Yellot, 1971), in which the RT is determined either by the time needed to complete, for example, a discrimination process, or until a predetermined deadline has arrived, whichever comes first. Ruthruff (1996) tested this model with two different decision tasks (relative brightness judgment and lexical decision task) and rejected it as an adequate explanation for the data. In the following, we incorporate these ideas in testing the hypotheses. However, our primary concern here is to compare different hypotheses about how payoffs may be incorporated into the decision process and not to contrast models for deadline experiments. Nevertheless, we include the χ^2 values as a measure of goodness of fit for both the more traditional view with evidence boundaries only and the one with qualitatively two different boundaries, for evidence and for time (Table 5).

In order to facilitate comparisons between the alternative hypotheses, the underlying sequential process is assumed to be a Wiener process, $X(t)$, with two absorb-

ing boundaries for all three hypothesis (e.g., Diederich & Busemeyer, 2003; Karlin & Taylor, 1975; Tuckwell, 1995). The process begins with an initial state of evidence, denoted $X(0)$, and, without loss of generality, $X(0) > 0$ represents an initial state favoring “same,” $X(0) < 0$ represents an initial state favoring “different,” and $X(0) = 0$ is the neutral state. The increment of information sampled at any moment in time is denoted $v(t)$, where $v(t) > 0$ indicates an increment favoring “same,” and $v(t) < 0$ indicates an increment favoring “different.” The evidence is accumulated from one moment in time, t , to the next moment in time, $t + h$, by summing the current state with the new increment: $X(t + h) = X(t) + v(t + h)$. The mean of the increments, $\mu \cdot h = E[v(t)]$, is called the *drift rate*, and the variance of the increments, $\sigma^2 \cdot h = \text{Var}[v(t)]$, is called the *diffusion rate*. The mean divided by the standard deviation, $E[v(t)]/SD[v(t)]$, is a measure of discriminability closely related to d' in the classic signal detection model. Intuitively, the drift rate reflects the tendency to approach one choice alternative over the other. The diffusion rate reflects the noisiness of the accumulation process.

This process continues until the magnitude of the cumulative evidence exceeds a threshold criterion or it stops before that when an internal deadline is met. First, consider the former. The process stops and a “same” response is initiated as soon as $X(t) > \theta_s = \theta$, or it stops and chooses “different” as soon as $X(t) < \theta_d = -\theta$. The decision criteria (absorbing boundaries) are set by the decision maker prior to the decision task and, among other things, depends on the time available for making a decision. Specifically, the criterion boundary is an increasing function of the time limit. Thus, we assume three evidence-based criteria with $|\theta_1| \leq |\theta_2| \leq |\theta_3|$ for the three time limits: 200, 350, and 500 msec, respectively. Alternatively, the process may terminate when an internal deadline, t_{id} , is met. A “same” response is initiated at time t_{id} if $X(t_{id}) > 0$, and a “different” response will be initiated at time t_{id} if $X(t_{id}) < 0$. Again, the internal deadline is set by the decision maker and increases as the experimenter-set deadline increases. Thus, we assume three internal deadlines with $t_{id1} < t_{id2} < t_{id3}$ for the three time-limit conditions, 200, 350, and 500 msec, respectively.

Figure 4 represents the basic model with $X(0) = 0, \mu = .5$, and $\sigma^2 = 1$. In this example, two responses are initiated by evidence bounds. On one trial, there is enough evidence to initiate a “same” response (the trajectory hits the decision criterion [“same”]); on another trial, there is enough evidence for a “different” response (the trajectory hits the decision criterion [“different”]). On the two remaining trials, the trajectories hit the internal deadline t_{id} before either reaches the evidence bound and, depending on the magnitude of accumulated evidence at that time (>0 or <0), favors a “same” or “different” response.

The difference between the three hypotheses concerns the assumption about how the payoffs influence stimulus processing. In the following, each approach is briefly described. To ease comparison, we assume the same number of parameters for all three hypotheses.

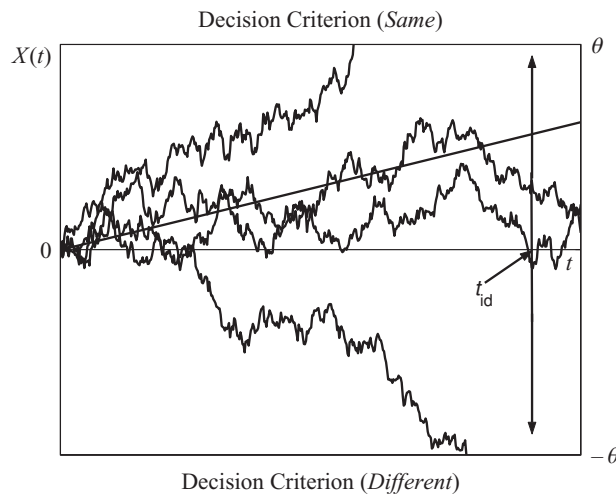


Figure 4. The trajectories symbolize the information accumulation process evolving over time for each single-choice trial. The solid line represents the mean drift. θ and $-\theta$ are the criteria for responding “same” and “different,” respectively.

Bound-change hypothesis. The unbiased Wiener process as described in Figure 4 has symmetric boundaries around the starting position. The Wiener process with bias does not start halfway between the decision boundaries, rather it starts closer to either one (see also Link & Heath, 1975). Here, it is assumed that, depending on the payoffs, a bias for saying “same” or “different” is induced, prior to stimulus processing. Figure 5 depicts the case in which *same* stimuli are presented. Obviously, the drift rate is the same—indicated by the same slope of the drift rate—regardless of the payoffs inducing the response bias. That is, information accumulation processing for the stimuli itself is not influenced by the payoffs; the starting position

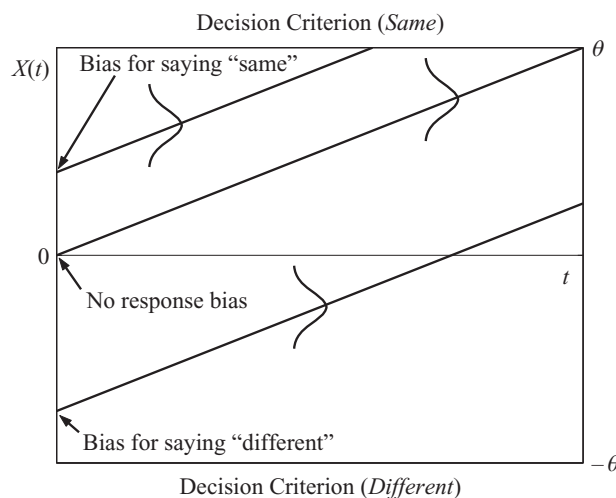


Figure 5. Bound-change hypothesis. The payoff matrix induces an a priori bias. The slope of the mean drift rate is the same regardless of the payoffs. Note that each drift rate comes from a normal distribution indicated by the bell curve.

of the process, however, is. The unbiased Wiener process (here, the one to account for the neutral payoffs) predicts the same RT for both the more frequently chosen alternative and the less frequently chosen alternative, regardless of the drift. However, RT and accuracy data often show a speeded response for the more frequently chosen alternative (see Luce, 1986, and Townsend & Ashby, 1983, for reviews). One way to account for that is to assume that the drift rates themselves are random variables (e.g., Laming, 1968; Ratcliff, 1978, 1981). The drift variability is often interpreted as being due to variability in the stimulus from trial to trial or variability in attention from trial to trial. Although this variability cannot be differentiated from the inherent variability of the stochastic process itself, it improves the fit of the model considerably.

The same reasoning applies when *different* stimuli are presented. The biases are independent of *same* and *different* stimuli (i.e., the bias for saying “same” given *same* stimuli is the same as that for saying “same” given *different* stimuli).

The parameters for this hypothesis include the following: The drift rates for *same* and *different* stimuli are approximately Gaussian distributed with means μ_s and μ_d , respectively, and variance η^2 . For simplicity, variances are assumed to be the same for both distributions: b_1 is the bias for saying “same” induced by payoff matrix SAME [i.e., the process starts at $X(0) = b_1$], and b_2 is the bias for saying “different” induced by payoff matrix DIFFERENT [i.e., the process starts at $X(0) = b_2$]. For the no-bias situation (payoff matrix NEUTRAL), the process starts at $X(0) = 0$.

Drift-rate-change hypothesis. The second hypothesis states that the payoffs actually influence the rate of information accumulation (Ashby, 1983; Ratcliff, 1981). That is, the drift rate depends not only on the stimuli but also on the values of the payoff matrix. As before, to account for different RTs for the more frequently chosen alternatives and the less frequently chosen alternative, a variable drift across trials is assumed. That is, on each trial, the information is accumulated according to a Wiener process with a different drift drawn from a distribution of drifts. Figure 6 presents the situation for *same* stimuli. Obviously, the slope of the drift depends on the payoff matrix in a characteristic way: The higher the payoff for a particular response, the steeper the drift rate for this particular response bias.

The parameters for this approach are the following: μ_s is the mean of the approximately Gaussian distributed drift rate for *same* stimuli, assuming no bias (payoff NEUTRAL); μ_d is the mean of the approximately Gaussian distributed drift rate for *different* stimuli, assuming no bias (payoff NEUTRAL); $\mu_s + b_1$ is the mean of the approximately Gaussian distributed drift rate for *same* stimuli with a bias for saying “same,” induced by payoff matrix SAME; $\mu_d + b_1$ is the mean of the approximately Gaussian distributed drift rate for *different* stimuli with a bias for saying “same,” induced by payoff matrix SAME; $\mu_s + b_2$ is the mean of the approximately Gaussian distributed drift rate for *same* stimuli with a bias for saying “different,” induced by payoff matrix DIFFERENT; $\mu_d + b_2$ is the mean

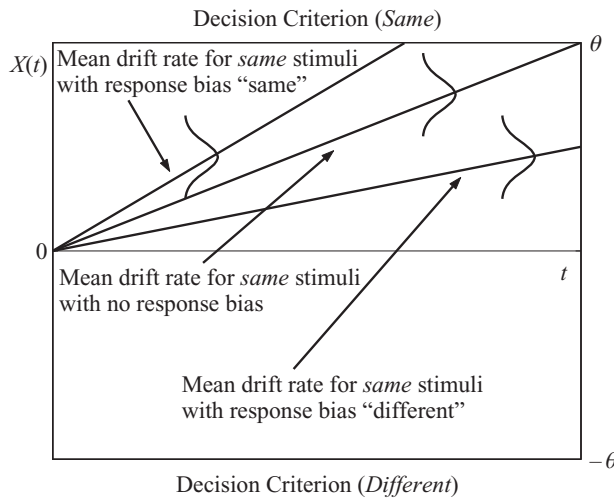


Figure 6. Drift-rate-change hypothesis. The response bias is mapped onto the mean drift rate and is indicated by the slope. Note that each drift rate comes from a normal distribution indicated by the bell curve.

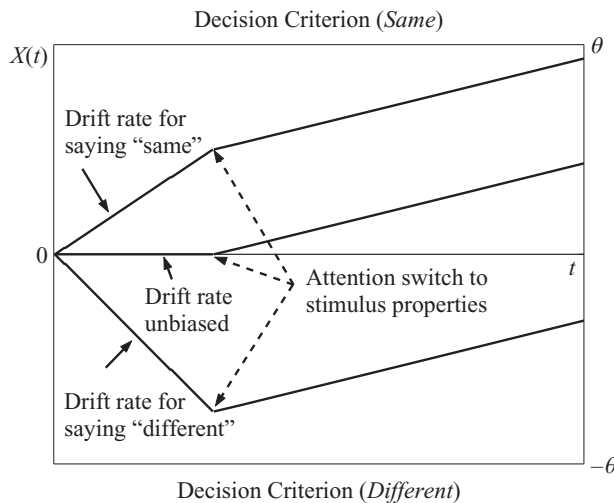


Figure 7. Two-stage-processing hypothesis. The entire process consists of two subprocesses—the first one processing the payoffs, the second processing the stimulus properties. For this hypothesis, it is not necessary to postulate variability in drift.

of the approximately Gaussian distributed drift rate for *different* stimuli with a bias for saying “different,” induced by payoff matrix DIFFERENT. The variance η^2 is assumed to be the same for all six distributions.

Two-stage-processing hypothesis. The third hypothesis assumes a Wiener process with changing drift within a trial (e.g., Diederich, 1995, 1997; see also Heath, 1981, 1992, and Ratcliff, 1980). The entire process consists of two subprocesses and switches from processing the payoff matrix (attribute i) to processing the stimulus properties (attribute j) at a particular rate ω_{ij} . At any particular time point, the process may be operating on the process for the payoffs. During the next moment, either it continues to

operate on the process for the payoffs with probability ω_{11} or attention switches with probability $\omega_{12} = 1 - \omega_{11}$ to operate on the process for the stimuli. That is, the process is a mixture of two subprocesses, and the attributes are processed in a serial manner (for details, see Bussemeyer & Diederich, 2002). Note two important aspects: Different from the previous approaches (1) the process is not time homogeneous (i.e., the drift rate changes within the course of one trial), and (2) it is not necessary to assume variability in drift to account for differences in RTs for correct and incorrect responses in the no-bias condition.

As shown in Figure 7, the process starts with the payoffs and then switches to process the stimulus properties. Unlike the hypothesis with a priori bias, the processing of the payoffs is assumed to occur over time.

The parameters are the following: μ_{b1} is the drift rate for processing payoff matrix SAME; μ_{b2} refers to the drift rate for processing payoff matrix DIFFERENT; the drift rate for processing matrix NEUTRAL (no bias) is set to 0. μ_s is the drift rate for “same” responses and μ_d is the drift rate for “different” responses. ω_{12} refers to the probability of switching from the first to the second attribute.

In addition to the parameters specific for each hypothesis, the following parameters for each approach are $\theta_1, \theta_2, \theta_3$, which refer to the evidence-based decision boundaries, and $t_{id1} < t_{id2} < t_{id3}$, which refer to the internal deadlines for the three time-limit conditions, 200, 350, and 500 msec, respectively.

For all three approaches, mean RT is assumed to be linearly related to the processing time (i.e., decision time, DT) predicted by the hypothesis $RT = s \cdot DT + R$, where s is a time unit scale factor and R is the residual, nondecisional time component (e.g., motor RT; for a discussion, see Luce, 1986).

Table 4 summarizes the hypothesis-specific parameters for all three approaches. Note that the μ s for the bound-change and drift-rate-change hypotheses refer to the mean of the mean drifts, whereas, for the two-stage-processing hypothesis, they refer to the drift rate or one mean drift.

Fits to the Data

For each approach, 13 parameters (i.e., the five hypothesis-specific parameters listed in Table 4, the three parameters

Table 4
Parameters for the Bound-Change, Drift-Rate-Change, and Two-Stage-Processing Hypotheses

Bound Change	Drift Rate Change	Two Stage Processing
b_1	$\mu_s + b_1$	μ_{b1}
b_2	$\mu_s + b_2$	μ_{b2}
μ_s	$\mu_d + b_1$	μ_s
μ_d	$\mu_d + b_2$	μ_d
η^2	η^2	ω_{12}

Note— μ_s refers to the mean drift rate for processing *same* stimuli; μ_d refers to the mean drift rate for processing *different* stimuli; b_1 refers to a bias induced by payoff matrix SAME; b_2 refers to a bias induced by payoff matrix DIFFERENT; η is the variability in the mean drift; ω denotes the probability of switching from one subprocess to the other. Bold μ indicates the mean drift rate itself as a random variable.

for the three time limits, the three parameters for the internal deadline, the scaling factor, and the nondecisional residual time) were estimated from 54 data points (6 probabilities and 12 mean RTs, corresponding to the $[i = \text{“same,” “different”}]$ responses to the $K = 3$ payoff matrices and the $[j = \text{same, different}]$ stimulus conditions for each of the $N = 3$ time limits) by minimizing the Pearson χ^2 statistic

$$\chi^2 = \sum_{n=1}^N \sum_{k=1}^K \left\{ \sum_i \sum_j \frac{[\bar{T}_{ij}(k,n) - \hat{T}_{ij}(k,n)]^2}{\sigma_{\bar{T}_{ij}(k,n)}^2} + \sum_j \frac{[P_{jj}(k,n) - \hat{P}_{jj}(k,n)]^2}{\sigma_{P_{jj}(k,n)}^2} \right\} \quad (1)$$

using the FMINSEARCH routine of MATLAB. Here, $\bar{T}_{ij}(k,n)$ and $\hat{T}_{ij}(k,n)$ are, respectively, the observed values and the fitted values of the mean time to make response i to stimulus j with payoff matrix k and time limit n ; $P_{jj}(k,n)$ and $\hat{P}_{jj}(k,n)$ are the observed and fitted values of the probability of a correct response to stimulus j with payoff matrix k and time limit n ; and $\sigma_{\bar{T}_{ij}(k,n)}$ and $\sigma_{P_{jj}(k,n)}$ are the respective standard errors. The χ^2 values for each hypothesis are found in Table 5.

Assuming that the observed values are normally distributed around the population means we utilize $\chi^2(6 \cdot K \cdot N - n)$, with n the number of estimated parameters, as a statistical test of goodness of fit, with $\chi^2(41)$ approx 65, $p = .10$. According to this criterion, only the two-stage processing hypothesis gave a satisfying fit for Group 2. Nevertheless, we can use the χ^2 as a qualitative measure to compare them.²

The observed choice probabilities and choice response time and the respective predictions of the three hypotheses for both groups are shown in Figures 8–11.

The fits of the hypotheses differ substantially, most obviously for choice probabilities and for the RTs of Group 1. For this group, choice probabilities differed as a function of both time limits and payoffs. The drift-rate-change hypothesis (dashed line) predicts different probabilities depending on payoffs but cannot account for different choice probabilities as a function of time limits. That is, the predicted choice probabilities are flat. For longer time limits, the predicted RTs are severely underestimated in 4 of 12 RT patterns but give a reasonable account for the remaining patterns. The bound-change hypothesis (dash-dotted line) predicts different choice probabilities as a function of time limits and payoffs but underestimates and overestimates them in a nonsystematic fashion. Predicted RTs increase as a function of RT of time limits but tend to be too fast.

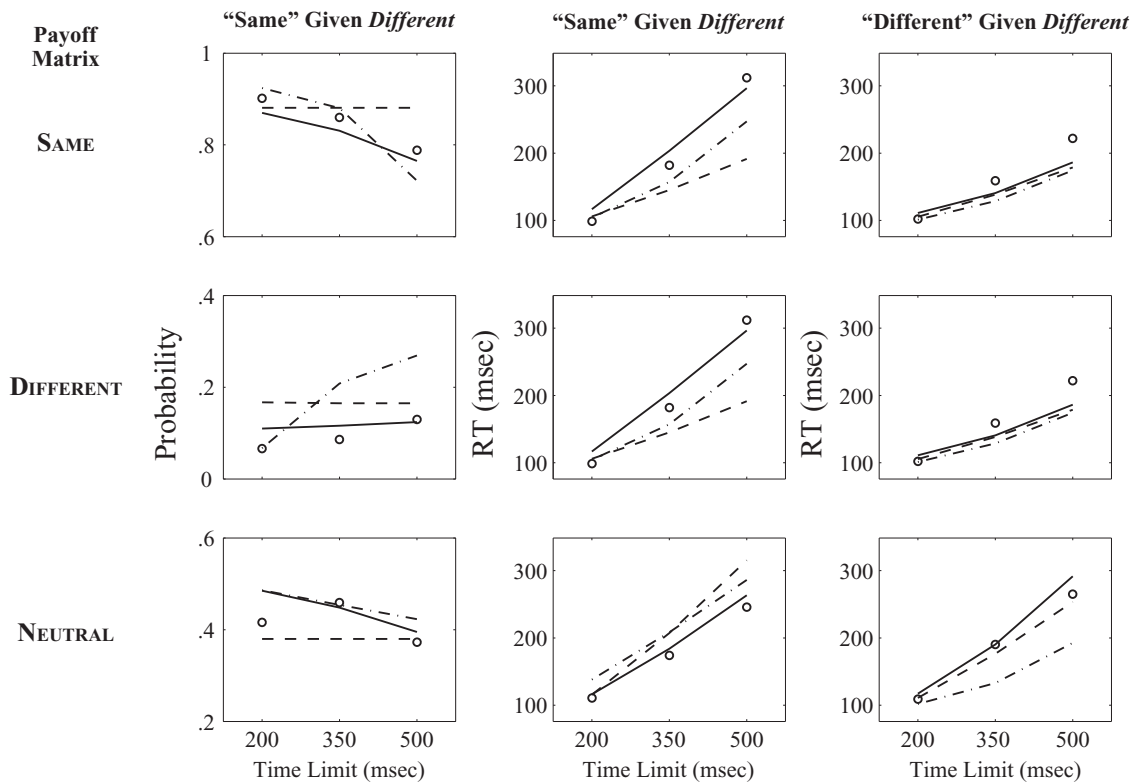


Figure 8. Probabilities for responding “same” and choice response times (RTs) given *different* stimuli for Group 1. Different rows refer to different payoff matrices (i.e., the upper row refers to payoff matrix SAME; the middle row refers to payoff matrix DIFFERENT; the bottom row refers to payoff matrix NEUTRAL). The dashed line represents the predictions of the drift-rate-change hypothesis; the dash-dotted line represents the predictions of the bound-change hypothesis; and the solid line represents the predictions of the two-stage-processing hypothesis.

Table 5
Pearson χ^2 Statistic Values as a Measure of Goodness of Fit for the Bound-Change, Drift-Rate-Change, and Two-Stage-Processing Hypotheses for Groups 1 and 2

Group	Hypothesis	Probability	RT, Incorrect	RT, Correct
1	Bound change	343 (495)	465 (1,004)	1,032 (1,063)
	Drift rate change	322 (322)	465 (817)	868 (3,277)
	Two stage processing	74 (79)	251 (260)	286 (286)
2	Bound change	175 (196)	120 (150)	121 (131)
	Drift rate change	102 (100)	48 (301)	26 (258)
	Two stage processing	22 (24)	25 (83)	17 (42)

Note— $\chi^2(41)$ approx 65, $p = .10$. χ^2 in parentheses refers to the model assuming evidence boundaries only.

The two-stage-processing hypothesis (line) also predicts different choice probabilities as a function of time limits and payoffs. Although far from perfect, it overall accounts best for the observed choice probability patterns. The same holds for the predicted choice RTs. It gives the closest account of all three approaches. This is also reflected in the relative size of the χ^2 values presented in Table 5. Interestingly, all three hypotheses predict almost identical correct RTs for *different* stimuli with payoff matrices SAME and DIFFERENT and for *same* stimuli with payoff matrix SAME.

Group 2 basically ignored the payoffs, but choice probabilities and RT varied as a function of time limits. All three approaches gave a good account for the RTs, both for correct and incorrect responses, but they differed with respect to choice probabilities. Both drift-rate-change and bound-change hypotheses had difficulties to account for choice probability patterns. The choice probabilities are monotonically decreasing with increasing time (i.e., the proportion of incorrect responses decreases as more time is available). Contrary to the trend in the data, the drift-rate-change hypothesis predicts constant probabili-

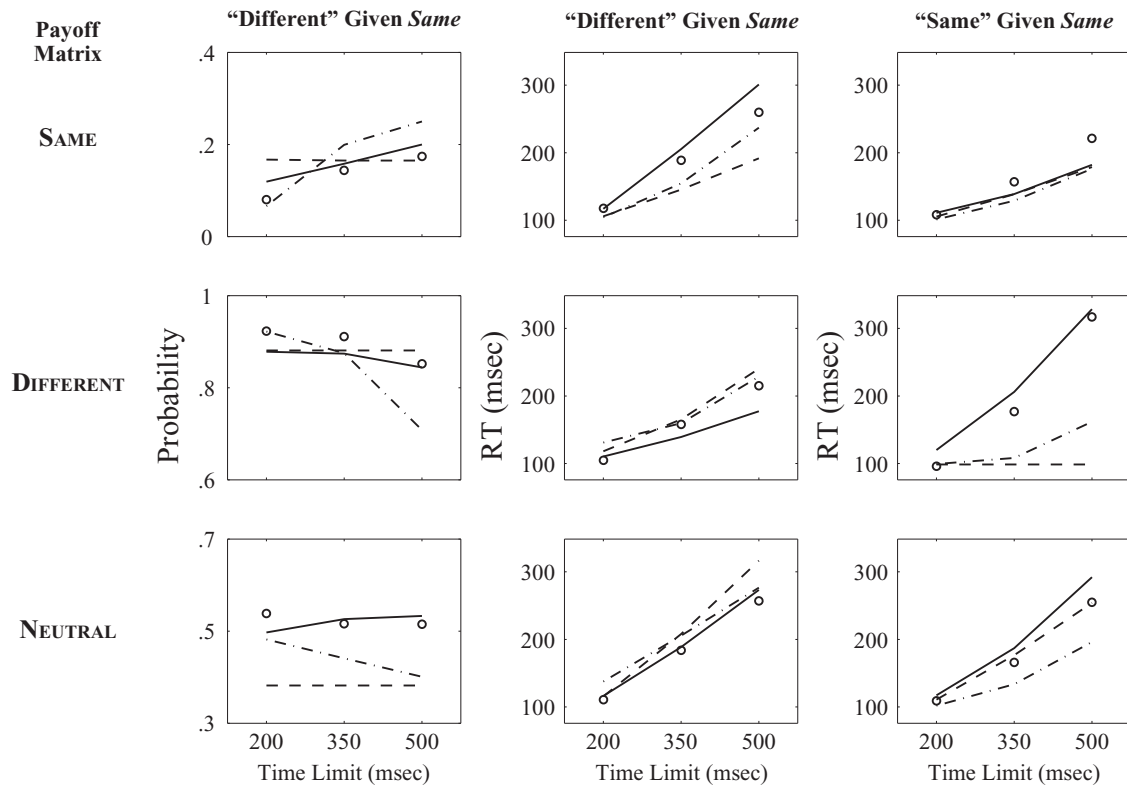


Figure 9. Probabilities for responding “different” and choice response times (RTs) given *same* stimuli for Group 1. Different rows refer to different payoff matrices (i.e., the upper row refers to payoff matrix SAME; the middle row refers to payoff matrix DIFFERENT; the bottom row refers to payoff matrix NEUTRAL). The dashed line represents the predictions of the drift-rate-change hypothesis; the dash-dotted line represents the predictions of the bound-change hypothesis; and the solid line represents the predictions of the two-stage-processing hypothesis.

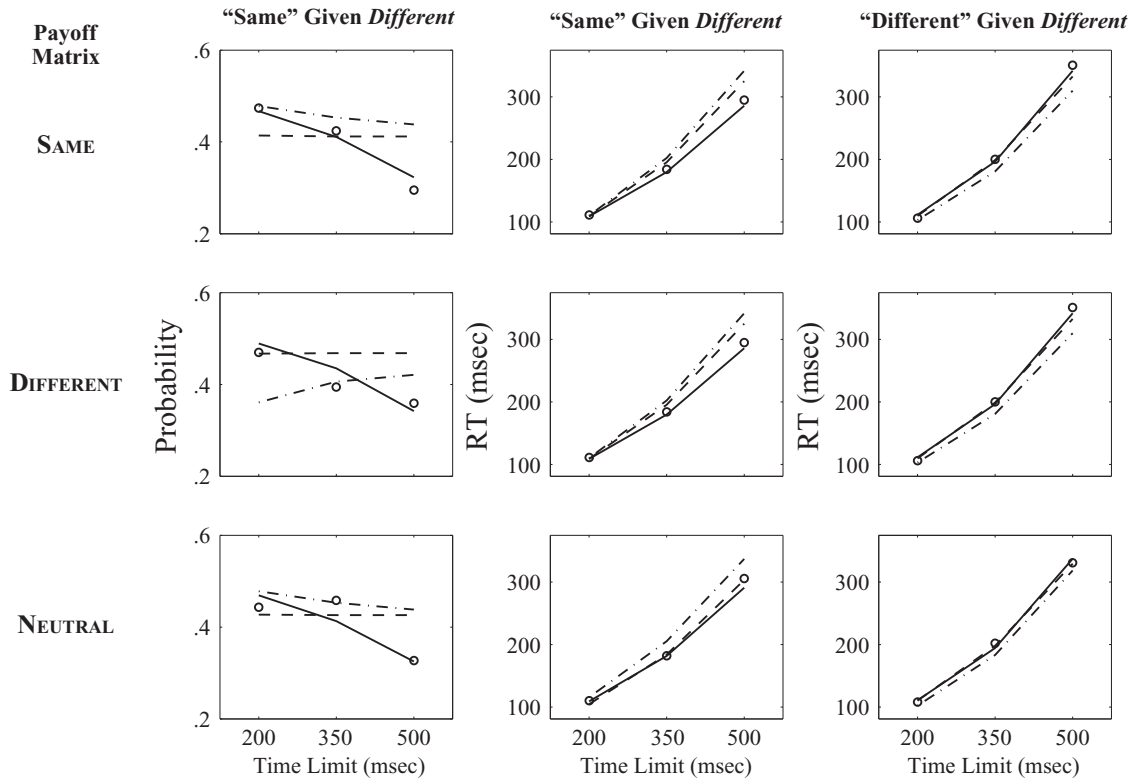


Figure 10. Probabilities for responding “same” and choice response times (RTs) given *different* stimuli for Group 2. Different rows refer to different payoff matrices (i.e., the upper row refers to payoff matrix SAME; the middle row refers to payoff matrix DIFFERENT; the bottom row refers to payoff matrix NEUTRAL). The dashed line represents the predictions of the drift-rate-change hypothesis; the dash-dotted line represents the predictions of the bound-change hypothesis; and the solid line represents the predictions of the two-stage-processing hypothesis.

ties across time limits and the bound-change hypothesis predicts even increasing probabilities for two conditions: for *same* stimuli with payoff matrix SAME and for *different* stimuli with payoff matrix DIFFERENT. The two-stage-processing hypothesis accounts for the patterns but slightly underestimates them for *same* stimuli for the 350-msec time limit. Comparing the χ^2 values of both groups confirms that all three hypotheses were more successful in accounting for the data of Group 2.

Signal Detection Theory Interpretation

The following alternative interpretation was suggested by one of the reviewers. Ideal observers are expected to show larger sensitivity to the payoff matrix when d' decreases. When d' is close to 0, the payoff matrix is expected to drive choice behavior. Furthermore, the members of Group 1 had lower d' than those of Group 2. This is in accordance with the two-stage hypothesis, which allows the decision maker to concentrate solely on the payoff matrix and to ignore the stimuli when initiating the choice.³

According to signal detection theory, an ideal statistical observer maximizes expected long-run rewards. To achieve this goal, the observer chooses a decision criterion, c , that maximizes the number of correct responses and minimizes the cost (or maximizes the gain) associated with the alter-

natives. It can be shown that, at the optimal criterion c^* , likelihood ratio β is determined by the ratio of the probability of an A trial to a B trial, $P(A)/P(B)$, and by the difference between the gain of a correct response and the cost of an error for each of the stimulus types. With V_{aA} and V_{bB} indicating the gains with correct responses, and with V_{aB} and V_{bA} indicating the cost with incorrect responses (lower-case letters indicate the response), the optimal bias β^* is

$$\log \beta^* = \log \frac{V_{bB} - V_{aB}}{V_{aA} - V_{bA}} - \log \frac{P(B)}{P(A)}. \quad (2)$$

Since the ideal observer knows the form of the distribution and the proportion of A and B trials, the relationship among d' , the optimal criterion c^* , and the optimal bias $\log \beta^*$ can be determined. Assuming an equal-variance Gaussian model, the optimal bias β^* is (see, e.g., Wickens, 2002, p. 34)

$$\log \beta^* = d' \left(c^* - \frac{1}{2} d' \right) = d' c^* - \frac{1}{2} d'^2. \quad (3)$$

In the present experiment, $P(A) = P(B)$ for all conditions. Therefore, combining Equation 2 and Equation 3,

$$\log \frac{V_{bB} - V_{aB}}{V_{aA} - V_{bA}} = d' c^* - \frac{1}{2} d'^2. \quad (4)$$

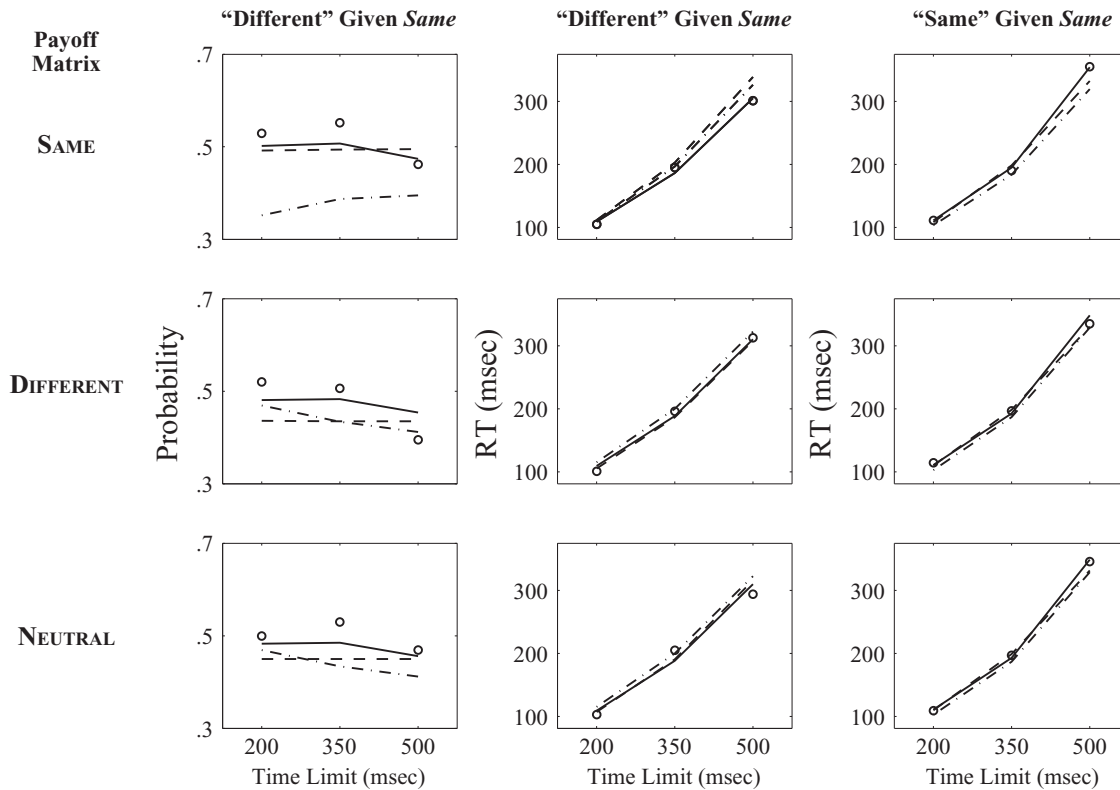


Figure 11. Probabilities for responding “different” and choice response times (RTs) given *same* stimuli for Group 2. Different rows refer to different payoff matrices (i.e., the upper row refers to payoff matrix SAME; the middle row refers to payoff matrix DIFFERENT; the bottom row refers to payoff matrix NEUTRAL). The dashed line represents the predictions of the drift-rate-change hypothesis; the dash–dotted line represents the predictions of the bound-change hypothesis; and the solid line represents the predictions of the two-stage-processing hypothesis.

For small d' ,

$$\frac{\log \frac{V_{bB} - V_{aB}}{V_{aA} - V_{bA}}}{d'} \approx c^* = -Z(\text{false alarm}) \quad (5)$$

(i.e., the choice probabilities are largely determined by the payoffs).

Table 6 presents the d' for both groups, separated according to time limits and payoff matrix. Except for the no-time-limit condition, all d' s are close to 0. The d' s tend to be larger for Group 2 than for Group 1. It seems that for Group 1 choice behavior was mainly driven by the payoff matrix for the time pressure conditions, and members of this group did not try to maximize their outcomes by maximizing the proportion of correct responses. Group 2 basically ignored the payoffs (see above). It seems that members of this group tried to maximize the outcomes by maximizing the number of correct responses, which was difficult to achieve with small d' s, as indicated in Table 6.

Maddox and colleagues (e.g., Bohil & Maddox, 2003a, 2003b; Maddox & Bohil, 2003, 2004; Maddox & Dodd, 2001) proposed and tested a hypothesis that postulates that the decision criterion placement results from a competition between reward and accuracy maximization

(COBRA). It is assumed that the decision criterion used by the observer on any given trial is determined from a weighted combination of the observer’s estimate of the reward-maximizing decision criterion and the accuracy-maximizing criterion. The reward-maximization decision criterion is influenced, among other things, by the discriminability of the stimuli and by payoff matrix entries. COBRA postulates that the observers attempt to maxi-

Table 6
 d' Conditioned on Payoff Matrix and Time Limit for Groups 1 and 2

Time Limit	Payoff Matrix	Group 1	Group 2
200 msec	SAME	0.11	0
	DIFFERENT	0.08	0.03
	NEUTRAL	0.31	0.14
350 msec	SAME	0	0.06
	DIFFERENT	0.02	0.25
	NEUTRAL	0.11	0.03
500 msec	SAME	0.14	0.63
	DIFFERENT	0.08	0.53
	NEUTRAL	0.36	0.53
None	SAME	1.45	1.85
	DIFFERENT	1.42	1.68
	NEUTRAL	1.46	1.88

mize expected reward but that they also place importance on accuracy maximization. The results of Group 1 seem to violate this prediction: Members of this group did not try hard to be accurate for payoff matrices SAME and DIFFERENT for short deadlines.

Results of Group 2 suggest that they tried to maximize the number of correct responses but were not very successful in maximizing their overall rewards because of a low ability to discriminate between stimuli. Members of Group 1 received 1,885 points, on average, across all four time limits (omitting punishments for violating the deadline). Members of Group 2 received 1,052 points, on average. However, considering only the no-time-limit condition, Group 1 received 778 points, and Group 2 received 828 points, providing partial support for COBRA.

DISCUSSION

Sequential sampling models provide an improvement over the static signal detection model by their ability to account for speed-accuracy trade-offs as well as providing measures of discriminability and response bias. Although many empirical studies have tested the response bias assumptions of signal detection theory, very little work has examined the way that payoffs affect response bias in sequential sampling models. The purpose of this article was to present the first comprehensive test of three different hypotheses concerning the way that payoffs affect response bias in sequential sampling models of perceptual decision making. The bound-change hypothesis states that payoffs affect the distance of the starting position to each bound. The drift-rate-change hypothesis states that payoffs affect the mean drift rate or, in other words, the way evidence is encoded. The two-stage-processing hypothesis states that payoffs determine the drift rate during the first stage and that stimulus information influences the second stage of processing. A large experiment was conducted that factorially manipulated stimulus information, payoffs, and processing time to produce a total of 18 conditions. Each of the three hypotheses was evaluated by formulating it as a particular version of a sequential sampling model and by estimating 13 parameters and evaluating the quantitative fit of the hypothesis to the 54 observed data points (18 choice probabilities, 36 conditional mean RTs). The main results of these analyses are as follows. The bound-change hypothesis produced the poorest account of the results, and the drift-change hypothesis produced a slightly better fit; but the two-stage-processing hypothesis provided the most accurate account. This suggests that individuals first process the payoff information for a period of time and later switch to the stimulus information. While processing the payoff information, they may occasionally cross a threshold and choose before the stimulus display even appears. If the latter does not occur, then they start from the state determined by the payoff and process the stimulus until a threshold is reached. These conclusions may be restricted to the method used in the present study to manipulate the payoffs. In order to make the payoff information very salient to the decision

maker, the payoff condition was varied randomly from one trial to the next, which forced the decision maker to pay attention to the new payoff information at the beginning of each trial. However, if the payoff information is held constant for a block of trials, then this may affect the processing of the payoffs and, consequently, the relative fits of the three hypotheses. Future research will be conducted to test this possibility.

REFERENCES

- ASHBY, F. G. (1983). A biased random walk model for two choice reaction times. *Journal of Mathematical Psychology*, *27*, 277-297.
- BAMBER, D. (1969). Reaction times and error rates for "same"-"different" judgments of multidimensional stimuli. *Perception & Psychophysics*, *6*, 169-174.
- BOHIL, C. J., & MADDOX, W. T. (2003a). On the generality of optimal versus objective classifier feedback effects on decision criterion learning in perceptual categorization. *Memory & Cognition*, *31*, 181-198.
- BOHIL, C. J., & MADDOX, W. T. (2003b). A test of the optimal classifier's independence assumption in perceptual categorization. *Perception & Psychophysics*, *65*, 478-493.
- BUSEMEYER, J. R., & DIEDERICH, A. (2002). Survey of decision field theory. *Mathematical Social Sciences*, *43*, 345-370.
- DIEDERICH, A. (1995). Intersensory facilitation of reaction time: Evaluation of counter and diffusion coactivation models. *Journal of Mathematical Psychology*, *39*, 197-215.
- DIEDERICH, A. (1997). Dynamic stochastic models for decision making with time constraints. *Journal of Mathematical Psychology*, *41*, 260-274.
- DIEDERICH, A., & BUSEMEYER, J. R. (2003). Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *Journal of Mathematical Psychology*, *47*, 304-322.
- EDWARDS, W. (1965). Optimal strategies for seeking information: Models for statistics, choice reaction times, and human information processing. *Journal of Mathematical Psychology*, *2*, 312-329.
- GREEN, D. M., SMITH, A. F., & VON GIERKE, S. M. (1983). Choice reaction time with a random foreperiod. *Perception & Psychophysics*, *34*, 195-208.
- GREEN, D. M., & SWETS, J. A. (1966). *Signal detection and psychophysics*. New York: Wiley.
- HEATH, R. A. (1981). A tandem random walk model for psychological discrimination. *British Journal of Mathematical & Statistical Psychology*, *34*, 76-92.
- HEATH, R. A. (1992). A general nonstationary diffusion model for two-choice decision making. *Mathematical Social Sciences*, *23*, 283-309.
- KARLIN, S., & TAYLOR, H. M. (1975). *A first course in stochastic processes* (2nd ed.). New York: Academic Press.
- KORNBROT, D. E. (1988). Random walk models of binary choice: The effect of deadlines in the presence of asymmetric payoffs. *Acta Psychologica*, *69*, 109-127.
- KRUEGER, L. E. (1978). A theory of perceptual matching. *Psychological Review*, *85*, 278-304.
- LAMING, D. (1968). *Information theory of choice reaction times*. New York: Academic Press.
- LINK, S. W., & HEATH, R. A. (1975). A sequential theory of psychological discrimination. *Psychometrika*, *40*, 77-105.
- LUCE, R. D. (1986). *Response times*. New York: Oxford University Press.
- MADDOX, W. T. (2002). Toward a unified theory of decision criterion learning in perceptual categorization. *Journal of the Experimental Analysis of Behavior*, *78*, 567-595.
- MADDOX, W. T., & BOHIL, C. J. (2003). A theoretical framework for understanding the simultaneous base-rate and payoff manipulations on decision criterion learning in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, *29*, 307-320.
- MADDOX, W. T., & BOHIL, C. J. (2004). Probability matching, accuracy maximization, and a test of the optimal classifier's independence as-

sumption in perceptual categorization. *Perception & Psychophysics*, **66**, 104-118.

MADDOX, W. T., & DODD, J. L. (2001). On the relation between base-rate and cost-benefit learning in simulated medical diagnosis. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **27**, 1367-1384.

RAPOPORT, A., & BURKHEIMER, G. J. (1971). Models for deferred decision making. *Journal of Mathematical Psychology*, **8**, 508-538.

RATCLIFF, R. (1978). A theory of memory retrieval. *Psychological Review*, **85**, 59-108.

RATCLIFF, R. (1980). A note on modeling accumulation of information when the rate of accumulation changes over time. *Journal of Mathematical Psychology*, **21**, 178-184.

RATCLIFF, R. (1981). A theory of order relations in perceptual matching. *Psychological Review*, **88**, 552-572.

RATCLIFF, R., & ROUDER, J. N. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception & Performance*, **26**, 127-140.

RUTHRUFF, E. (1996). A test of the deadline model for speed-accuracy tradeoffs. *Perception & Psychophysics*, **58**, 56-64.

SWENSSON, R. G. (1972). The elusive tradeoff: Speed vs accuracy in visual discrimination tasks. *Perception & Psychophysics*, **12**, 16-32.

SWETS, J. A., TANNER, W. P., JR., & BIRDSALL, T. G. (1961). Decision processes in perception. *Psychological Review*, **68**, 301-340.

TOWNSEND, J. T., & ASHBY, F. G. (1983). *Stochastic modeling of elementary psychological processes*. Cambridge: Cambridge University Press.

TUCKWELL, H. C. (1995). *Elementary applications of probability theory: With an introduction to stochastic differential equations* (2nd ed.). London: Chapman & Hall.

VAN ZANDT, T., COLONIUS, H., & PROCTOR, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, **7**, 208-256.

WICKENS, T. D. (2002). *Elementary signal detection theory*. New York: Oxford University Press.

YELLOTT, J. I. (1971). Correction for fast guessing and the speed-accuracy

tradeoff in choice reaction time. *Journal of Mathematical Psychology*, **8**, 159-199.

NOTES

1. Blocks of a no-time-limit condition were also included in the experiment, and we intended to compare it with the time-limit conditions as well. Unfortunately, due to a programming error, the stimulus in this condition was presented for 500 msec, followed by a dark time in which the decision was made. Since this condition was different from the other three, we did not include it in the fits. The choice probabilities and choice RTs are presented in the Appendix.
2. One reviewer felt that the bound-change hypothesis did not get fair treatment because the parameter for the bias induced by the neutral payoff matrix was set to 0. For Group 1, the relative frequency for incorrect responses when *same* stimuli were presented was .52 across all time limits, and the one when *different* stimuli were presented was .42. For Group 2, the relative frequencies were .5 and .41, respectively. For both stimulus conditions (*same* and *different*), we assumed only one parameter, and, for parsimony, setting the parameter to 0 seemed justifiable. The same reason applied to the two remaining hypotheses. Nevertheless, we fitted the bound-change hypothesis again with an additional parameter for the no-bias situation. The improvement of the fit was marginal. Allowing variability in the starting position may improve the fit but adds still another parameter.
3. Further evidence for ignoring the stimulus under short time limits seems to be given by the recorded eye movements in this study. However, we cannot completely rule out that the participants ignored the stimulus on some trials and ignored the payoffs on others. In this case, no interaction would occur between payoffs and stimulus information; therefore, payoffs would not affect stimulus processing. None of the above models here were designed to describe this situation. Such a mixture of two processes—one for payoffs, operating on some trials, and one for stimulus information, operating on the remaining trials—is a possibility, but it was not pursued in the present study.

APPENDIX
Observed Choice Probabilities, Mean Response Times (RTs), and Standard Errors (SEs) for Incorrect and Correct Responses for the No-Time-Limit Condition

Stimulus	Payoff Matrix	Prob _{incorrect}	Group 1				Group 2				
			RT, Incorrect		RT, Correct		RT, Incorrect		RT, Correct		
			M	SE	M	SE	Prob _{incorrect}	M	SE	M	SE
<i>Different</i>	SAME	.338	679	21.1	726	13.0	.147	747	50.1	652	11.2
	DIFFERENT	.148	823	39.9	665	10.5	.189	683	27.0	661	11.3
	NEUTRAL	.258	754	27.7	696	11.7	.161	664	36.3	683	10.6
<i>Same</i>	SAME	.150	796	38.8	689	11.5	.211	670	30.9	707	11.8
	DIFFERENT	.354	710	25.2	756	14.5	.211	651	36.8	716	11.4
	NEUTRAL	.208	786	40.1	733	13.7	.186	645	31.6	707	11.2

Note—Due to an error in the program for stimulus presentation, these data were not included in the model fits.

(Manuscript received May 17, 2004;
revision accepted for publication March 11, 2005.)