

## Judgment and Decision Making: Behavioral Approaches

Edmund Fantino

University of California, San Diego

The area of judgment and decision making has given rise to the study of many interesting phenomena, including reasoning fallacies, which are also of interest to behavior analysts. Indeed, techniques and principles of behavior analysis may be applied to study these fallacies. This article reviews research from a behavioral perspective that suggests that humans are not the information-seekers we sometimes suppose ourselves to be. Nor do we utilize information effectively when it is presented. This is shown from the results of research utilizing matching to sample and other behavioral tools (monetary reward, feedback, instructional control) to study phenomena such as the conjunction fallacy, base-rate neglect, and probability matching. Research from a behavioral perspective can complement research from other perspectives in furthering our understanding of judgment and decision making.

*Key words:* decision making, matching to sample, base-rate neglect, conjunction fallacy, probability learning, observing

---

Our research has focused on choice behavior since my first experiment (on self-control in the pigeon) as a graduate student in 1961. Interest in choice has gone hand in hand with interest in conditioned reinforcement, which has long occupied a central role in an appreciation of human behavior. The reasons why some conditioned reinforcers are chosen over others led us to develop delay-reduction theory and to study behavioral analogues to foraging behavior. Of present relevance, however, is the debate about the conditions under which an informative stimulus serves as a conditioned reinforcer. This debate took place in the context of what is known as *observing behavior*. We shall see that research on observing suggests that, under some conditions, humans are not the information-seekers we might suppose ourselves to be. This

conclusion caused us to have an interest in how well we utilize information once we have acquired it, a pursuit that led naturally to the fascinating world of logical fallacies developed so thoroughly by specialists in judgment and decision making. Given our background, we would hope to understand these fallacies, and human choice more generally, within a consistent behavioral framework. The present article summarizes some of our findings, and it identifies how a behavioral approach has something to offer, in conjunction with cognitive, quantitative, and social psychology approaches, to a fuller understanding of decision making.

### *Observing*

Research in decision making has emphasized how effectively and appropriately information is utilized. A perhaps more basic issue concerns whether people seek information in the first place. This issue has been dealt with in the observing-response literature. There are two basic questions in the area of observing: (a) Do subjects seek information when it cannot influence reward outcomes and has no other present utility? (b) If so, why?

The classic observing-response procedure was developed by Wyckoff

---

Invited address at the annual meeting of the Society for Judgment and Decision Making, Philadelphia, November 23, 1997. I thank David Case and Greg Zarow for helpful comments on the original version. Preparation of this article was supported by NSF Grant IBN 9870900 and NIMH Grant MH57127 to the University of California, San Diego.

Reprints may be obtained from Edmund Fantino, Department of Psychology-0109, University of California, San Diego, LaJolla, California 92093-0109 (E-mail: efantino@ucsd.edu).

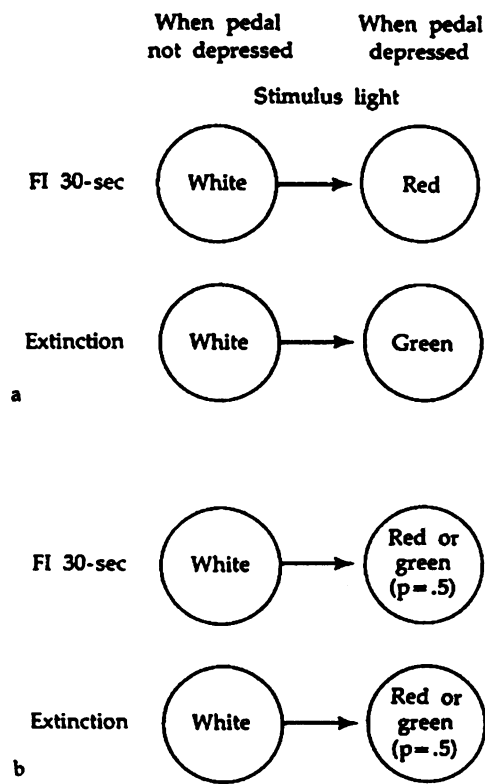


Figure 1. Typical observing-response procedure as developed by Wyckoff (1952). (a) In the correlated condition, the subject's observing response changes a mixed schedule into an analogous multiple schedule, as shown. The red and green stimulus lights are correlated with an FI 30-s schedule and extinction, respectively. (b) In the uncorrelated condition, observing responses also produce a red or green light. In this case, however, the lights are uncorrelated with the schedules in effect.

(1952) and is shown in Figure 1. In the "correlated" or experimental condition, the subject's observing response (for Wyckoff, the depression of a pedal by a pigeon) changes a mixed schedule into an analogous multiple schedule, as shown. That is, if the pedal is not depressed (no observing) a white light is associated with the stimulus key whether the food schedule in effect is fixed-interval (FI) 30 s or extinction (no food available). When the pedal is depressed, however (observing), red and green stimulus lights are correlated with the FI and extinction schedules, respectively, as shown. In other words,

an effective observing response (in this experiment depression of the pedal) converts the white light to either red or green, depending upon the schedule in effect. When the FI schedule is in effect, key pecks to the white or red light produce food upon the first response following completion of the 30-s interval. When extinction is in effect, key pecks to the white or green light are ineffective. It is critical to emphasize that pedal pressing has no effect on these outcomes. Observing has an informative (or in behavioral terms, discriminative) function only. Next note the arrangement in the bottom half of Figure 1. In this control (or "uncorrelated") condition, pedal depression again produces a change in the key-light from white to either red or green, but in this condition there is no informative or discriminative function: The red and green keylights are uncorrelated with the schedule in effect on the food key. In such an uncorrelated condition, pigeons will exhibit relatively little pedal depression (presumably maintained by sensory reinforcement, i.e., by stimulus change). The fundamental observation is whether significantly more pedal depression is maintained in the correlated condition. Wyckoff's results demonstrated a large positive effect. Subsequent studies with other species have confirmed the robustness of observing: People, monkeys, rats, pigeons, and fish all observe, despite the fact that observing produces no change in the scheduled rate of primary reinforcement. This then leads to the second basic question about observing: Why does it occur?

Since the phenomenon of observing was demonstrated, there have been two principal theories to account for it. One, the information or uncertainty-reduction hypothesis, stipulates that informative stimuli are reinforcers. This theory makes good intuitive sense: Organisms that attend to currently neutral stimuli may be able to profit later when the stimuli become relevant for reinforcement. The second theory, the conditioned reinforcement hypothesis,

stipulates that observing is maintained by production of a stimulus correlated with positive reinforcement. This theory also makes sense intuitively: Stimuli paired with a higher likelihood of reward (or a reduction in average time to reward) are likely to be conditioned reinforcers; stimuli paired with a lower likelihood of reward (or an increase in average time to reward) are unlikely to be conditioned reinforcers. Bloomfield (1972) captured the critical difference between the theories when he noted that, according to the uncertainty-reduction view both "good news" and "bad news" should maintain observing (because both are informative), whereas according to the conditioned reinforcement view only good news should maintain observing.

A series of elegant studies from Jim Dinsmoor's laboratory at Indiana University (e.g., Dinsmoor, 1983; Mulvaney, Dinsmoor, Jwaideh, & Hughes, 1974) seemed to decide the issue in favor of the conditioned reinforcement hypothesis, at least when pigeons are the subjects (a review of this early literature, including the theories, may be found in Fantino & Logan, 1979, pp. 197–207). Although some more recent studies involving primates suggest that under some conditions bad news (i.e., a stimulus correlated with extinction) will maintain observing, the preponderance of evidence suggests that this will occur only when the bad news has some usefulness (e.g., in saving effort or in signaling time to pursue other activities).

For example, David Case and I conducted a series of studies, with students from the University of California–San Diego (UCSD) as subjects, that differed in two important ways from those of Wyckoff and of Dinsmoor and his colleagues (in addition to the species difference, of course). First, reinforcement (points, sometimes backed by money) was delivered independently of responding. That is, we used response-independent schedules of reinforcement (e.g., a variable-time schedule in which reinforcers occur indepen-

dently of any responses after random intervals). Such a schedule is analogous to the more common variable-interval schedule, except in the latter a response is required after the interval elapses. Jenkins and Boakes (1973) had previously used response-independent schedules in a study of observing with pigeons. By using response-independent schedules, we eliminated the possibility (occurring with more common response-dependent schedules) that production of the S- (bad news) would permit the subject to save effort by not responding on the extinction schedule. Second, we gave our subjects a choice of two observing responses, because in our experience, choice is a more sensitive measure of response strength than is absolute rate of responding. Also, Mulvaney et al. (1974) had used a choice measure in a study of observing with pigeons. In some conditions, subjects could respond on a key that occasionally produced only good news when appropriate (i.e., when the reward schedule was in effect) or on a key that produced only bad news when appropriate (i.e., when extinction was in effect). Responses on either observing schedule were equally informative, but only one should be reinforcing according to the conditioned reinforcement view. In the most interesting condition, responses on one key could produce only bad news, whereas responses on the other key produced only stimuli uncorrelated with the schedules in effect (no news). According to the information view, bad news should be preferred to unreliable news; but according to the conditioned reinforcement view, relative to bad news, no news is good news. As shown in Figure 2, our results supported the conditioned reinforcement position.

The top set of panels shows the mean and individual-subject data for 6 UCSD students in one of several replications of this experiment (these results are from Experiment 2 of Fantino & Case, 1983). The bars on the left indicate preference when no news was pitted against no news (i.e., a control

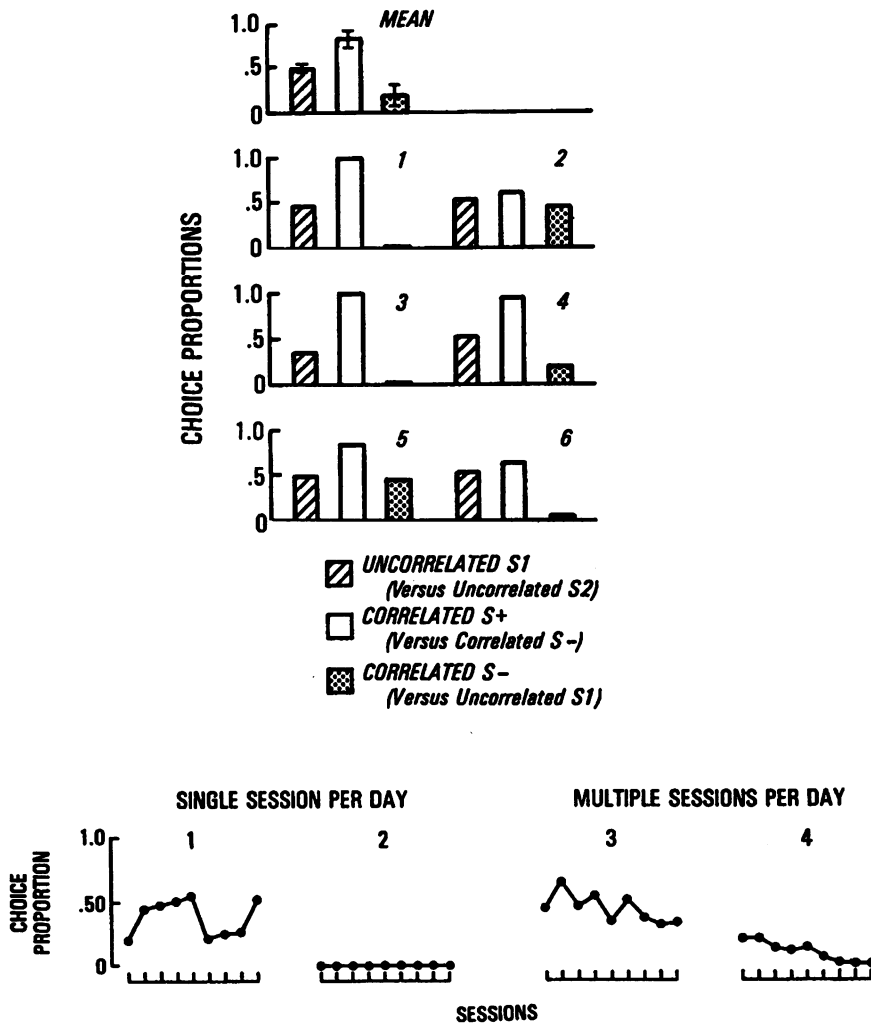


Figure 2. Top: Mean and individual choice proportions for each of three conditions from Experiment 2 of Fantino and Case (1983). The numbers 1 through 6 refer to each of 6 subjects. Choice proportions are based on observing responses made prior to the production of observing stimuli on a given trial. The left bar in each panel gives the choice proportion maintained by one stimulus uncorrelated with points compared to the other stimulus, also uncorrelated with points (baseline). The center bar gives the choice proportion maintained by the S+ compared to the S-. The right bar gives the choice proportion maintained by the S- compared to the stimulus uncorrelated with points. Error bar length is one standard deviation. Bottom: Individual choice proportions for each of 4 subjects (numbered 1 through 4) from Experiment 4 of Fantino and Case (1983). Preference was calculated with respect to the stimulus correlated with extinction (relative to the stimulus uncorrelated with reinforcement). Subjects 1 and 2 were studied in single sessions per day, and Subjects 3 and 4 were studied in three sessions per day.

condition). As expected, preference was near .50 (or indifference). The second set of open bars indicates preference for good news over bad news, and the third set of bars indicates preference expressed in terms of bad news versus no news. There is no suggestion

that bad news was preferred to unreliable news. This was true even when the instructions made it explicitly clear that responding had no effect on the earning of points and when only subjects who correctly answered a set of questions designed to assess their ac-

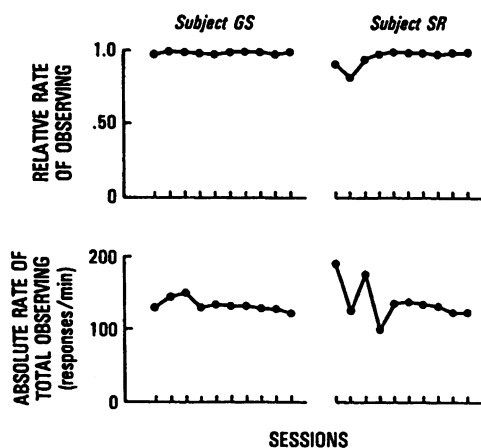


Figure 3. Choice proportions (top) and total absolute rate of observing (bottom) for Subjects GS and SR in each session. Choice proportions were calculated with respect to the stimulus uncorrelated with reinforcement (i.e., no news). Subjects chose between producing this stimulus versus a stimulus correlated with extinction (bad news). (From Fantino, Case, & Altus, 1983)

curate understanding of the instructions were allowed to participate further. Nor did preference for bad news develop with greater exposure to the procedure. In a later experiment in Fantino and Case (1983), 4 additional subjects were studied for nine sessions. As shown in the bottom portion of Figure 2, there was no evidence of an increase in preference for bad news over no news across sessions for these subjects (again the figure plots preference for individual subjects in terms of bad news; thus, preference for bad news would involve choice proportions above .50 and an increase in such preference over sessions would be reflected by an upward slope of the data, which is not evident).

Fantino, Case, and Altus (1983) replicated the same general results with children from three different age groups (4–5, 9–10, and 13–14 years old). For 2 children the key comparison of preference for no news versus bad news was extended over 10 sessions. As shown in Figure 3, which plots preference for no news over bad news (top) and absolute rate of observing (bottom), although these 2 children

(9 to 10 year olds) observed at a high rate (upwards of 100 responses per minute) there was no suggestion of preference for bad news over no news, nor did the opposite preference appear to diminish over sessions.

If subjects were permitted to save effort by producing the stimulus correlated with extinction, then preference for bad news over no news might emerge. Case, Fantino, and Wixted (1985) did not obtain such a preference but, for each of their 6 subjects (again UCSD students) the degree of selection of bad news over no news increased in a response-dependent procedure when compared to a response-independent procedure (i.e., there was a shift from preference for no news when no effort could be saved to indifference when it could be). But one important point should be clear: Whereas the utility of bad news may sometimes modulate its reinforcing potency, good news is preferred to no news even when it has no utility.

It can be argued, of course, that our tasks are artificial and that our results might be quite different in a more natural setting. David Case, Bertram Ploog, and I (1990) reasoned that for UCSD students nothing could be more natural than playing computer games. Thus, David modified the popular Star Trek game so that it could serve as the context for asking the same questions about preference for no news versus bad news under conditions in which the bad news could or could not be readily utilized. For example, different commands (Battle Computer, Tricorder) might provide different types of information (bad news, no news) about the availability of reinforcement (Klingons in the area to be destroyed) and, in some conditions, the information that reinforcement is unavailable (no Klingons in area) permitted refueling of the Enterprise space ship (a benefit). For example, if the schedule in effect had not yet arranged for an observing response to be effective, the message, "Command not available—try again later," was displayed. When observing

responses were effective, they changed the indicator for "battle status" in the upper right part of the screen appropriately. Thus, Mr. Spock might deliver the message: "Captain, my instruments suggest that Klingons are approaching. Beware!"

We found answers comparable to those presented earlier: No news was preferred to bad news when the latter could not be utilized; but bad news was actually preferred when it permitted refueling (of course, in the larger scheme of the game this news was not bad). We are currently assessing the generality of these findings in the area of medical diagnoses: Under what conditions will subjects seek unwelcome information about their medical condition?

### *Fallacies*

If humans appear to be more interested in how favorable stimulation is than in its informative value, how does that argue for their use of information when confronted with conflicting sources of information (problems in multiple stimulus control) as in some of the classic fallacy problems? Let us consider three such areas in which decision making is typically nonoptimal: (a) our response to compound statements or stimuli as in the conjunction fallacy; (b) base-rate neglect; and (c) probability matching or learning.

*Conjunctions.* Subjects evincing the conjunction fallacy report that the conjunction of two events is more rather than less likely to have occurred than one of the events alone. For example, using an example from Tversky and Kahneman (1982b), an individual named Bill is at least as likely to play jazz for a hobby as he is to play jazz for a hobby and happen to be an accountant as well. Yet when Bill is described as a rather dull individual (e.g., as "intelligent but unimaginative, compulsive, and generally lifeless"), the statement "Bill plays jazz for a hobby" is judged less likely by most subjects than the conjunction, the statement

"Bill is an accountant and plays jazz for a hobby." In fact, of course, this conjunction can be no more probable than the less probable of the simple statements "Bill is an accountant" and "Bill plays jazz for a hobby." Can this simple error be appreciated within a behavioral perspective? For example, will the error occur at all if the problem is presented with clear instructions, or with suitable rewards, or with repeated trials? The robustness of the effect is well documented in Tversky and Kahneman (1983) and in scores of subsequent studies, including our own attempts to eliminate it (see Stolarz-Fantino & Fantino, 1990). For example, we reasoned that if our instructions provided a clear set (discriminative stimuli) favoring logical responding, the conjunction fallacy would be eliminated. However, the instruction "Your judgments should be made in terms of their probability and not simply in terms of whatever intuitive appeal is generated by the description above" actually resulted in a larger proportion of our subjects showing the fallacy. It should be noted that our subjects were again UCSD students. These students are selected from the top 12% of high school graduates in California, and subjects in other studies of the fallacy have been comparably elite. Although the occurrence of the fallacy is variable (e.g., Stolarz-Fantino, Fantino, & Kulik, 1996, found large differences across academic classes) it is unquestionably robust. Let me illustrate with two examples from Stolarz-Fantino et al. (1996).

Would subjects with explicit training in logic not commit the fallacy? We studied the students in the penultimate week of a logic course at UCSD, taught by a distinguished philosopher (Patricia Churchland) who introduced me by saying, "You have been studying logic all quarter. Dr. Fantino is now here from the Department of Psychology to give you a test of reasoning." Of 92 subjects, 43% committed the conjunction fallacy.

Would the fallacy be committed when there was no "framing descrip-

tion" to bias the subjects' responses? For example, we asked subjects the following: "Bill is a 34-year-old individual. You know nothing else about him. Mark an X next to the alternative you think is more likely: Bill plays jazz for a hobby; or Bill is an accountant and plays jazz for a hobby." In two experiments of this type (including one in which subjects saw only the one problem without a frame) at least 41% of our subjects committed the fallacy even without the framing description.

These two experiments show that it is not difficult to reveal patterns of illogical responding even in highly educated subjects. Why? There are several possibilities, three of which I would like to explore at least briefly here. First, subjects may respond to compound statements by taking an average of the components they are comprised of. Second, subjects may not be sufficiently motivated to take the problem seriously. Third, subjects may have a natural tendency to respond to compounds more vigorously than consideration of the individual components might warrant (a tendency not necessarily incompatible with the first possibility).

In support of the first possibility, Fantino, Kulik, Stolarz-Fantino, and Wright (1997) applied Anderson's (1981) functional measurement methodology to the judgments of subjects rating conjunctions in which the likelihood of the components was varied. Results were consistent with the possibility that subjects averaged the likelihood of the component statements in arriving at judged likelihood of the conjunction. An advantage of a weighted averaging model is its applicability even to cases in which the fallacy occurs without the framing description.

What if we make it worthwhile for subjects to respond logically? Daniel Zizzo, Stephanie Stolarz-Fantino, Julie Wen, and I have recently completed a series of studies that have departed from prior work in ways that make them more in keeping with the behavioral tradition: We provided feedback

with respect to the logical correctness of subjects' responses, repeated trials, and, in some conditions, a handsome monetary reward (\$3 per correct answer). Four versions of this experiment were conducted. Although we were able to sometimes find statistically significant effects of feedback and of monetary reward, the size and replicability of the effects were behaviorally uninteresting. Moreover, there was no evidence in any of the four experiments for improvement over the six trials. Clearly feedback and monetary reward do not eliminate the conjunction fallacy.

Upon reflection, however, perhaps this lack of a clear effect is unsurprising. We are exploring a complex situation in which subjects have rich histories dealing with the types of situations described, and in that context six trials are few trials. This reflection takes us to a more fundamental question about how subjects respond to complex stimuli. The finding that the conjunction fallacy occurs at a high rate in the absence of a framing description suggests a tendency to overestimate the likelihood of compound or conjunctive events. This leads to the following question: How would subjects react to novel compound stimuli encountered in a thoroughly novel behavioral setting? We can gain a tentative answer to this question by leaving the area of the conjunction fallacy for the moment and considering the results of another experiment with UCSD students, this one using a matching-to-sample (MTS) procedure (Fantino & Savastano, 1996).

Although the stimuli used in our experiment were not analogous to those in conjunction fallacy problems, they were selected with a resemblance in mind. The conjunctive statements of the fallacy problems typically combine high-likelihood and low-likelihood events. We created compounds by combining stimuli correlated with high and low probabilities of reinforcement. Here, unlike research on the conjunction fallacy, there was no logical con-

straint on reaction to the compound. Would subjects treat these compounds as they might a stimulus correlated with a particularly high, intermediate, or low probability of reinforcement? Our subjects responded for three sessions in an MTS procedure in which the sample was usually one of four semicircles of color (red, green, blue, or yellow). Two colors (say red and green) were associated with a high probability of reinforcement for a correct match ( $p = .80$ ; incorrect matches were never reinforced); the other two colors (say blue and yellow) were associated with a low probability of reinforcement for a correct match ( $p = .20$ ). In the first session all subjects received 360 MTS trials. In the second session the 10 "untrained" subjects received an additional 440 trials. For 10 "trained" subjects, however, on one fifth of their trials the sample was a compound stimulus that combined two of the semicircles, one associated with a high probability of reinforcement and the other with a low probability (say red and blue). When this compound was the sample, a correct match was reinforced only 10% of the time (i.e.,  $p = .10$ ). Would this training with a single exemplar compound affect responding to a novel compound encountered subsequently? This question was answered by comparing all subjects' responses to a novel compound (green and yellow) in the third and final session. During this third session we interspersed probe trials in which no sample stimulus occurred and no reinforcement was awarded. On these trials subjects chose between two alternatives: the novel compound stimulus and either one of its constituent stimuli. Instructions prepared subjects for these probes and asked them to "Always choose the one that you think is more likely to earn you a point"; subjects were told that points could be earned during probe trials, but that they would be tallied at the end of the session and that the usual reinforcer message would not occur after these trials. The results were clear: Untrained sub-

jects preferred the compound over either individual stimulus, demonstrating what might be termed *summation*; trained subjects did not prefer the novel compound. Specifically, whereas 8 of 10 subjects in the untrained group chose the compound more than either constituent stimulus (the other 2 treated the compound as intermediate), only 1 of 10 subjects in the trained group did so (5 selected the compound less often than either simple stimulus; 4 treated the compound as intermediate).

This result suggests three conclusions: (a) Subjects tend to respond more to compounds than components, a tendency that may help us to understand the conjunction fallacy. (b) Behavior in this novel setting is readily modifiable in the sense that subjects given training with a single exemplar compound correlated with a lower rate of reinforcement than either of the individual component stimuli no longer responded to the novel compound more than to its component stimuli. (c) MTS is a procedure developed in the behavioral laboratory that may have useful and broad applicability to decision making, a point that is better illustrated with the next topic.

In summary, research on the conjunction fallacy suggests that it is a robust phenomenon, occurring in highly educated subjects, including those with training in logic. Although framing descriptions typically enhance the tendency to commit this fallacy, it occurs with surprising frequency in college students, even without the frame (over 40%). It is likely that the fallacy depends, at least in part, on a tendency to respond more to more complex or compound stimuli than to their simpler components, and to a rich history dealing with complex stimuli.

*Use of base rates.* Base-rate neglect refers to another robust phenomenon in judgment and decision-making research, again one first developed by Kahneman and Tversky (1973). When assessing the probability of a future event, people often ignore background information in favor of case-specific



information. For example, consider the well-known "taxi-cab problem," modified from that presented by Tversky and Kahneman (1982a):

A cab was involved in a hit and run accident at night. Two cab companies, the Green and the Blue, operate in the city. You are given the following data: (a) 67% of the cabs in the city are Blue and 33% are Green. (b) a witness identified the cab as Green. The court tested the reliability of the witness under the same circumstances that existed on the night of the accident and concluded that the witness correctly identified each one of the two colors 50% of the time and failed 50% of the time. What is the probability that the cab involved in the accident was Green rather than Blue?

Obviously in this example the witness is no help at all (i.e., the witness is not reliable). Subjects should say that the probability that the cab is green is 33% (the base rate of green cabs in the city). Instead subjects ignore the base rates and judge the likelihood that the cab is green around 50%, or equal to witness accuracy. Adam Goodie and I wondered about the robustness of base-rate neglect. Although many studies have reported such neglect, they typically did so in a single paper-and-pencil question with no special care taken to insure an attentive and motivated subject. As behaviorists we wondered if we would find base-rate neglect in a behavioral task, and again an MTS procedure appeared to provide a convenient way to address this question. In fact, whereas the use of this procedure with compound stimuli was only somewhat analogous to the conjunction fallacy problem, as typically studied, the MTS procedure allows us to mimic the base-rate problem rather precisely. After all, the base-rate problem is one of multiple stimulus control: control by the sample cue and control by the base rates (or probabilities of reinforcement for choosing either alternative, independent of the sample cue). The MTS procedure allows us to separately manipulate these two sources of stimulus control with repeated trials in a behavioral setting. The sample corresponds to the witness in the taxi problem; the probabilities of reinforcement for se-

lecting the comparison stimuli correspond to the base rates, or incidence of taxi types. Would base-rate neglect occur under such circumstances? If so, could we appreciate it in terms of competing sources of stimulus control? Our procedure was simple. The sample in a MTS task was either a blue or green light. After the sample was terminated, two comparison stimuli appeared: These were always a blue and a green light. Subjects were instructed to choose either. We could present subjects with repeated trials rapidly (from 150 to 400 trials in a less than 1-hr session, depending on the experiment).

We could readily manipulate the probability of reinforcement for selecting either color after a blue sample and after a green sample. Consider the condition corresponding to that in the cab problem above. Let the blue and green samples be equiprobable. Following a blue sample, selection of the blue comparison stimulus is reinforced on 67% of trials and selection of the green comparison stimulus is reinforced on 33% of trials. Following a green sample, selection of the green comparison stimulus is reinforced on 33% of trials and selection of the blue comparison stimulus is reinforced on 67% of trials. In other words the sample in this case has no discriminative (or informative) function, just as the witness testimony had no function in the cab problem. Reinforcement consisted of points (in one experiment the points were backed by money, but the money had no effect on performance).

If our UCSD subjects responded optimally they should have come to select the blue comparison stimulus on every trial, regardless of the sample color, thereby obtaining reward on 67% of trials. However, there is a rich literature on probability matching (or probability learning) which shows that when humans are presented with repeated identical binary choices, each associated with a consistent payoff likelihood, they match their choices to the arranged probabilities (e.g., Humphreys, 1939) instead of maximizing

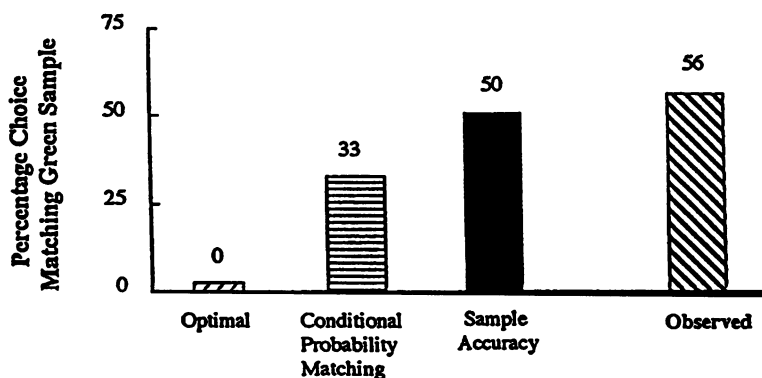


Figure 4. The percentage of trials on which the green sample is matched (when matching green is reinforced on 33% of trials) if a subject were choosing optimally, according to probability matching, if a subject were matching overall sample accuracy, and the observed degree of matching green (data from Goodie & Fantino, 1995).

their payoffs by always choosing the stimulus with the higher reinforcement likelihood. If our subjects probability matched, they should choose blue on 67% of trials (although nonoptimal, this would not be termed base-rate neglect). Finally, if our subjects responded as subjects do on a single paper-and-pencil question such as the cab problem, they should be sensitive to sample accuracy and match the sample 50% of the time, a decidedly nonoptimal strategy.

The three predictions outlined above and the actual data are presented in Figure 4, which gives percentage choice of the green comparison stimulus following a green sample. Note that there is strong evidence for base-rate neglect. Thus, base-rate neglect occurs in a behavioral task and over several hundred trials. Does it occur with pigeons as subjects? Alex Hartl and I (Hartl & Fantino, 1996) studied pigeons in a comparable MTS task and found no evidence of base-rate neglect. Indeed there was no evidence of probability matching. Instead, pigeons responded optimally in this task. What might account for the drastic difference in the behavior of pigeons and college students? We have speculated that humans have acquired strategies for dealing with matching problems that are misapplied in our MTS problem (e.g., Stolarz-Fantino & Fantino, 1995). For

example, from early childhood we learn to match like shapes and colors at home, in school, and at play (e.g., in picture books and in playing with blocks and puzzles), to the extent that we may display insensitivity to changes in the reinforcement contingencies in our present environment sometimes associated with rule-governed, as opposed to contingency-shaped, behavior. In order to assess this possibility, Adam Goodie and I explored the possibility that base-rate neglect is a learned phenomenon. If it is, we should be able to eliminate base-rate neglect by having the sample stimuli be physically unrelated to the comparison stimuli. Thus, we repeated our earlier experiment (Goodie & Fantino, 1995) with an MTS procedure in which the sample stimuli were line orientations and the comparison stimuli were again the colors blue and green. This change eliminated base-rate neglect in keeping with the learning hypothesis (Goodie & Fantino, 1996). Instead our subjects' choices were well described as probability matching. To further assess the learning hypothesis, we next introduced an MTS task in which the sample and comparison stimuli were physically different but were related by an extensive history: The samples were the words *blue* and *green*; the comparison stimuli the colors blue and green. A robust base-rate

neglect was reinstated. These and other experiments led us to conclude with some confidence that base-rate neglect results from preexisting associations between cues and options.

The MTS procedure permits the ready study of variables influencing base-rate neglect, or, more generally, the control of choices by sample accuracy and base rates. Adam Goodie has devised several studies along these lines. I present two such examples, one each from two sets of studies in his dissertation (Goodie, 1997). In one we asked: Will our subjects "catch on" if we extend the number of trials beyond the usual 600 trials that we employed in some of our prior research? In an experiment that replicated our finding that base-rate neglect is a function of the relatedness of sample and comparison stimuli (and also explored other variables), we assessed choice for 1,600 trials. As one might predict, we were rewarded for our patience (and that of our subjects!). The degree of selecting the stimulus correlated with the lower likelihood of reinforcement after its matching sample declined very gradually from above 50% (large base-rate neglect) to well under 30% (below probability matching). Thus, base-rate neglect not only disappears with a sufficient number of trials, but perhaps eventually probability matching would also be eliminated. But the decline is extremely gradual (less than 1% per 100 trials), which probably explains why it was not detected in the many prior studies of probability matching in the previous literature (only one of which extended training beyond 450 trials). I might also add that even after 1,600 trials there was an effect of relatedness: Subjects with related samples and comparison stimuli continued to match at a significantly higher rate than subjects with unrelated samples and comparison stimuli. The effects of prior learning are persistent indeed. Moreover, the eventual disappearance of base-rate neglect in no way minimizes its reality or importance: Life rarely offers 1,600 trials.

MTS also permits us to further assess cases in which potential control by sample accuracy and base rates are pitted against one another. Consider the situation depicted in Figure 5 (top). Note first that when the sample is blue (right side of Figures 5a and 5b) the likelihood that blue is correct and the likelihood that green is correct both equal .50. To repeat, this is true in both 5a and 5b. The situations differ, however, after a green sample: In 5a, green is correct after green with a probability of .80 and blue is correct after green with a probability of .20; in 5b, the situation is reversed with green correct after green with a probability of .20 and blue correct after green with a probability of .80. One might say that in 5a matching is generally reinforced after a green sample, whereas in 5b countermatching (i.e., selecting blue) is generally reinforced after a green sample. How will this dramatic change affect responding after the blue sample (right side of 5a and 5b), that is, after the sample for which the contingencies are unchanged? There are three possibilities, each of which is quite reasonable. First, there should be no change in responding following the blue sample because the contingencies after blue remain unchanged (the right sides of 5a and 5b are identical). Such an outcome would suggest control by the conditional probabilities of reinforcement; that is, subjects' responses to either type of sample (here, blue) would not be affected by events following the other type of sample (here, green). The second possibility is that there should be greater selection of blue (i.e., matching blue) because the base rates for blue overall have doubled from 33% in 5a to 66% in 5b. That is, if we ignore the samples, we see that selection of blue is reinforced on one third of trials in 5a but on two thirds of trials in 5b. Such an outcome would suggest control by the base rates of reinforcement in the entire situation. Third, there should be greater selection of green (i.e., countermatching) because although blue is correct twice as often,





Figure 5. Top: The between-groups design from a study by Goodie (1997), in which the groups differ only in the contingencies of reinforcement for choices following a green cue. This was a matching-to-sample task in which the sample was either green (on 56% of trials) or blue (on 44% of trials). On every trial the sample is followed by a choice between green and blue. The likelihood that a choice of green or blue was reinforced differed after a green sample between Conditions 5a and 5b but was the same (always 50%) after a blue sample in both 5a and 5b. For example in Condition 5a (marked “67%/67%” because the overall sample accuracy is 67% and the base rate of reinforcement with respect to green is also 67%) following a green sample, choice of green is reinforced on 80% of the trials and choice of blue is reinforced on 20% of the trials. In Condition 5b (marked “33%/33%” because both overall sample accuracy and the base rate of reinforcement with respect to green are 33%) following a green sample, choice of green is reinforced in 20% of the trials and choice of blue is reinforced in 80% of the trials. These reinforcement probabilities are indicated to the right of  $p(s^{R+})$ . Bottom: The results show that subjects in the 67%/67% group all matched blue samples more often than subjects in the 33%/33% group, even though they were associated with identical contingencies of reinforcement.

sample accuracy has been deflated (from 67% overall to 33%), thereby promoting less matching to blue. Such an outcome would suggest control by sample accuracy in the entire situation. Indeed, this control would be rather powerful because blue would be selected less often even though it was correct twice as often.

The results are shown for each of 12 subjects (6 exposed to Condition 5a and 6 to Condition 5b) in the bottom of Figure 5. Each filled bar in the bottom of the figure shows the proportion of matching blue (selecting blue after a blue sample) for 1 subject in the corresponding conditions shown immediately above. The two open bars present the mean proportion of matching blue for each group of 6 subjects. There was no overlap between groups. Matching to blue declined in 5b as compared to 5a, indicating another decisive "win" for the control exerted by sample accuracy (and therefore another instance in which base rates exerted little control). Other manipulations in the same set of studies supported the same conclusion.

Can we diminish control by the sample? For example, we can give subjects experience with the base rates of reinforcement while omitting the sample. If so subjects will begin to be influenced by the base rates. David Case, Adam Goodie, and I did such a study. We gave our UCSD subjects 400 trials of binary choice without a sample and then presented 400 trials of MTS (with the same comparison stimuli as in the first block of 400 trials). In both between-subjects and within-subject comparisons, we found behavior more sensitive to base rates in the MTS procedure after this "forced exposure" to the base rates of reinforcement. Note too that when the sample is omitted the procedure becomes a prototypical probability-learning procedure. It is to probability learning that we now turn.

*Probability learning.* Whereas the normal MTS procedure, used above, is directly analogous to the base-rate problems of cognitive psychology,

when the sample is eliminated from the procedure, it becomes analogous to probability learning. Probability learning experiments involve repeated binary choices, as in selection between green and blue comparison stimuli. The scores, even hundreds, of experiments on probability learning in the cognitive literature correspond to a comparable plethora of experiments involving concurrent stimuli in the behavior-analytic literature. As we noted above, people tend to probability match in this situation, rather than behave optimally (which would entail always choosing the stimulus with the greater likelihood of reinforcement). Other than conducting over 1,600 trials, how might we eliminate probability learning and encourage more optimal responding? Prior literature suggests that instructions are one route to improving performance: If subjects are told that trials are random and patternless and that there is a simple solution that should be applied on all trials, then subjects tend to select the richer outcome at levels significantly higher than probability matching. The pattern point is important, because in some of our research fully half of our subjects spontaneously ask if the task is one of pattern detection (including some subjects who had already been told that it is not). I conclude by briefly summarizing some preliminary research from our laboratory that may help us better appreciate the maintenance of probability learning. These unpublished experiments have been largely designed and carried out by Greg Zarow, whose insights and methodological innovations are worth sharing with you.

In our ongoing work, Greg and I have manipulated motivation and also what I term the salience of the probabilities correlated with the outcomes. Our preliminary (but statistically significant) findings from four of these sets of studies follow.

First, we told subjects the actual probabilities involved in a task in which the choices were presented on a computer screen. In another task the

probabilities were self-evident, as in the rolling of a die (e.g., in one condition Greg required subjects to color a die so that it had four black and two white sides). In both cases subjects quickly came to choose the more probable outcome more than 67% of the time (between 79% and 88% of trials, depending on the format); in the case of this die, the mean was 80%. This result may appear trivial but in at least one respect it is intriguing. Subjects, when asked, correctly report the probabilities involved even when they have not been informed about them. This was found in prior work and was confirmed by one of Greg's studies; in fact, our 25 subjects could identify the probabilities rather precisely. In this case the actual probability of the more likely outcome was .75 and the subjects' mean estimated likelihood was 75.1 with a standard error of 2.2. Yet over 192 trials, on the average, these subjects' choices displayed probability matching. It appears, therefore, that when we call attention to the probabilities (by telling subjects what they are or by having them be self-evident as in the case of the die) we sharply modify their behavior beyond what occurs when subjects merely experience the probabilities and can correctly estimate them.

Second, what occurs if between two 96-trial blocks of probability learning (predicting colors on a computer screen with unannounced probabilities of .75 and .25) subjects are asked to recommend a strategy for a hypothetical other subject (we gave subjects no feedback or new instructions after they made their recommendation)? The answer is they appear to do better on the second block, their selection of the richer outcome jumping from 75% to 85% immediately after the break (more than control subjects who were merely asked the time of day). Also, among the subset of the subjects in this study who recommended exclusive preference (i.e., optimal choice) for the hypothetical subject, selection of the

richer outcome jumped from 79% to 87%.

Third, competition with other subjects (or prior experience with competition) also appears to enhance performance. In within-subject comparisons, competition increased the tendency to maximize when it was compared with solo game playing. Subjects with competitive experience in this procedure were also compared with subjects who lacked this experience: The between-groups measure of maximizing we used in this study was 51% in a solo session for the subjects without competitive experience and 87% in a solo session for the subjects with earlier competitive experience.

Finally, when subjects were told that the top 5% of 100 subjects could leave the study early (still earning full credit), their selection of the richer outcome also soared. After performing optimally in 55%, 52%, and 45% of blocks in three prior conditions, subjects given the "top 5%" instruction performed optimally on 81% of trial blocks. Prior research has also suggested that increased motivation may produce somewhat more optimal performance (e.g., Smith & Walker, 1993, using monetary rewards).

Thus, despite its apparent procedural simplicity, a binary-choice procedure produces complex reactions in human subjects. Perhaps the fact that it is simpler and less structured than matching to sample makes it even more susceptible to the effects of unspecified experience in the subjects' rich histories of decision making and problem solving. When the task is arranged in a manner that makes the probabilities involved more salient, behavior in the task becomes more optimal. Our preliminary ongoing research shows a sharp increase in maximizing when the probabilities are made more salient by informing subjects of them (or having subjects construct the probabilities), by having subjects reflect on the probabilities (recommending action for others), or by increasing motivation (competitive situation or providing the oppor-

tunity for an early exit from the experiment). Without such intervention it may take thousands of trials for near-optimal behavior to gradually emerge.

### *Correspondence Between Cognitive and Behavioral Tasks*

The research discussed above utilizes procedures that were designed to correspond to those common in research by cognitive psychologists. If it is granted that we have crafted our procedural correspondences tolerably well, an obvious next question is whether behavior in corresponding situations is indeed governed by the same principles. For example, can behavior in the base-rate neglect problems that are studied by cognitive psychologists be understood in terms of competing sources of stimulus control, where the degree of control is determined by both proximal reinforcement contingencies (such as the rates of reinforcement associated with matching the sample and with selection of each of the comparison stimuli) and more distal contingencies (such as a history of rule following, including those embedded in the "demand characteristics" of the task)? In principle, I believe the answer is yes. But in practice it sometimes may be so difficult to specify the controlling contingencies that a behavioral account may add little to our appreciation of a particular problem. However, even in such cases, I suggest that a behavioral perspective is the one with the most promise for untangling the historical and contemporary determinants of decision making.

### *Conclusions*

The research we have reviewed does not paint a rosy picture in terms of human decision making. People are not always seekers of information (the observing literature), and they often do not utilize information in a logical (conjunction fallacy, base-rate neglect) or optimal (base-rate neglect, probability learning) fashion. The examples of nonoptimal behavior that we have dis-

cussed are all robust, complexly determined, and susceptible to study from a behavioral perspective. We hope that research from this perspective can help us to better appreciate the extent to which, and the conditions under which, we seek and effectively utilize information (i.e., make effective use of discriminative stimuli). If so, behavioral research can supplement research from other perspectives in furthering our understanding of the variables that control decision making.

### REFERENCES

- Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.
- Bloomfield, T. M. (1972). Reinforcement schedules: Contingency or contiguity? In R. M. Gilbert & J. R. Millenson (Eds.), *Reinforcement: Behavioral analyses* (pp. 165–208). New York: Academic Press.
- Case, D. A., Fantino, E., & Wixted, J. (1985). Human observing: Maintained by negative informative stimuli only if correlated with improvement in response efficiency. *Journal of the Experimental Analysis of Behavior*, *43*, 289–300.
- Case, D. A., Ploog, B. O., & Fantino, E. (1990). Observing behavior in a computer game. *Journal of the Experimental Analysis of Behavior*, *54*, 185–199.
- Dinsmoor, J. A. (1983). Observing behavior and conditioned reinforcement. *Behavioral and Brain Sciences*, *6*, 693–704.
- Fantino, E., & Case, D. A. (1983). Human observing: Maintained by stimuli correlated with reinforcement but not extinction. *Journal of the Experimental Analysis of Behavior*, *40*, 193–210.
- Fantino, E., Case, D. A., & Altus, D. (1983). Observing reward-informative and -uninformative stimuli by normal children of different ages. *Journal of Experimental Child Psychology*, *36*, 437–452.
- Fantino, E., Kulik, J., Stolarz-Fantino, S., & Wright, W. (1997). The conjunction fallacy: A test of averaging hypotheses. *Psychonomic Bulletin and Review*, *4*, 96–101.
- Fantino, E., & Logan, C. A. (1979). *The experimental analysis of behavior: A biological perspective*. San Francisco: Freeman.
- Fantino, E., & Savastano, H. I. (1996). Humans' responses to novel stimulus compounds and the effects of training. *Psychonomic Bulletin and Review*, *3*, 204–207.
- Goodie, A. S. (1997). *Base-rate neglect under direct experience*. Doctoral dissertation, University of California, San Diego. Dissertation Abstracts International.
- Goodie, A. S., & Fantino, E. (1995). An ex-

- perentially derived base-rate error in humans. *Psychological Science*, 6, 101-106.
- Goodie, A. S., & Fantino, E. (1996). Learning to commit or avoid the base-rate error. *Nature*, 380, 247-249.
- Hartl, J. A., & Fantino, E. (1996). Choice as a function of reinforcement ratios in delayed matching to sample. *Journal of the Experimental Analysis of Behavior*, 66, 11-27.
- Humphreys, L. G. (1939). Acquisition and extinction of verbal expectations in a situation analogous to conditioning. *Journal of Experimental Psychology*, 25, 294-301.
- Jenkins, H. M., & Boakes, R. A. (1973). Observing stimulus sources that signal food or no food. *Journal of the Experimental Analysis of Behavior*, 20, 197-207.
- Kahneman, E., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237-251.
- Mulvaney, D. E., Dinsmoor, J. A., Jwaideh, A. R., & Hughes, L. H. (1974). Punishment of observing by the negative discriminative stimulus. *Journal of the Experimental Analysis of Behavior*, 21, 37-44.
- Smith, V. L., & Walker, J. M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry*, 31, 245-261.
- Stolarz-Fantino, S., & Fantino, E. (1990). Cognition and behavior analysis: A review of Rachlin's *Judgment, Decision, and Choice*. *Journal of the Experimental Analysis of Behavior*, 54, 317-322.
- Stolarz-Fantino, S., & Fantino, E. (1995). The experimental analysis of reasoning: A review of Gilovich's *How We Know What Isn't So*. *Journal of the Experimental Analysis of Behavior*, 64, 111-116.
- Stolarz-Fantino, S., Fantino, E., & Kulik, J. (1996). The conjunction fallacy: Differential incidence as a function of descriptive frames and educational context. *Contemporary Educational Psychology*, 21, 208-218.
- Tversky, A., & Kahneman, D. (1982a). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgement under uncertainty: Heuristics and biases* (pp. 153-160). Cambridge, U.K.: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1982b). Judgments of and by representativeness. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgement under uncertainty: Heuristics and biases* (pp. 84-98). Cambridge, U.K.: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293-315.
- Wyckoff, L. B., Jr. (1952). The role of observing responses in discrimination learning: Part 1. *Psychological Review*, 59, 431-442.