

Recommendation or Evaluation? Task Sensitivity in Information Source Selection

ANDREW D. GERSHOFF
SUSAN M. BRONIARCZYK
PATRICIA M. WEST*

Ironically, when consumers turn to an information source to assist in decision making, they are faced with the added responsibility of having to make a decision about the information source itself. A normative model is presented that shows that consumer assessment of an information source should include both the task for which an information source is sought and the appropriate probability of success for the task. A key distinction between tasks is whether or not the consumer knows the agents' (or prospective agents') ratings of the alternative in question at the time of his or her assessment of the information source. A series of experiments is presented that examines consumer assessment of the diagnosticity of one type of information source, consumer agents (e.g., movie critics and stock analysts) across the three common tasks of seeking recommendations, seeking evaluations, and choosing between agents who have provided conflicting advice. Results show that consumers frequently select inferior agents for providing recommendations and choose product alternatives that should be avoided because of a failure to recognize when a task calls for a conditional rather than overall assessment of agent prior performance. A final study attempts to isolate the reasons for these shortcomings by examining the process underlying consumer diagnosticity assessment of information sources. Implications and future research to address and improve consumer assessment of information sources across tasks are discussed.

Countless sources of information are available to reduce uncertainty in the choice process and to assist consumers in making predictions about a future consumption event. These information sources range from product attribute information, such as recent company profits to predict future stock performance or genre to predict personal preference for a film, to the opinions of others, such as an analyst's rating of a stock or a film critic's rating of a movie. However, all information sources are not created equal in terms of predictive ability, and thus, consumers must be selective when relying on information sources.

One way to evaluate the diagnosticity of an information source as a predictor of future events is to examine how well it has performed in the past. For instance, Film Critic

A, who has regularly shared a consumer's preferences in the past, is likely to be perceived as providing a more useful film rating than Film Critic B, who has only occasionally shared the consumer's preferences. Yet, this general assessment of an information source's diagnosticity fails to appreciate that the diagnosticity assessment should be task specific.

For instance, Film Critic B may be the superior information source for a subset of the consumer's tastes, such as in a particular genre. Similarly, if the consumer is seeking a recommendation for a movie this coming Friday night, the relevant prior performance to examine for this task is not a film critic's overall success rate but rather the critic's success rate when predicting a positive movie experience. Film Critic B may therefore be the superior match for this task if s/he is a hard to please critic.

The purpose of this research is to examine consumer accuracy in assessing the predictive ability of information sources across decision tasks. Our focus is on the assessment of diagnosticity of the opinions of others, commonly referred to as consumer agents (Solomon 1986; West 1996).¹ A nor-

*Andrew D. Gershoff is assistant professor of marketing at Columbia University, Graduate School of Business, Uris Hall, New York, NY 10027 (ag644@columbia.edu). Susan M. Broniarczyk is associate professor of marketing at the University of Texas at Austin, TX 78712 (Susan.Broniarczyk@bus.utexas.edu). Patricia M. West is associate professor of marketing at the Ohio State University, Columbus, OH 43210 (west@cob.ohio-state.edu). This article is based on the first author's doctoral dissertation, co-chaired by the second and third authors. The authors wish to acknowledge Greg Allenby, Wayne Hoyer, Jonathan Koehler, Barbara Mellers, and William Swann, as well as the reviewers and editors, for their helpful input.

¹The term "agent" is used here to refer to either paid or unpaid individuals, organizations, or systems that assist consumers in the decision-making process through market simplification and guidance. Although sim-

mative model is presented that shows that consumer assessment of an information source should include both the task for which the information source is sought and the appropriate probability of success for the task. In a series of experiments, we then examine three distinct tasks in the decision-making process for which a consumer might seek outside information: product recommendation, evaluation, and product choice. The results show that consumers often do not select the most predictive information source or make appropriate product decisions, because they fail to consider the conditional probability of information source success.

TASK DEPENDENCY OF INFORMATION ASSESSMENT

Consumers commonly turn over some of the decision-making tasks of information search and integration to another who is either more experienced, more knowledgeable, or more capable of expending time and effort in the product search and evaluation process (Solomon 1986, 1999). These others act as agents for the consumer and may include movie critics, financial analysts, and organizations such as *Consumer Reports* magazine (Beatty and Smith 1987; Berning and Jacoby 1974; Kiel and Layton 1981; Urbany, Dickson, and Wilkie 1989). Extant research has focused on consumers' reliance on characteristics of the agent that serve as cues that indicate the agent's ability to perform the task. Prior work has concluded that, depending on the product category, consumers prefer to receive guidance from others with more expertise (Feick and Price 1987; Katz and Lazarsfeld 1955; Rogers 1983; Solomon 1987; Woodside and Davenport 1974) or from those who are perceived to be similar in terms of relevant social, demographic, or self-described characteristics (Brown and Reingen 1987; Feick and Higie 1992; Price, Feick, and Higie 1989; Reingen and Kernan 1986). Presumably, consumers use these agent characteristics as cues to infer the diagnosticity of an agent based on beliefs about the relationship between the particular agent characteristic and performance.

Although in many situations it may be easy for consumers to rely on such cues, it may not always be possible nor in their best interest to do so. In many instances, consumers have very little information available to differentiate information sources. For example, the Internet site Yahoo.com provides stock ratings from multiple investment firms, but it provides no information about the relative expertise of these firms. Similarly, consumers are unlikely to have much information about the similarity between themselves and the movie critics for the *New York Times* or *USA Today*. Thus, consumers must rely on the past performance of agents to evaluate their suitability. Furthermore, if the consumer is seeking an agent who will perform best for a given task, a measure of the agent's actual performance, rather than in-

ference, should be used as a basis for selection. Our focus here is on consumer assessment of the diagnosticity of an agent based on past performance.

Consumers face a number of distinct tasks in their interactions with agents. Three particular tasks that we examine here include determining which agent should provide a recommendation, determining which agent should provide an evaluation of an alternative under consideration, and making a decision in the face of conflicting agent opinions. Such tasks, which arise depending on the decision situation, are seemingly quite similar, yet a key difference among them has important implications for how to assess both source diagnosticity as well as the value of information provided. In particular, these tasks differ in whether or not the consumer knows the agent's (or prospective agent's) evaluation of the alternative when assessing the information source.

Consider a consumer who seeks to invest in the stock market. If the consumer is a naive investor or in the early stages of the decision-making process, s/he may possess minimal knowledge about the set of product alternatives available, the attributes possessed by these alternatives, and the decision criteria to evaluate these alternatives. To simplify the decision process, the consumer may turn to an agent such as a stock analyst for a recommendation. When seeking a recommendation, the consumer has no particular alternative under consideration. The agent serves the consumer by simultaneously providing an alternative together with a positive evaluation for that alternative (Rosen and Olshavsky 1987). When assessing the predictive ability of the agent for a recommendation task, the consumer knows that the agent will provide a positive rating of an alternative; thus, assessing diagnosticity requires only considering the accuracy of the analyst's positive ratings in the past. The accuracy of the agent's past negative ratings becomes irrelevant (see fig. 1, panel a).

If the consumer is more knowledgeable or has progressed farther into the decision-making process, s/he may be actively considering a particular alternative as an investment. The consumer may turn to an information source for an evaluation of this stock. In this case, the consumer's task is to determine which information source is most diagnostic for an evaluation under consideration and hence whose opinion should be sought out first. The agent serves the consumer by providing either a positive or negative rating of the alternative (Solomon 1986). Prior to choosing an agent to provide an evaluation, however, the consumer does not know what the agent's rating will be. Therefore, assessing diagnosticity requires considering the accuracy of both the agent's positive and negative ratings in the past (see fig. 1, panel b).


Of course, the ultimate goal for the consumer is to use the information from the agent to make a product choice. Significant research has shown that consumers seek relatively little prepurchase information (Beatty and Smith 1987; Olshavsky and Granbois 1979). Thus, many consumers may simply seek out the opinion of a single agent before proceeding to a final product choice. Clearly, it is in their

ilar, this is distinct from the term "surrogate consumer," which includes individuals who tend to be paid for their services and perform market manipulation tasks, such as providing access to outlets, negotiation, and completion of transactions (Solomon 1986, 1999).

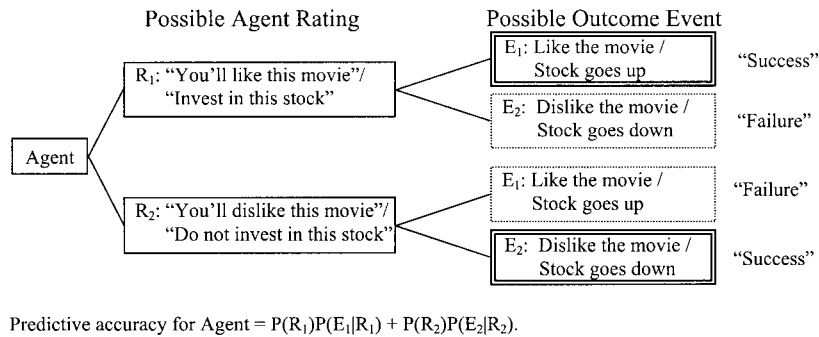
FIGURE 1

TASK DEPENDENCY IN ASSESSING THE DIAGNOSTICITY OF AN INFORMATION SOURCE

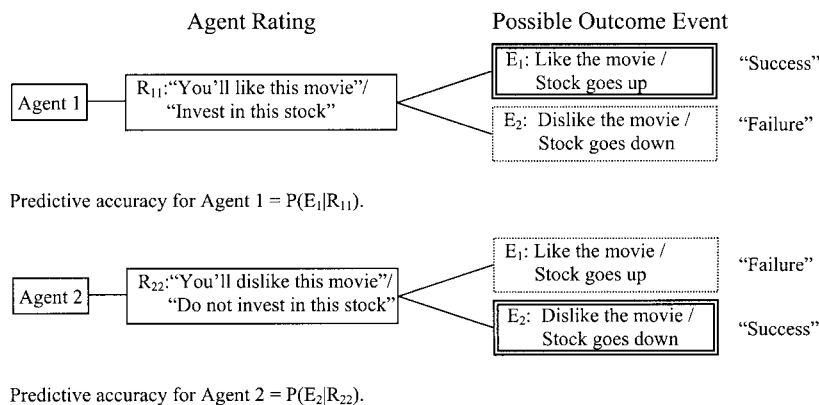
(a) Selecting an Agent for a Recommendation



(b) Selecting an Agent for an Evaluation



(c) Deciding between Two Conflicting Agent Ratings



best interests in this case to select the most diagnostic agent as their information source. However, if a consumer is motivated, s/he may seek the opinions of multiple information sources (West and Broniarczyk 1998).

Consider a situation where a consumer has consulted two analysts for their evaluations of a stock and has found that their opinions conflict (i.e., the first analyst provided a positive rating and the second analyst provided a negative rating). Here the consumer's task is not to determine which agent to approach but instead to choose whether to invest or not invest in the stock. Because the consumer knows the agents' ratings at the time of the investment decision, s/he

should focus on only the relevant past performance information associated with each agents' rating. This suggests considering the accuracy of the first analyst's positive ratings and the second analyst's negative ratings only (see fig. 1, panel c).

To summarize, a key distinction between tasks is whether or not the consumer knows the agents' (or prospective agents') ratings of the alternative in question at the time of his or her assessment of the information source. When selecting an agent for a recommendation or choosing between conflicting agents, the consumer does know the agents' ratings. In contrast, when seeking an evaluation from an in-

formation source, the consumer must make an assessment without knowing what the agents' ratings will be.

The next section will develop a normative model that demonstrates the importance of this distinction in consumer assessment of source diagnosticity and use of information provided. Following the normative model will be a review of the psychological literature on individuals' abilities to properly assess information diagnosticity with particular emphasis on conditional probability assessment. Then three experiments will be presented that examine consumer sensitivity to task dependence in their assessment of agent diagnosticity.

NORMATIVE MODEL

Ironically, consumers turn to agents to reduce effort in decision making but, because agents' abilities may differ, consumers are saddled with the new task of choosing the best agent, or in the case of multiple agents, taking an action based on conflicting opinions. Three factors need to be considered in this process. Of foremost interest in the current study is the appropriate assessment of source diagnosticity, reflecting the probability that the rating provided by the agent will accurately predict the outcome in question. In addition, the consumer must consider the cost of obtaining the information and the expected gain or loss associated with the choice outcome. Evaluating the expected value of a single agent entails careful consideration of possible outcome states (see Mann 1972). Such an evaluation can be described by

$$EV(\text{Agent}) = \sum_i P(R_i) \sum_k [P(E_k | R_i) \pi(E_k)], \quad (1)$$

where $P(R_i)$ is the probability associated with an agent's rating, $P(E_k | R_i)$ is the conditional probability of experiencing an event given the agent has provided a particular rating, and $\pi(E_k)$ is the expected gain or loss associated with the event assuming that the consumer followed the agent's advice. Comparing the expected value of multiple independent agents entails

$$EV(\text{Agent selection}) = \max_{1 \leq i \leq n} [EV(\text{Agent}_i) - c(\text{Agent}_i)], \quad (2)$$

where $c(\text{Agent}_i)$ is the cost associated with obtaining an agent rating. (In the case where agent ratings are not independent, one must correct for redundancy in ratings.)

Our focus is on consumer ability to assess the predictive accuracy of an agent across decision tasks. To isolate this aspect of agent assessment and simplify the decision, we examine situations in which the cost of obtaining a rating is the same for all agents considered. In addition, we consider only situations where the gains and losses associated with "success" (the information source accurately predicts an outcome event) and with "failure" (the information source inaccurately predicts an outcome event) are equivalent. This

TABLE 1
PRIOR PERFORMANCE INFORMATION

Outcome event	Agent's rating	
	Positive rating (R_1)	Negative rating (R_2)
Positive event (E_1)	cell a (success)	cell b (failure)
Negative event (E_2)	cell c (failure)	cell d (success)

reduces the computation of expected value to an assessment of the predictive accuracy, given the decision task and prior performance information.

In the case of a two-outcome event, as illustrated in figure 1, prior performance can be easily represented in a 2×2 contingency table (see table 1). Ratings provided by an agent are represented in columns R_1 and R_2 . These ratings are intended to predict an outcome event, for example, liking or disliking of a movie or movement in price of a stock, represented in rows E_1 and E_2 . The four possible combinations of ratings and outcomes can be classified as "successes" (cells a and d) or "failures" (cells b and c) depending on the accuracy of the agent rating in predicting the outcome event.² The nature of the decision task determines the relevance of prior performance information needed for assessing predictive accuracy (see table 2). The probability of a successful decision is determined using Bayes Theorem (see the appendix).

Our normative model provides a useful guideline for assessing the task dependent posterior probability of a success and failure for a given agent, or for a decision based on conflicting opinions. This is accomplished by using prior performance information to directly form the appropriate sample space (Gavanski and Hui 1992). The prescriptive implications of this analysis reveal the following:

- 1) When obtaining a recommendation, the assessment of agent diagnosticity should be restricted to the subset of prior performance data associated with positive predictions only (see panel A of table 2). A consumer should select the available agent who has the highest probability of success at predicting positive events: $\max [(cell\ a_i)/(cell\ a_i + cell\ c_i)]$.
- 2) When an evaluation of an existing product alternative is sought, the assessment of agent diagnosticity should be based on prior performance data associated with both positive and negative predictions (see panel B of table 2). In this case a consumer should select the available agent who has the highest Overall Probability of Success (OPS): $\max [(cell\ a_i + cell\ d_i)/(cell\ a_i +$

²Signal Detection Theory offers an alternative vocabulary for describing the four cells represented in table 1 that is useful when considering situations where the gains or losses associated with success and failure are not equivalent (Birnbaum and Mellers 1983; McFall and Treat 1999). Successes are classified as "hits" (cell a) and "correct rejections" (cell d), while failures are classified as "misses" (cell b) and "false alarms" (cell c).

TABLE 2
TASK-RELEVANT PRIOR PERFORMANCE INFORMATION

A. Selecting between Two Agents for a Recommendation				
Outcome event	Agent 1 will provide a positive rating Positive rating (R_{11})		Agent 2 will provide a positive rating Positive rating (R_{21})	
Positive event (E_1)	a_1		a_2	
Negative event (E_2)	c_1		c_2	
$P(\text{success} \text{recommendation})$	$a_1/(a_1 + c_1)$		$a_2/(a_2 + c_2)$	
B. Selecting between Two Agents for an Evaluation				
Outcome event	Agent 1 may provide either a positive or negative rating		Agent 2 may provide either a positive or negative rating	
	Positive rating (R_{11})	Negative rating (R_{12})	Positive rating (R_{21})	Negative rating (R_{22})
Positive event (E_1)	a_1	b_1	a_2	b_2
Negative event (E_2)	c_1	d_1	c_2	d_2
$P(\text{success} \text{evaluation})$	$(a_1 + d_1)/(a_1 + b_1 + c_1 + d_1)$		$(a_2 + d_2)/(a_1 + b_2 + c_2 + d_2)$	
C. Deciding between Two Conflicting Agent Ratings				
Outcome event	Agent 1 has provided a positive rating		Agent 2 has provided a negative rating	
	Positive rating (R_{11})	Negative rating (R_{12})	Positive rating (R_{21})	Negative rating (R_{22})
Positive event (E_1)	a_1	b_1	a_2	b_2
Negative event (E_2)	c_1	d_1	c_2	d_2
	$P(\text{positive event})/P(\text{negative event}) = P(E_1 R_{11})/P(E_2 R_{22}) = (a_1/c_1) \times (b_2/d_2) \times P(E_2)/P(E_1)$			

NOTE.— $R_{..}$ represents either a positive or negative rating by either agent 1 or agent 2; $E_{..}$ represents either a positive or negative outcome event experienced by the consumer; and $a, b, c,$ and d represent the cells in a 2×2 contingency table (see table 1).

cell $b_i + \text{cell } c_i + \text{cell } d_i$].

- When faced with conflicting opinions from two agents, the consumer should consider the posterior probabilities associated with a positive and negative event occurring by updating the prior probabilities of the events using the appropriate conditional probabilities for each agent (see panel C of table 2). The consumer should then act in accordance with the event having the higher posterior probability. Equivalently, if the consumer has no prior information about the probability associated with a positive and negative event occurring (i.e., $P(E+) = P(E-) = .50$), then he/she can act in accordance with the agent whose Conditional Probability of Success (CPS) is higher: $\max[(\text{cell } a_1)/(\text{cell } a_1 + \text{cell } c_1), (\text{cell } d_2)/(\text{cell } b_2 + \text{cell } d_2)]$. See the appendix.

The normative model is a prescriptive model of what data subjects should be using to assess information source diagnosticity. We do not contend that our model is a descriptive model. Rather, the model derivation demonstrates that a simpler way to assess agent diagnosticity is to use an agent's rating to directly form a sample space that includes only the subset of outcomes specified by the rating and directly assess the probability from this revised sample (Gavanski and Hui 1992). Specifically, our normative model reveals that, once the agent has given a rating of the product, the appropriate

assessment of agent diagnosticity is restricted to the subset of prior performance data consistent with that rating. Consumers must recognize that certain tasks mandate the agent providing an a priori rating and, hence, assessment of this agent's diagnosticity should involve a conditional probability assessment.

CONDITIONAL PROBABILITY LITERATURE

The psychology literature is replete with research that examines individuals' assessment of probabilities (for a review, see Fischhoff and Beyth-Marom [1983]). Research examining conditional probability assessment reveals that people frequently stray from a normative solution because they fail to identify the set of possible outcomes associated with the conditionalizing event. Numerous studies have shown that subjects utilize a Positive Hits (PH) strategy where they rely exclusively on cell a of a 2×2 contingency table when assessing covariation (Arkes and Harkness 1983; McKenzie 1994; Nisbett and Ross 1980; Shaklee and Mims 1982; Shaklee and Tucker 1980). A structurally similar outcome was found by Pollatsek et al. (1987). In this case, subjects mistakenly substituted calculations of the joint probability of two events, representing the co-occurrence of

the two events $P(A \cap B)$, for the conditional probability, which is the probability of one event taking place knowing that a second event has already occurred, $P(A|B)$.

In the current context, these errors would translate into a failure to appreciate how the nature of the task dictates the relevance of prior performance information. A simple comparison of the frequency of instances in cell a for two agents (cell a_1 vs. cell a_2) or a comparison based on the joint probability, $P(E+ \cap R+)$, will not always yield the normatively correct response for any of the three tasks we are considering.

Other researchers have found that, when required to provide a conditional probability of A given B, subjects commonly substitute its inverse, B given A (Bar-Hillel 1983; Dawes 1988; Eddy 1982; Gigerenzer and Murray 1987; Meehl and Rosen 1955; Nisbett and Ross 1980; Wyer 1977). In our case, an inverse conditional error would imply calculating the probability that an agent provided a positive rating when the experienced outcome was known to be positive, $P(R+ | E+) = \text{cell } a / (\text{cell } a + \text{cell } b)$, rather than computing the probability of experiencing a positive outcome once it is known that an agent has provided a positive rating, $P(E+ | R+) = \text{cell } a / (\text{cell } a + \text{cell } c)$.

This poor performance on explicit conditional probability assessment tasks has been attributed to ambiguity of problem wording (Einhorn and Hogarth 1986; Macchi 1995; Pollatsek et al. 1987) and to the presentation format of the prior history information (Gigerenzer and Hoffrage 1995; Kleiter 1994). In their critique of the biases associated with probability assessment, Gigerenzer and Hoffrage (1995) found that these biases were greatly attenuated when information in the problem was represented in the more natural form of frequencies of successes and failures instead of the experimentally used probability form. This would suggest that if agent prior performance data is presented in frequency form (i.e., contingency table), subjects may be able to correctly assess conditional probabilities.

The aforementioned research was generally conducted in an experimental setting where the explicit subject task was to assess conditional probabilities from a given set of relevant information. Our recommendation and conflicting agents tasks additionally require that subjects recognize the need to select only a subset of past agent performance, thereby ignoring other instances of agent performance that are irrelevant to the task. Numerous studies have demonstrated that irrelevant information can influence judgment, particularly when the irrelevant information is relevant for other, similar problems (Castellan 1973; Glazer, Steckel, and Winer 1992; Hilton and Fein 1989; Hutchinson and Alba 1991). Past instances where an agent has provided a negative rating, although irrelevant for a recommendation task, are in fact quite relevant for assessing the overall ability of the agent, which is appropriate in the case of an evaluation task. Thus, consumers may find it difficult to exclude these task-irrelevant instances when assessing information source diagnosticity.

This inability to ignore irrelevant information is mani-

fested in conditional probability tasks by subjects forming inappropriate sample spaces. Research by Slowiaczek et al. (1992) has shown that subjects were good at asking diagnostic questions but were poor at assessing the diagnosticity of answers to these questions. Subjects apparently failed to appreciate that, when assessing the diagnosticity of the answers, only the subset of data consistent with each answer was now relevant. Instead, they tended to rely on strategies consistent with assessment of overall diagnosticity.

This reliance on overall diagnosticity may be driven by how consumers categorize their past interactions with information sources. Consumers may naturally categorize outcomes associated with an agent as successful or unsuccessful, because this categorization has general applicability for many domains encountered in the world (Rosch 1972; Rosch and Mervis 1975; Troler and Hamilton 1986). Experiments in probabilistic inference have shown that subjects rely on their own natural categories when forming sample spaces for probability assessment rather than the appropriate categories for the task (Brase, Cosmides, and Tooby 1998; Sherman, McMullen, and Gavanski 1992).

If subjects rely on the natural category of successes and failures, then they may erroneously substitute the overall probability that an agent is successful for the appropriate conditional probability. This sample space should be restricted to include only positive successes and positive failures when an agent provides a positive rating (for a recommendation task or when considering conflicting opinions) and negative successes and negative failures when the agent provides a negative rating (when considering conflicting opinions).

To summarize, it is predicted that most subjects will perform well when selecting an agent for providing an evaluation of an alternative, because computing the normatively appropriate OPS is consistent with subjects' natural tendencies. On the other hand, subjects are predicted to fare poorly in tasks such as seeking an agent to provide a recommendation or making a decision when two agents provide conflicting evaluations of an alternative. Each of these tasks requires consideration of the CPS at predicting only positive or negative events. It is expected that subjects will be more likely to rely on an OPS strategy rather than the appropriate CPS strategy and will frequently select inferior agents resulting in inferior product decisions.

We will present three studies that examine consumer sensitivity to task dependence in information assessment. Study 1 is a series of experiments that compare consumer assessment of agent diagnosticity for the tasks of recommendation versus evaluation. Study 2 will examine the task of product choice given conflicting agent opinions. Study 3 will attempt to isolate the process underlying consumer assessment of information source diagnosticity.

STUDY 1

The purpose of study 1 is to examine how individuals select an agent for the tasks of providing an evaluation versus a recommendation when presented with historical

frequency data of potential agents' successes and failures. The three-part execution was employed primarily to test robustness of findings across product category and information presentation formats. Since prior research has shown that consumers prefer agents similar to themselves for subjectively evaluated product categories (Feick and Higie 1992; Price et al. 1989), both subjective (movies in studies 1a and 1c) and objective (stocks in study 1b) product categories were examined. In addition, information presentation format included both 2×2 summary tables (studies 1a and 1b) and serial presentation (study 1c).

The method used in the studies presented here was first introduced by Shaklee and Tucker (1980), and it has been used subsequently by numerous others (Arkes and Harkness 1983; Harkness, DeBono, and Borgida 1985; Shaklee and Mims 1982). In these studies, subjects were exposed to a deliberately designed series of contingency estimation tasks that allowed a decision rule to be inferred through observation of the pattern of responses. In studies 1a, 1b, and 1c, subjects' decision strategies were similarly inferred from the pattern of choices they made when exposed to a series of options between agents whose prior histories are made available. The primary focus was to determine whether individuals' choices involving agents would involve conditional (CPS) or overall (OPS) probability strategies that are normatively appropriate for the decision-making task.

Study 1a

To examine subjects' strategies for selecting agents, participants in study 1a were provided with a scenario in which they were asked to guide a target customer who frequented a particular video store to one of the store's two on-duty clerks to act as an information agent. In the evaluation condition, subjects were told that the target customer had picked up a movie that he knew nothing about. The task was to select a clerk to provide an evaluation of whether he would like the movie under consideration. In the recommendation condition, subjects were told that the target customer had no movie in mind. The task was to select a clerk to provide a recommendation for a movie that he would enjoy.

Following the instructions, each subject was exposed to three agent choices in which they were asked to select the best clerk for the target customer. The instruction to select a clerk to provide a recommendation or an evaluation was repeated twice on each page. In each agent choice, subjects were provided with two 2×2 matrices of preference histories for 100 movies seen by both the target consumer and each of the two video store clerks. A pretest was conducted to ensure understanding of the information presented in the 2×2 summary tables. Subjects provided with an example summary table reported the number of movies where the clerk and the customer agreed, the number of movies the customer liked out of those that the clerk liked, and the number of movies that the customer did not like out of those that the clerk did not like. Sixty-four of 65 subjects (98.5%) accurately answered these questions. The matrices for the agent performance histories and comparative success rate

values for each of the clerks in the three agent choices are presented in table 3. The first two agent choices were deliberately designed to discriminate between use of OPS and CPS strategies. In addition, the design allowed for discrimination between these and other potential strategies, such as reliance on the PH strategy or on an Inverse Conditional Probability of Success (ICPS) strategy.

For example, selection of Clerk 1 in both Agent Choice 1 and Agent Choice 2 rules out a strategy consistent with consideration of the CPS, as well as other strategies including PH and the ICPS. This choice pattern would be consistent with consideration of the OPS strategy, because the overall probability is greatest for Clerk 1 in both agent choices.

A third agent choice was also included to rule out the explanation that subjects are aware of and capable of using a CPS strategy but merely prefer an OPS when it is possible to use it. This choice was designed to examine whether subjects would recognize and use the normatively appropriate CPS strategy in the recommendation condition when the OPS is the same for both agents and is therefore not useful to discriminate between agents. Finally, after making all of their choices, subjects responded to a question that asked them to examine a list of potential strategies and select the one they used to choose between agents. These measures allowed for a double check of the choice measures as well as provided data for insight into other potential strategies used by subjects. A similar method of examining outcome and process has been used in Bayesian reasoning experiments (Gavanski and Hui 1992; Gigerenzer and Hoffrage 1995).

Method. One hundred and ninety-one subjects were recruited from introductory undergraduate marketing courses to participate in exchange for extra credit. On arrival, subjects were given one of the stimuli booklets and randomly assigned to the recommendation or evaluation condition. To motivate accuracy, subjects were informed that the top performer would receive a \$20 gift certificate at a local movie theater. Each booklet consisted of a title page, an instruction page that described the recommendation or evaluation task, and the three agent choices. The order of presentation of agents on each page and the order of the pages in the stimulus booklets were counterbalanced. Following the agent choices, subjects selected the strategy they felt best represented the way they had evaluated the agents. Finally, half of all subjects received a manipulation check to ensure that subjects understood their agent task.

Results. Ninety-six subjects completed the evaluation task, and 95 subjects completed the recommendation task. Eighty-five percent of subjects correctly identified the task they were asked to complete. Analyses including all subjects were compared with analyses including only subjects who correctly identified the task. No comparable differences were found. Thus, the results of all subjects' patterns of Agent Choice 1 and Agent Choice 2 are presented in table 4 and reported below.

TABLE 3
STIMULI MATRICES AND POTENTIAL MEASURES OF AGENT ABILITY FOR STUDY 1a

Consumer	Agent choice 1			
	Clerk 1		Clerk 2	
	Thumbs up	Thumbs down	Thumbs up	Thumbs down
Matrices presented to subjects:				
Thumbs up	35 movies	15 movies	15 movies	35 movies
Thumbs down	15 movies	35 movies	3 movies	47 movies
Potential measures of diagnosticity used by subjects:				
Overall probability of success	70% = (35 + 35) / 100		62% = (15 + 47) / 100	
Conditional probability of success	70% = 35 / (35 + 15)		83% = 15 / (15 + 3)	
Positive hits	35		15	
Inverse conditional probability of success	70% = 35 / (35 + 15)		30% = 15 / (15 + 35)	
	Agent choice 2			
Matrices presented to subjects:				
Thumbs up	35 movies	15 movies	47 movies	3 movies
Thumbs down	15 movies	35 movies	35 movies	15 movies
Potential measures of diagnosticity used by subjects:				
Overall probability of success	70% = (35 + 35) / 100		62% = (15 + 47) / 100	
Conditional probability of success	70% = 35 / (35 + 15)		57% = 47 / (47 + 35)	
Positive hits	35		47	
Inverse conditional probability of success	70% = 35 / (35 + 15)		94% = 47 / (47 + 3)	
	Agent choice 3			
Matrices presented to subjects:				
Thumbs up	47 movies	3 movies	15 movies	35 movies
Thumbs down	35 movies	15 movies	3 movies	47 movies
Potential measures of diagnosticity used by subjects:				
Overall probability of success	62% = (47 + 15) / 100		62% = (15 + 47) / 100	
Conditional probability of success	57% = 47 / (47 + 35)		83% = 15 / (15 + 3)	
Positive hits	47		15	
Inverse conditional probability of success	94% = 47 / (47 + 3)		30% = 15 / (15 + 35)	

Within the evaluation condition, the most frequent response (74%) was consistent with the use of the normative OPS strategy. Moreover, use of this strategy was exhibited significantly more than any of the nonnormative strategies ($Z = 7.97, p < .001$).³ Of the minority of subjects who did not exhibit a normative choice pattern, 8.3% chose consistent with the CPS strategy and 14.6% chose consistent with either a PH or an ICPS strategy.

Our main prediction was that, in the recommendation condition, which required a conditional probability assessment, subjects would be more likely to rely on an OPS rather than the normative CPS. Consistent with this prediction, the largest proportion of subjects (47.4%) in the recommendation con-

dition used a strategy consistent with the OPS. Moreover, the proportion using the OPS strategy was significantly greater than the 22.1% of subjects choosing consistent with the normative CPS strategy ($Z = 2.77, p < .006$).

Thus, significantly more subjects used a normative strategy when assessing agent diagnosticity for an evaluation (74%) than for a recommendation (22.1%) task ($\chi^2(1) = 51.42, p < .001$). The majority of subjects in the recommendation task appeared to fail to realize that the task required a conditional assessment of probability. Yet, it should be noted that there was some indication of consumer sensitivity to task dependence, with more subjects choosing consistent with the CPS strategy in the recommendation task than in the evaluation task (22.1% vs. 8.3%; $Z = 2.65, p < .01$, respectively).

The third agent choice was included in the recommendation condition to provide a strong test of aversion to the normative CPS strategy. Specifically, the third agent choice allowed for teasing apart whether results of the first two agent choices were due to a true failure to recognize and use the normative CPS strategy from a simple preference for an OPS strategy. If subjects had selected the agent with the greater conditional probability of success in the third

³To test for significant differences in proportions when more than two categories are possible, a Z-score for the difference is calculated by the following formula:

$$Z = \frac{\hat{\delta}}{SE \hat{\delta}}, \text{ where } P_1 - P_2 \text{ and}$$

$$SE \hat{\delta} = \sqrt{\frac{P_1(1 - P_1)}{n} + \frac{P_2(1 - P_2)}{n} + \frac{2P_1P_2}{n}}$$

TABLE 4
AGENT CHOICE PATTERNS FOR STUDIES 1a, 1b, AND 1c

Strategy consistent with choice pattern	Study 1a: Movies with matrix presentation		Study 1b: Stocks with matrix presentation		Study 1c: Movies with serial presentation	
	Evaluation (n = 96)	Recommendation (n = 95)	Evaluation (n = 61)	Recommendation (n = 59)	Evaluation (n = 22)	Recommendation (n = 24)
Overall probability of success (%)	74.0 ^a	47.4	59.0 ^a	37.3	77.3 ^a	79.2
Conditional probability of success (%)	8.3*	22.1 ^{a, *}	13.1*	33.9 ^a	4.5*	12.5 ^{a, *}
Positive hits/inverse conditional probability of success(%)	14.6*	24.2*	21.3*	13.6*	18.1*	8.3*
Other (%)	3.1*	6.3*	6.6*	15.3*	0*	0*

^aIndicates the normative strategy for the task.

*Indicates difference from Overall Probability of Success Strategy within task ($p < .05$).

agent choice, an argument could be made that they were aware and capable of using the CPS strategy in all agent choices but that they preferred to rely on the OPS strategy. In fact, less than half of the subjects (37.9%) preferred the agent with the greater CPS ($Z = 2.37, p < .02$). This suggests that most subjects simply do not consider the appropriate conditional probability for this task.

The process measure results are presented in table 5. An examination of subjects' self-reports together with their choice patterns reveals that subjects who report using OPS (77%) and CPS (64%) strategies tend to select agents consistent with the strategy reported and that this did not differ by condition ($Z = .87, p = .383$). Consequently, process measure results are similar to the choice measures. In the evaluation condition, the most frequently reported strategy was the normative OPS (39.6%). Subjects indicated using this OPS strategy significantly more than the PH (15.6%) or CPS (2.1%) strategies (p 's $< .001$).

The pattern of self-reports in the recommendation condition was strikingly similar to the evaluation condition, indicating a failure to appreciate the task dependent nature of information assessment. The largest proportion of subjects (25.3%) reported using a strategy consistent with the OPS. This was greater than both the next most frequently reported strategy of PH (12.6%; $Z = 2.05, p < .05$) and the normative CPS strategy (11.6%; $Z = 2.25, p < .02$). Thus, a significantly greater proportion of subjects reported using a normative strategy when seeking evaluations (39.6%) than when seeking recommendations (11.6%; $\chi^2(1) = 19.63, p < .001$).

Study 1a revealed that, when seeking an evaluation, most subjects will use and report a normative strategy consistent with finding the agent with the greatest OPS. Although, when seeking a recommendation, some subjects do consider the conditional probability, most do not. Instead, subjects tend to seek recommendations nonnormatively by relying on a strategy consistent with selecting the agent with the highest OPS.

Study 1b

Method. The method for study 1b was largely similar to that of study 1a with three exceptions. First, in place of

video store clerks, the agents were described as stock analysts who had provided their expectations of the direction of stock movement in the past. Outcomes were described as objective positive or negative movements in the stocks' values rather than as subjective ratings of movies. Subjects either selected an agent for an evaluation of a stock that they were considering buying more of, or selling; or they selected an agent to provide a recommendation of a stock they should buy. Second, the history of agent ratings was increased from 100 to 500 to address subjects' potential avoidance of a CPS strategy due to a perceived insufficient sample of prior agent performance. Third, the task was shortened to include only the first two agent choice tasks.

Results. One hundred and twenty subjects participated in study 1b in exchange for extra credit in their undergraduate introductory marketing course. Nearly all subjects (94.2%) correctly indicated the task that they had been instructed to complete, with no difference between conditions. The results of subjects' choice patterns are presented in table 4.

In the evaluation condition, use of a strategy consistent with the normative OPS was more frequent (59%) than any other pattern ($Z = 3.62, p < .001$). Of the minority of subjects who did not exhibit a normative choice strategy, 13.1% chose consistent with a CPS strategy and 21.3% chose consistent with either a PH or ICPS strategy.

In the recommendation condition, 37.3% of subjects exhibited a choice consistent with an OPS strategy, 33.9% with a CPS strategy, and 13.6% with either a PH or ICPS strategy. Thus, the evidence again clearly indicates that the majority of subjects do not focus on the appropriate conditional probability when selecting an agent for a recommendation ($Z = 2.49, p < .01$). The proportion using the normative CPS strategy was not significantly different from the proportion using the nonnormative OPS strategy ($Z = .31, p = .756$).

Consistent with the normative model, we do observe increased evidence of the CPS strategy in the recommendation as compared with the evaluation condition (33.9% vs. 13.1%; $Z = 2.69, p < .01$); thus, there is some indication of consumer sensitivity to task dependence in assessing agent

TABLE 5
SUBJECTS' REPORTED EVALUATION PROCESS FOR STUDIES 1a AND 1b

Strategy reported	Study 1a: Movies with matrix presentation		Study 1b: Stocks with matrix presentation	
	Evaluation (n = 96)	Recommendation (n = 95)	Evaluation (n = 61)	Recommendation (n = 59)
Overall probability of success (%)	39.6 ^a	25.3	44.3 ^a	37.3
Conditional probability of success (%)	2.1*	11.6 ^{a, *}	24.6	28.8 ^{a, *}
Positive hits (%)	15.6*	12.6*	3.3*	6.8*
Inverse conditional probability of success (%)	6.3*	11.6*	4.9*	8.5*
All other/ unidentified (%) ^b	36.4	38.9	23.9	18.6

^aIndicates the normative strategy for the task.
^bNo single other reported strategy was greater than 12% in either study.
^{*}Indicates difference from Overall Probability of Success Strategy within task ($p < .05$).

diagnosticity. However, as compared with those in the evaluation condition, subjects in the recommendation condition exhibited weak sensitivity to the task, with significantly fewer subjects using a normative strategy (33.9% vs. 59%; $\chi^2(1) = 7.60, p < .01$).

The process measure results, presented in table 5, provide even stronger evidence of subjects' failure to recognize the conditional nature of a recommendation task. In the recommendation condition, significantly more subjects reported use of the nonnormative OPS strategy (37.3%) than the normative CPS strategy (28.8%; $Z = 2.49, p < .01$). Moreover, the proportion who reported using a CPS strategy did not differ between the recommendation and evaluation conditions (28.8% vs. 24.6%; $Z = .52, p > .5$). Again, subjects demonstrated self-insight, with the majority of subjects selecting agents consistent with their reported strategy (OPS, 64%; CPS, 77%).

The results of study 1b are largely consistent with those of study 1a. A greater proportion of individuals use a normative strategy when the task requires consideration of an OPS. For a task that requires consideration of a CPS, however, only a minority of subjects use a normative strategy. Instead, the largest proportion of subjects rely on the OPS. However, the extent of nonnormative agent selection appears to have been attenuated by the objectively evaluated outcome associated with the stock analyst scenario and the increase in the number of prior instances of agent ratings.

Study 1c

Method. The 2 × 2 summary table presentation format in studies 1a and 1b was deliberately selected because it has been shown to be associated with the least amount of error by subjects in studies of covariation assessment (Nisbett and Ross 1980). The intent was to provide a strong test of subjects' use of CPS and OPS strategies by presenting information in a format that would best facilitate subjects' performance. However, summary tables may represent a departure from the way information is likely to be available in real world settings. For this reason, study 1c was conducted to test the robustness across information presentation

formats. Specifically, study 1c differed from study 1a by presenting subjects with agent performance information arranged in serial format. Each instance of the agents' ratings and the subsequent consumer outcomes were individually presented without the benefit of summary table organization, which is closer to how the information might be available through natural word-of-mouth experiences. In addition, only choice measures were collected in study 1c.

The serial presentation format also allows subjects to form their own categories from the prior performance data. If subjects' naturally categorize the data in terms of "successes" and "failures," rather than "positive successes" and "negative successes," we would expect even fewer subjects' to adopt the normatively appropriate CPS strategy for agent choice in the recommendation task than in studies 1a and 1b.

Results. The results of study 1c are presented in table 4. In the evaluation condition, 77.3% of subjects selected agents in a manner consistent with the use of the normative OPS strategy significantly more frequent than any other pattern ($Z = 3.57, p < .001$). But in the recommendation condition, as was found in both studies 1a and 1b, most subjects appeared insensitive to the conditional nature of the recommendation task when assessing agent diagnosticity. In the recommendation condition, the majority of subjects exhibited a choice pattern consistent with the OPS (79.2%), and this was significantly larger than that using the normative CPS (12.5%) strategy ($Z = 4.76, p < .001$). Furthermore, there was no significant increase in subjects' use of the CPS strategy when seeking a recommendation (12.5%) than when seeking an evaluation (4.5%; $Z = .96, p < .3$). If anything, these results suggest that a serial, rather than summary, presentation of agent prior performance is likely to exacerbate subjects' failure to focus on the appropriate conditional probability when selecting an agent for a recommendation, which is consistent with the notion that subjects' naturally categorize prior performance information in terms of successes and failures.

Discussion of Studies 1a, 1b, and 1c

Studies 1a, 1b, and 1c provided subjects with tasks that required consideration of either an OPS or a CPS through

scenarios involving selection of agents to provide either evaluations or recommendations. Across both subjectively and objectively evaluated product categories, larger and smaller sets of histories of agent ratings, and summary table and serial presentation format, subjects were largely insensitive to the differences in the probability requirements of the task.

In all three studies, the majority of the subjects in the evaluation condition made their choices in a manner that was consistent with consideration of the normative overall probability of success. Although studies 1a and 1b do provide evidence that some subjects are sensitive to the conditional nature of a recommendation task, the majority of subjects in all three studies did not make their selections in a normative manner. The largest proportion of subjects made choices consistent with seeking the greatest OPS.

Thus, subjects are quite proficient at selecting a better agent when agents' ratings are unknown at the time they make their selection, as is the case in the evaluation condition. However, most subjects appear to not be sensitive to the conditionalizing information when the agents' ratings are available at the time of selection, as in the recommendation condition. These findings are consistent with prior research in the assessment of diagnosticity (Slowiaczek et al. 1992). Still, it would be premature to suggest that such an effect is robust across all tasks in which conditionalizing information is known a priori.

The tasks presented in studies 1a, 1b, and 1c required subjects to evaluate prospective agents and then to choose only one to provide either an evaluation or a recommendation, depending on condition. While such a task is common in the real world, at other times, people seek out the advice of more than one agent. When two agents' ratings conflict, the consumer must decide how to act. In this situation, as in the recommendation situation, each agent's rating is known at the time the consumer makes the choice. Thus, it is normatively appropriate to consider the conditional probability of each agent's success, given the ratings they have provided.

Two factors about situations involving conflicting agents may cue subjects to the conditional nature of the task and result in an increase in the proportion of subjects using a normative CPS strategy. First, the fact that the agents are providing different ratings, one positive and one negative, means that the agents themselves are clearly providing different information. This difference in information provided by the agents may cue subjects to consider that different prior information may be useful for evaluating each agent. Second, both recommendations and evaluations require approaching an agent to acquire a piece of unknown information. For an evaluation, the unknown information is a positive or negative rating. For a recommendation, it is an alternative that carries a positive rating. Potentially, the conditional nature of a recommendation is hidden because the positive rating is merely implicit in the alternative provided by the agent. In contrast, a rating provided by an agent in an evaluation is explicit. The positive or negative rating is

not assumed; it is stated. Thus, it may be more likely to cue the normatively appropriate conditional processing strategies. Study 2 examines these possibilities in a conflicting agent task.

STUDY 2

Method

The stimuli in study 2 were similar to those used in studies 1a and 1b, with subjects making a series of choices to reveal their selection strategy using information presented in 2×2 tables. Both a video store clerk and a stock analyst scenario were used for purposes of generalization. In the video store scenario, subjects had to decide whether a customer should or should not see a particular movie under consideration. Subjects based their decisions on ratings of the movie provided by two agents. One agent was described as providing a positive rating of the movie being considered. The other provided a negative rating. As in studies 1a and 1b, histories of agent successes and failures were provided. In the stock analyst scenario, subjects had to decide whether a stock under consideration should be bought or sold, and they were provided with conflicting ratings of the stock by two analysts. Thus, rather than selecting an agent as they did in study 1, subjects selected a see/don't see or buy/sell action based on the information provided by the agents.

Once again, subjects' pattern of choices allowed for discrimination among possible strategies (see table 6). For example, selection to "not see" the movie in choice 1 but to "see" the movie in choice 2 would be consistent with an OPS strategy, because the overall probability is greatest for the agent providing a negative rating in choice 1 and a positive rating in choice 2. Conversely, a "see" in choice 1 but "not see" in choice 2 pattern would be consistent with a CPS strategy. After choosing an action, subjects provided a self-reported process measure by selecting from a list of possible agent- information-use strategies.

Results

Fifty-eight subjects participated in the video store condition, and 60 subjects participated in the stock analyst condition, for a total of 118 subjects. The results are presented in table 7. No significant difference between the distributions of subjects' use of choice strategies in the stock analyst and video clerk conditions was found ($p > .10$), and so results were pooled across the conditions.

Only 22% of subjects selected actions that were consistent with consideration of the agents' CPS. Most subjects (78%) exhibited patterns consistent with nonnormative strategies ($Z = 6.08, p < .001$). The largest proportion of subjects (58.5%) made choices in a manner consistent with consideration of agents' OPS, and this was a significantly greater proportion than that of subjects using a CPS strategy ($Z = 4.84, p < .001$). Thus, despite our making the conditional nature of the task more explicit, more than twice as many subjects continued to rely on an overall assessment

TABLE 6

STIMULI MATRICIES AND POTENTIAL MEASURES OF AGENT ABILITY FOR STUDY 2 (VIDEO CLERK PRESENTATION)

Consumer	Agent choice 1			
	Clerk 1 said "thumbs up"		Clerk 2 said "thumbs down"	
	Thumbs up	Thumbs down	Thumbs up	Thumbs down
Matrices presented to subjects:				
Thumbs up	16 movies	34 movies	35 movies	15 movies
Thumbs down	3 movies	47 movies	14 movies	36 movies
Potential measures				
of diagnosticity used by subjects:				
Overall probability of success	63% = (16 + 47) / 100		71% = (35 + 36) / 100	
Conditional probability of success	84.2% = 16 / (16 + 3)		70.6% = 36 / (36 + 15)	
Positive hits	16		35	
Inverse conditional probability of success	32% = 16 / (16 + 34)		70% = 35 / (35 + 15)	
	Agent choice 2			
Matrices presented to subjects:				
Thumbs up	34 movies	16 movies	46 movies	4 movies
Thumbs down	14 movies	36 movies	35 movies	15 movies
Potential measures				
of diagnosticity used by subjects:				
Overall probability of success	70% = (34 + 36) / 100		61% = (46 + 15) / 100	
Conditional probability of success	70.8% = 34 / (34 + 14)		78.9% = 15 / (15 + 4)	
Positive hits	34		46	
Inverse conditional probability of success	68% = 34 / (35 + 15)		92% = 46 / (46 + 4)	

of agent diagnosticity than on the appropriate assessment of agent diagnosticity restricted to the subset of prior performance data consistent with the agent's rating.

Self-reports of decision making (see table 8) also revealed that subjects were not appropriately sensitive to the conditional nature of choosing between conflicting agents. Even though the conditionalizing information was made explicit, subjects were equally likely to report using an OPS (36.4%) strategy as the normative CPS (32.2%) strategy ($Z = .551, p = .581$). It is of interest that there was a greater consistency between strategy self-reports and actual strategy use for subjects reporting an OPS (86%) than a CPS (53%) strategy ($Z = 3.29, p < .001$). This suggests that some subjects who may have intended to use the normative strategy were unable to execute it. Study 3 will examine this possibility.

Discussion

Together, the results of studies 1a, 1b, 1c, and study 2 suggest that consumers' reliance on the OPS is quite robust. Study 1 demonstrates this effect in the commonly occurring tasks of seeking evaluations, where it is normative, and for seeking recommendations, where it is not normative. The result is shown to be robust in both subjectively and objectively evaluated product categories and across information presentation formats. Study 2 shows that reliance on the OPS occurs in another common choice environment involving the use of agents. Despite having both an explicit presentation of the conditionalizing information and conflicting agents providing a potential cue to the differences

in the usefulness of prior information, subjects continued to rely on the OPS when making a product-purchase decision.

A number of potential reasons exist to account for subjects' difficulty selecting agents (studies 1a, 1b, and 1c) or assessing agent-provided ratings (study 2) when the appropriate measure of agent ability is associated with a conditional probability. One reason is that, as compared with objectively evaluated categories, people prefer an agent who is overall more similar to the target for subjectively evaluated alternatives (Feick and Higie 1992; Price et al. 1989). Contrary to differences observed between the scenarios in studies 1a and 1b, the results of study 2 are not consistent with this explanation. The subjective outcome of film preference in the video clerk scenario (29.3%) had more subjects choosing normatively than in the objective outcome of the stock analyst scenario (15%), although this difference did not reach significance ($Z = 1.87, p = .06$).

Three other potential reasons include lack of motivation, failure to ignore irrelevant information, and failure to select diagnostic information. First, lotteries in the prior studies may not have provided sufficient incentive to motivate subjects to perform well. Second, subjects may have generated an appropriate strategy for agent selection and were motivated to select the necessary information, but the presence of irrelevant information sidetracked otherwise good strategy use (Edgell et al. 1996; Hilton and Fein 1989; Hutchinson and Alba 1991). Finally, similar to findings in studies of hypotheses testing, subjects may have failed to generate an appropriate strategy for the task leading them to be unaware of the diagnostic information needed to assess agent ability (Doherty et al. 1979; Doherty et al. 1996). Study 3

TABLE 7
AGENT CHOICE PATTERNS FOR STUDY 2

	Conflicting video clerk task (<i>n</i> = 58)	Conflicting stock analyst task (<i>n</i> = 60)	Pooled tasks (<i>n</i> = 118)
Overall probability of success (%)	55.2	61.7	58.5
Conditional probability of success (%)	29.3 ^{a,*}	15.0 ^{a,*}	22.0 ^{a,*}
Positive hits/inverse conditional probability of success (%)	10.3*	8.3*	9.3*
Other (%)	5.2*	15.0*	10.2*

^aIndicates the normative strategy for the task.

*Indicates difference from Overall Probability of Success Strategy within task ($p < .05$).

was designed to identify which of these potential reasons may account for subjects' failure to recognize the conditional assessment of agent diagnosticity necessary in recommendation and conflicting agent tasks.

STUDY 3

To address the motivation issue, subjects in this study were provided with a more substantial incentive than in prior studies. In addition, subjects were only presented with the prior performance information that they requested. Observing subjects' information requests would provide insight into their ability to generate appropriate strategies. Further, those who did generate an appropriate strategy could apply it without being exposed to irrelevant information.

Method

Eighty-five student subjects participated individually in study 3 in exchange for extra credit in their marketing classes and a monetary reward for their performance. The study was a 2 (task) \times 2 (product scenario) repeated measures design. Unlike in study 1, task (evaluation and recommendation) and product scenario (video clerk and stock analyst) were manipulated within-subject. The video store clerk (study 1a) and stock analyst scenarios (study 1b) were included to increase generalizability of results and to reduce any potential carryover effect between the two tasks. Both the order of the task factor and the product scenario factor were counterbalanced across subjects.

For each scenario, subjects received a booklet containing a set of instructions explaining the scenario and the agent selection task. Subjects were told that they could receive up to \$4 for correctly selecting the better of two possible agents. Subjects were then presented with the summary tables for the two agents' prior performance. They were told that each agent had rated 100 alternatives; however, unlike in the previous studies, the tables were empty except for the labels "cell *a*," "cell *b*," "cell *c*," and "cell *d*" in the cells (see table 1). Subjects were told that they could request cell information from the experimenter but that each cell of information requested would reduce the possible \$4 earnings by 50 cents. Thus, subjects had both an incentive to select the better agent, an incentive to select only the information they

deemed necessary for the task, and a disincentive to select task-irrelevant information. Subjects were told that they could buy as many or as few cells as they liked, and in any order. The description of the agent selection scenario and the set of instructions were verbally repeated by the experimenter, and subjects were given the opportunity to ask questions to ensure they understood the task. Subjects then selected desired cell information, receiving both agents' frequency of instances for each cell requested. The experimenter recorded the order in which the cells of information were requested by each subject. Following agent selection, subjects were asked to verbally describe why they selected the information they did and to explain how they selected their preferred agent. Subjects' answers to these questions were tape-recorded for later analysis.

Results

The information selection results are presented in table 9. Allowing subjects to select as many or as few cells of information as they desired results in 16 possible single or multiple cell information selection combinations. The CPS can be assessed by selecting only cell *a* and cell *c*. Because subjects were provided with the total number of alternatives ($n = 100$) that had been rated, the OPS can be assessed by selection of either cells *a* and *d* or cells *b* and *c*. (Cell *b* and cell *c* would allow for assessment of the overall probability of failure, which is equivalent to $1 -$ the overall probability of success.) Thus, regardless of whether the subject utilizes a CPS or OPS strategy, just two cells of information are required. As a result, there is no monetary advantage found in the selection of the information necessary for assessing the OPS as compared with the CPS.

Subjects' selection of appropriate information for each task did not differ by scenario ($\chi^2(1) = .04, p = .850$), so results are presented and discussed collapsed across scenario. In addition, the order in which subjects completed the tasks was found to be unrelated to subjects' selection of information ($\chi^2(1) = 1.76, p = .184$).

In both the evaluation and recommendation conditions, selection of the information necessary for assessment of the OPS was more common than selection of any other information (evaluation condition: 25.9%; $Z = 2.41, p < .05$; recommendation condition: 33%; $Z = 2.42, p < .05$). Con-

TABLE 8
SUBJECTS' REPORTED EVALUATION PROCESS FOR STUDY 2

	Conflicting video clerk task (n = 58)	Conflicting stock analyst task (n = 60)	Pooled tasks (n = 118)
Overall probability of success (%)	36.2	36.7	36.4
Conditional probability of success (%)	31.0 ^a	33.3 ^a	32.2 ^a
Positive hits (%)	6.9*	1.7*	4.2*
Inverse conditional probability of success (%)	3.5*	6.7*	5.1*
Other (%) ^b	22.4	21.6	22.1

^aIndicates the normative strategy for the task.
^bNo single other reported strategy was greater than 10% in either study.
 *Indicates difference from Overall Probability of Success Strategy within task ($p < .05$).

sistent with the results of studies 1 and 2, subjects in the recommendation condition failed to recognize that assessment of agent diagnosticity should be restricted to the subset of prior performance data associated with positive predictions only. In fact, significantly fewer subjects in the recommendation condition selected only the necessary cell *a* and cell *c* information needed to assess the CPS (15.3%) than the OPS (33%; $Z = 2.33, p < .05$).

Subjects described 20 agent selection strategies (see table 10). Strategies that are mathematically or functionally equivalent were grouped together. For example, selecting the agent with the greater overall success (cell *a* plus cell *d*) is functionally equivalent to selecting the agent with the smaller overall failure (cell *b* plus cell *c*), because it will always lead to the selection of the same agent. After grouping, 11 different strategies were identified. Although many subjects (40% in the evaluation condition and 29% in the recommendation condition) selected more than two cells of information, no

subjects in either condition reported using strategies that utilized more than two cells. Thus, many subjects requested more information than they reported using in their agent choice. Once again, no effect was found for scenario ($\chi^2(1) = .90, p = .344$) nor order of task presentation ($\chi^2(1) = .324, p = .569$), and so results were collapsed.

In the evaluation condition, 30.6% of subjects reported utilizing a normatively appropriate OPS strategy, significantly greater than the next most commonly used strategies of PH (15.3%) and Least False Negatives (12.9%) (p 's $< .05$). In the recommendation condition, subjects' self-reports again exhibited a failure to appreciate the conditional nature of the task. The OPS strategy (38.8%) is also most common in the recommendation condition and was reported significantly more than the normatively appropriate CPS strategy (10.6%, $Z = 2.33, p < .05$). It is interesting that the next most commonly reported strategy in the recommendation condition was Least False Positives (20%).

TABLE 9
SUBJECTS' INFORMATION SELECTION FOR STUDY 3

Cells selected	Evaluation % (n = 85)	Recommendation % (n = 85)	Consistent with required information for strategy
None	1.2	.0	
<i>a</i> only	7.1	2.4	Positive hits
<i>b</i> only	2.4	2.4	
<i>c</i> only	3.5	10.6	
<i>d</i> only	.0	.0	
<i>a</i> and <i>b</i>	7.1	3.5	Inverse conditional probability
<i>a</i> and <i>c</i>	9.4	15.3 ^a	
<i>a</i> and <i>d</i>	21.2 ^a	22.4	Conditional probability
<i>b</i> and <i>c</i>	4.7 ^a	10.6	Overall probability
<i>b</i> and <i>d</i>	.0	.0	Overall probability ^b
<i>c</i> and <i>d</i>	3.5	3.5	
<i>a, b, and c</i>	10.6	8.2	
<i>a, b, and d</i>	5.9	7.1	
<i>a, c, and d</i>	9.4	8.2	
<i>b, c, and d</i>	3.5	.0	
<i>a, b, c, and d</i>	10.6	5.9	

NOTE.—Evaluation % and Recommendation % are the percentages of subjects who selected information from cells in the agent contingency tables.

^aIndicates the normative information selection for the task.
^bNote that *b* and *c* cells can be used to calculate the overall probability of success using $1 -$ the overall probability of failure.

TABLE 10
INFORMATION SOURCE SELECTION STRATEGIES FOR STUDY 3

Strategy	Subject described choosing the agent having:	Evaluation % (<i>n</i> = 85)	Recommendation % (<i>n</i> = 85)
Overall probability	Greater cell <i>a</i> + cell <i>d</i> Greater (cell <i>a</i> + cell <i>d</i>)/total Smaller cell <i>c</i> + cell <i>b</i> Smaller (cell <i>c</i> + cell <i>b</i>)/total	30.6 ^a	38.8
Conditional probability	Greater cell <i>a</i> /(cell <i>a</i> + cell <i>c</i>) Greater ratio of cell <i>a</i> to cell <i>c</i> Greater cell <i>a</i> and smaller cell <i>c</i> Smaller ratio of cell <i>c</i> to cell <i>a</i> Smaller of cell <i>d</i> /cell <i>a</i>	8.2*	10.6 ^{a,*}
Positive hits	Greater cell <i>a</i>	15.3*	5.9*
Inverse conditional probability	Greater ratio of cell <i>a</i> to cell <i>b</i> Greater cell <i>a</i> and smaller cell <i>b</i>	8.2*	4.7*
Negative inverse conditional probability	Greater cell <i>d</i> and smaller cell <i>c</i> Greater cell <i>d</i> /(cell <i>c</i> + cell <i>d</i>)	4.7*	4.7*
Positive predictions	Greater cell <i>a</i> + cell <i>c</i>	3.5*	1.2*
False negatives	Smaller cell <i>b</i>	3.5*	4.7*
False positives	Smaller cell <i>c</i>	12.9*	20.0*
Negative hits	Greater cell <i>d</i>	4.7*	2.4*
Balanced	More balanced cell <i>a</i> and cell <i>d</i>	1.2*	3.6*
Guessing or other	Don't know/just guessed	5.9*	3.5*

^aIndicates the normative strategy for the task.

*Indicates difference from Overall Probability of Success Strategy within task ($p < .05$).

Discussion

The results of study 3 suggest that subjects' overreliance on OPS strategies where CPS strategies are more appropriate is not due to either a lack of motivation or an inability to ignore task-irrelevant information. Providing a substantial monetary performance incentive still resulted in the majority of subjects in the recommendation condition failing to select only the information needed for the task and failing to utilize a CPS to select an agent. Removal of task-irrelevant information did not dissuade subjects from requesting it or from selecting an agent using an inappropriate strategy. It thus appears that subjects' shortcomings do not occur because they are not sufficiently motivated to use the best strategy or because otherwise good strategies become muddled when they are presented with irrelevant information. Instead, these shortcomings occur at a more basic level of strategy generation. Most subjects simply fail to generate appropriate strategies when the agent task requires consideration of conditional probabilities of agent success.

GENERAL DISCUSSION

Prior research has shown that people are subject to numerous types of difficulties when faced with tasks involving probability assessment (Bar-Hillel 1983; Brase et al. 1998; Einhorn and Hogarth 1986; Fischhoff and Beyth-Marom 1983; Gigerenzer and Hoffrage 1995; Pollatsek et al. 1987). The research presented here provides insight into probability assessment by consumers in a context involving the selection of agents to act as information sources. We examined three distinct tasks that are commonly encountered by consumers

at different stages in the decision process: choosing an agent to provide a recommendation, choosing an agent to provide an evaluation, and choosing an action when faced with conflicting agents' ratings.

A normative model was developed that specified how probability assessment from prior performance information is task dependent. Specifically, when ratings are not known at the time of choice, such as for evaluations, consideration of the OPS is required. Conversely, when the agent (or prospective agent) rating is known at the time of choice, such as for recommendations and choosing between conflicting agents, consideration of the CPS is required for optimal agent selection. A series of studies revealed that subjects failed to recognize the conditional nature of the tasks when the agent (or prospective agent) rating was known at the time of choice.

Three experiments in study 1 showed that subjects relied on the OPS strategy when it was normative for an evaluation task as well as when it was nonnormative for a recommendation task. In the recommendation task, when the agent's positive rating was known in advance, the majority of subjects failed to utilize the normative CPS strategy for both subjective (movies in studies 1a and 1c) and objective (stocks in study 1b) product categories as well as summary table (studies 1a and 1b) and serial (study 1c) information presentation formats.

A possible explanation for reliance on an OPS strategy is simply the implicit nature of the conditionalizing event in the recommendation task. Study 2 ruled out this explanation as the majority of subjects continued to rely on an OPS strategy even when the conditional nature of the task was made explicit by having subjects choose an action based on two conflicting agents' evaluations of a particular alter-

native. Study 2 demonstrated that failure to properly assess agent diagnosticity for a conditional task would lead to inferior agent selection but also, of greater importance, to inferior product choices.

Study 3 showed that, even when subjects were both enticed with substantial earnings and not constrained to select the agent in the presence of irrelevant information, the favored strategy for agent selection across tasks was still the OPS strategy and not the normative CPS strategy. Thus, the reason underlying subjects' poor performance appears to be a fundamental failure to recognize that when an agent's rating is known, the task requires a conditional probability assessment of agent diagnosticity. That is, most subjects failed to recognize that assessment of agent success should be restricted to the subset of prior performance data consistent with that agent's rating. Our results provide a number of insights to understanding of consumers' assessment of probability for information-source-selection tasks. First, a commonly replicated finding in the psychology literature is that, when asked to assess a conditional probability, $P(A|B)$, people will often provide its inverse, $P(B|A)$ (Bar-Hillel 1983; Dawes 1988; Eddy 1982; Gigerenzer and Murray 1987; Meehl and Rosen 1955; Nisbett and Ross 1980; Wyer 1977). Our results were not consistent with this finding. We found very little evidence of subjects using an ICPS strategy in any of the studies presented here that call for consideration of the conditional probability. Instead, our subjects favored an OPS strategy. This is consistent with studies that examine diagnosticity assessment in which subjects recognized the overall level of diagnosticity of a question but failed to recognize that conditionalizing information renders a subset of the data irrelevant for the task (Skov and Sherman 1986; Slowiaczek et al. 1992). Quite possibly, when assessing an agent, subjects relied on the OPS strategy because it is consistent with the way in which they naturally categorize an agent's successes and failures (Brase et al. 1998; Gavanski and Hui 1992; Rosch 1972). Study 1c provides some support for this. When the 2×2 summary tables were removed and subjects were required to structure the data themselves, there was a greater incidence of use of the OPS strategy. Thus, future research should examine how consumers form categories with respect to information-source successes and failures and the effects of these categories on selection strategy.

Second, recent studies have shown that many biases in probability assessment can be attenuated when the information is presented in terms of frequencies of past events rather than probabilities (Cosmides and Tooby 1996; Gigerenzer and Hoffrage 1995; Kleiter et al. 1997). While our intent was not to examine the effects of presentation format, our results indicate that a frequency presentation in these tasks still results in departure from normative strategies.

Finally, we note that subjects demonstrated sizable heterogeneity in the strategies used to select agents. This is particularly apparent in study 3, where subjects described using no less than 11 different strategies to select information and assess agent ability, suggesting that consumers

vary widely in the strategies they generate when faced with problems involving probabilities and, in particular, conditional probabilities. Still, although most of the strategies employed are not normative for the tasks presented, they do represent attempts by subjects to apply a rule to find the best agent for the task as opposed to guessing or some other random selection process. Further, in studies 1 and 2, we see that most of those who report using the OPS or CPS strategies demonstrated agent-selection patterns consistent with the reported strategy. This suggests that people generate and attempt to apply reasoned strategies to tasks involving probability assessment and that they may be quite capable of applying these strategies. Recent research reports that those subjects who tend to perform better in tasks of diagnostic data selection and covariation assessment also tend to score higher on tests of cognitive ability (Stanovich and West 1998). Potentially, consumers' agent-selection strategies may be similarly linked to cognitive ability. This should be addressed in future research. Beyond individual differences, it may also be possible for consumers to learn to correctly generate and apply appropriate decision rules for each agent-selection task (Nisbett 1993) or for task environments to be presented to fit the cognitive processes that people are most likely to use (Klayman and Brown 1993).

The fact that consumers regularly turn to agents for decision-making assistance makes our focus on the agent-selection process a relevant area of inquiry. Prior research has shown that consumers commonly turn to other individuals, systems, or organizations to act on their behalf to provide information, assistance, and opinions. Consumer reliance on agents is so pervasive that information provided by others makes up a large proportion of all prepurchase information and at times represents the only source of prepurchase information (Beatty and Smith 1987; Berning and Jacoby 1974; Kiel and Layton 1981; Olshavsky and Granbois 1979; Urbany et al. 1989). This reliance on agents may be a double-edged sword. Agents may be valuable as a means to simplify decision making but may be costly if consumers fail to select those agents who can best help them because they fail to consider the nature of the task involved.

We acknowledge that the modal OPS strategy is indeed an appropriate strategy for selecting an agent to provide an evaluation and, if used for recommendations or choosing between conflicting agents, will not always lead to inferior agent selection. Thus, if consumers do apply only a single strategy to all agent selection tasks, it may be that the OPS strategy is the most efficient, even if it does not always lead to selection of an optimal agent (McKenzie 1994). But the inability to choose a more appropriate strategy when necessary clearly does have negative consequences. The results presented here show that consumers' inability to recognize that information diagnosticity is task dependent may lead to choice of inferior agents and poor product decisions, even when the cost of doing so is high. This finding not only underscores the need for continued exploration of consumer agent selection but calls for further investigation of the

agent-selection process and methods for improving consumers' ability.

It is likely that in many agent-selection situations, in addition to prior performance of the agent, people will also have characteristics about the agent at their disposal (Burnkrant and Cousineau 1975; Feick and Higie 1992). How this information would be integrated should be considered. The striking similarity between the dichotomy of prior performance information versus characteristics about the agent and that of data-based versus theory-based processing in the covariation literature would suggest one starting point for future investigation (Bettman, Roedder John, and Scott 1986; Broniarczyk and Alba 1994). Another area for future investigation would be to consider how consumers make attributions about agents' abilities when they consider the characteristics of the agent, the agent's overall past performance, and the details of particular past instances of success and failure (Folkes 1988; Folkes, Koletsky, and Graham 1987; Weiner 2000).

In addition to consumer agent selection, our findings are relevant to other domains, including management choices involving marketing research information (Menon and Varadarajan 1992), accountant choices involving audit tools (Nelson, Libby, and Bonner 1995), and physician choices involving medical tests and symptoms (Burton 1994; Christensen-Szalanski and Bushyhead 1981). Like consumer agent selection, decisions in these domains involve probabilistic relationships between sources of information that can provide multiple ratings and multiple possible outcomes. If professionals in these areas rely on OPS strategies when they have conditionalizing information available in the form of market research results or the presence of particular symptoms, then they may make costly investment decisions or mistreat patients. Future research should examine whether professionals who regularly encounter these types of problems are also prone to rely on nonnormative strategies.

APPENDIX

COMPUTING THE CONDITIONAL PROBABILITY OF A SUCCESSFUL OUTCOME

Assessing the probability of a successful outcome based on prior performance information requires consideration of the task. Below, we consider how such assessment should work with various tasks.

A. AN EVALUATION TASK

In the case of an evaluation task, the odds of "success" given the relevant prior performance information (see table 1) can be computed using

$$\frac{P(\text{success}|\text{evaluation})}{P(\text{failure}|\text{evaluation})} = \frac{P(R_1)P(E_1|R_1) + P(R_2)P(E_2|R_2)}{P(R_1)P(E_2|R_1) + P(R_2)P(E_1|R_2)}, \quad (\text{A1})$$

where $P(R_i)$ is the probability that an agent will provide a particular rating, and $P(E_i|R_i)$ is the conditional probability that the consumer will experience a particular event or outcome, given the agent's rating.

When an information source is providing an evaluation, the source's rating is not known in advance. The probability of the agent's rating can be computed from the cell frequencies from table 1 as follows:

$$P(R_1) = \frac{a + c}{a + b + c + d},$$

$$P(R_2) = \frac{b + d}{a + b + c + d}. \quad (\text{A2})$$

According to Bayes Theorem, the conditional probabilities of success, given the agent's ratings, can be represented as

$$P(E_1|R_1) = \frac{P(R_1|E_1)P(E_1)}{P(R_1|E_1)P(E_1) + P(R_1|E_2)P(E_2)},$$

$$P(E_2|R_2) = \frac{P(R_2|E_2)P(E_2)}{P(R_2|E_2)P(E_2) + P(R_2|E_1)P(E_1)}. \quad (\text{A3})$$

Similarly, the conditional probabilities of failure given the agent's ratings can be represented as

$$\begin{aligned}
 P(E_2|R_1) &= \frac{P(R_1|E_2)P(E_2)}{P(R_1|E_1)P(E_1) + P(R_1|E_2)P(E_2)}, \\
 P(E_1|R_2) &= \frac{P(R_2|E_1)P(E_1)}{P(R_2|E_2)P(E_2) + P(R_2|E_1)P(E_1)}.
 \end{aligned}
 \tag{A4}$$

Assuming that prior performance information is a natural sampling (Kleiter 1994), the conditional probabilities of success and failure, given the agent’s ratings, can be computed from the cell frequencies as follows:

$$P(E_1|R_1) = \frac{\left(\frac{a}{a+b}\right)\left(\frac{a+b}{a+b+c+d}\right)}{\left(\frac{a}{a+b}\right)\left(\frac{a+b}{a+b+c+d}\right) + \left(\frac{c}{c+d}\right)\left(\frac{c+d}{a+b+c+d}\right)} = \frac{a}{a+c}.
 \tag{A5}$$

$$P(E_2|R_2) = \frac{\left(\frac{d}{c+d}\right)\left(\frac{c+d}{a+b+c+d}\right)}{\left(\frac{d}{c+d}\right)\left(\frac{c+d}{a+b+c+d}\right) + \left(\frac{b}{a+b}\right)\left(\frac{a+b}{a+b+c+d}\right)} = \frac{d}{b+d}.
 \tag{A6}$$

$$P(E_2|R_1) = \frac{\left(\frac{c}{c+d}\right)\left(\frac{c+d}{a+b+c+d}\right)}{\left(\frac{a}{a+b}\right)\left(\frac{a+b}{a+b+c+d}\right) + \left(\frac{c}{c+d}\right)\left(\frac{c+d}{a+b+c+d}\right)} = \frac{c}{a+c}.
 \tag{A7}$$

$$P(E_1|R_2) = \frac{\left(\frac{b}{a+c}\right)\left(\frac{a+c}{a+b+c+d}\right)}{\left(\frac{d}{c+d}\right)\left(\frac{c+d}{a+b+c+d}\right) + \left(\frac{b}{a+b}\right)\left(\frac{a+b}{a+b+c+d}\right)} = \frac{b}{b+d}.
 \tag{A8}$$

Therefore, the odds of success are

$$\frac{P(\text{success}|\text{evaluation})}{P(\text{failure}|\text{evaluation})} = \frac{\left(\frac{a+c}{a+b+c+d}\right)\left(\frac{a}{a+c}\right) + \left(\frac{b+d}{a+b+c+d}\right)\left(\frac{d}{b+d}\right)}{\left(\frac{a+c}{a+b+c+d}\right)\left(\frac{c}{a+c}\right) + \left(\frac{b+d}{a+b+c+d}\right)\left(\frac{b}{b+d}\right)} = \frac{a+d}{c+b}.
 \tag{A9}$$

B. A RECOMMENDATION TASK

In the case of a recommendation task, the information source’s rating is known in advance to be positive. Therefore, the $P(R_1) = 1$ and $P(R_2) = 0$, which simplifies the computation of the odds of success.

$$\frac{P(\text{success}|\text{recommendation})}{P(\text{failure}|\text{recommendation})} = \frac{P(E_1|R_1)}{P(E_2|R_1)} = \frac{\left(\frac{a}{a+c}\right)}{\left(\frac{c}{a+c}\right)} = \frac{a}{c}.
 \tag{A10}$$

C. CONFLICTING OPINIONS

When making a decision in the face of conflicting ratings from two agents, the consumer is faced with the task of deciding whether to purchase the product rather than selecting an agent. In this case, a “success” is defined as maximizing the probability of a positive event outcome; therefore,

$$\frac{P(\text{success}|\text{conflicting ratings})}{P(\text{failure}|\text{conflicting ratings})} = \frac{P(E_1|R_{11} \cap R_{22})}{P(E_2|R_{11} \cap R_{22})} = \frac{P(R_{11} \cap R_{22}|E_1)P(E_1)}{P(R_{11} \cap R_{22}|E_2)P(E_2)}, \quad (\text{A11})$$

where R_1 is the rating provided by the first agent and R_2 is the rating provided by the second agent. This can be computed using cell frequencies as

$$\frac{\left(\frac{a_1}{a_1+b_1}\right)\left(\frac{b_2}{a_2+b_2}\right)\left(\frac{a_1+b_1+a_2+b_2}{a_1+b_1+c_1+d_1+a_2+b_2+c_2+d_2}\right)}{\left(\frac{c_1}{c_1+d_1}\right)\left(\frac{d_2}{c_2+d_2}\right)\left(\frac{c_1+d_1+c_2+d_2}{a_1+b_1+c_1+d_1+a_2+b_2+c_2+d_2}\right)}, \quad (\text{A12})$$

where subscripts refer to cell frequencies for the corresponding agent. If we assume that prior performance information is based on a natural sampling (Kleiter 1994), then the equation reduces to

$$\frac{P(\text{success}|\text{conflicting ratings})}{P(\text{failure}|\text{conflicting ratings})} = \left(\frac{a_1 \times b_2}{a_2 + b_2}\right) \times \left(\frac{c_2 + d_2}{c_1 \times d_2}\right) = \frac{a_1}{c_1} \times \frac{b_2}{d_2} \times \frac{c_2 + d_2}{a_2 + b_2}. \quad (\text{A13})$$

Posterior odds greater than 1 indicate higher probability of E_1 as compared with E_2 . Posterior odds less than 1 indicate lower probability of E_1 as compared with E_2 . Therefore, in order to make a decision in the face of conflicting opinions, the consumer must consider the appropriate conditional probabilities associated with each agent rating and the probability associated with a positive and negative event.

When the prior odds are 1/1 ($P(E_1) = P(E_2) = .5$), then only the accuracy of each agent needs to be considered. In this case, the posterior odds reduce to

$$\frac{P(\text{success}|\text{conflicting ratings})}{P(\text{failure}|\text{conflicting ratings})} = \frac{a_1}{c_1} \times \frac{b_2}{d_2}. \quad (\text{A14})$$

Therefore, when prior odds are 1/1,

$$\frac{P(\text{success}|\text{conflicting ratings})}{P(\text{failure}|\text{conflicting ratings})} > 1, \text{ only if } \frac{a_1}{c_1} > \frac{d_2}{b_2} \text{ and } \frac{a_1}{a_1 + c_1} > \frac{d_2}{b_2 + d_2}; \quad (\text{A15})$$

$$\frac{P(\text{success}|\text{conflicting ratings})}{P(\text{failure}|\text{conflicting ratings})} < 1, \text{ only if } \frac{a_1}{c_1} < \frac{d_2}{b_2} \text{ and } \frac{a_1}{a_1 + c_1} < \frac{d_2}{b_2 + d_2}; \quad (\text{A16})$$

$$\frac{P(\text{success}|\text{conflicting ratings})}{P(\text{failure}|\text{conflicting ratings})} = 1, \text{ only if } \frac{a_1}{c_1} = \frac{d_2}{b_2} \text{ and } \frac{a_1}{a_1 + c_1} = \frac{d_2}{b_2 + d_2}. \quad (\text{A17})$$

So when prior odds are 1/1, comparing the appropriate conditional probability of success for each agent is equivalent to assessing the posterior odds of success.

[Received June 1999. Revised April 2001. David Glen Mick served as editor, and Joel Huber served as associate editor for this article.]

REFERENCES

- Arkes, Hal R. and Allan R. Harkness (1983), "Estimates of Contingency between Two Dichotomous Variables," *Journal of Experimental Psychology: General*, 112 (March), 117-135.
- Bar-Hillel, Maya (1983), "The Base Rate Fallacy Controversy," in *Decision Making under Uncertainty*, ed. Roland W. Scholz, Amsterdam: North-Holland, 39-61.
- Beatty, Sharon E. and Scott M. Smith (1987), "External Search Effort: An Investigation across Several Product Categories," *Journal of Consumer Research*, 14 (June), 83-95.
- Berning, Carol A. Kohn and Jacob Jacoby (1974), "Patterns of Information Acquisition in New Product Purchases," *Journal of Consumer Research*, 1 (September), 18-22.
- Bettman, James R., Deborah Roedder John, and Carol A. Scott (1986), "Covariation Assessment by Consumers," *Journal of Consumer Research*, 13 (December), 316-326.

- Birnbaum, Michael H. and Barbara A. Mellers (1983), "Bayesian Inference: Combining Base Rates with Opinions of Sources Who Vary in Credibility," *Journal of Personality and Social Psychology*, 45 (October), 792-804.
- Brase, Gary L., Leda Cosmides, and John Tooby (1998), "Individuation, Counting, and Statistical Inference: The Role of Frequency and Whole-Object Representations in Judgment under Uncertainty," *Journal of Experimental Psychology: General*, 127 (March), 3-21.
- Broniarczyk, Susan M. and Joseph W. Alba (1994), "Theory versus Data in Prediction and Correlation Tasks," *Organizational Behavior and Human Decision Processes*, 57 (January), 117-139.
- Brown, Jacqueline Johnson and Peter H. Reingen (1987), "Social Ties and Word-of-Mouth Referral Behavior," *Journal of Consumer Research*, 14 (December), 350-362.
- Burnkrant, Robert E. and Alain Cousineau (1975), "Informational and Normative Social Influence in Buyer Behavior," *Journal of Consumer Research*, 2 (December), 206-215.
- Burton, Paul R. (1994), "Helping Doctors to Draw Appropriate Inferences from the Analysis of Medical Studies," *Statistics in Medicine*, 13 (September), 1699-1713.
- Castellan, N. John (1973), "Multiple-Cue Probability Learning with Irrelevant Cues," *Organizational Behavior and Human Decision Processes*, 9 (February), 16-29.
- Christensen-Szalanski, Jay J. J. and James B. Bushyhead (1981), "Physicians' Use of Probabilistic Information in a Real Clinical Setting," *Journal of Experimental Psychology: Human Perception and Performance*, 7 (August), 928-935.
- Cosmides, Leda and John Tooby (1996), "Are Humans Good Intuitive Statisticians After All? Rethinking Some Conclusions from the Literature on Judgement under Uncertainty," *Cognition*, 58 (January), 1-73.
- Dawes, Robin M. (1988), *Rational Choice in an Uncertain World*, San Diego, CA: Harcourt Brace Jovanovich.
- Doherty, Michael E., Randall Chadwick, Hugh Garavan, David Barr, and Clifford R. Mynatt (1996), "On People's Understanding of the Diagnostic Implications of Probabilistic Data," *Memory and Cognition*, 24 (September), 644-654.
- Doherty, Michael E., Clifford R. Mynatt, Ryan D. Tweney, and Michael D. Schiavo (1979), "Pseudodiagnosticity," *Acta Psychologica*, 43 (March), 111-121.
- Eddy, David M. (1982), "Probabilistic Reasoning in Clinical Medicine: Problems and Opportunities," in *Judgement under Uncertainty: Heuristics and Biases*, ed. Danny Kahneman, Paul Slovic, and Amos Tversky, New York: Cambridge University Press, 249-267.
- Edgell, Stephen E., N. John Castellan, Jr., Robert M. Roe, Jon M. Barnes, Pak C. Ng, Randy D. Bright, and Laura A. Ford (1996), "Irrelevant Information and Probabilistic Categorization," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22 (November), 1463-1481.
- Einhorn, Hillel J. and Robin M. Hogarth (1986), "Judging Probable Cause," *Psychological Bulletin*, 99 (January), 3-19.
- Feick, Lawrence and Robin A. Higie (1992), "The Effects of Preference Heterogeneity and Source Characteristics on Ad Processing and Judgments about Endorsers," *Journal of Advertising*, 21 (Summer), 9-24.
- Feick, Lawrence and Linda L. Price (1987), "The Market Maven: A Diffuser of Marketplace Information," *Journal of Marketing*, 51 (January), 83-97.
- Fischhoff, Baruch and Ruth Beyth-Marom (1983), "Hypothesis Evaluation from a Bayesian Perspective," *Psychological Review*, 90 (July), 239-260.
- Folkes, Valarie S. (1988), "Recent Attribution Research in Consumer Behavior: A Review and New Directions," *Journal of Consumer Research*, 14 (March), 548-565.
- Folkes, Valeries S., Susan Koletsky, and John L. Graham (1987), "A Field Study of Causal Inferences and Consumer Reaction: The View from the Airport," *Journal of Consumer Research*, 13 (March), 534-539.
- Gavanski, Igor and Chun Hui (1992), "Natural Sample Spaces and Uncertain Belief," *Journal of Personality and Social Psychology*, 63 (November), 766-780.
- Gigerenzer, Gerd and Ulrich Hoffrage (1995), "How to Improve Bayesian Reasoning without Instruction: Frequency Formats," *Psychological Review*, 102 (October), 684-704.
- Gigerenzer, Gerd and David J. Murray (1987), *Cognition as Intuitive Statistics*, Hillsdale, NJ: Erlbaum.
- Glazer, Rashi, Joel H. Steckel, and Russell S. Winer (1992), "Locally Rational Decision Making: The Distracting Effect of Information on Managerial Performance," *Management Science*, 38 (February), 212-226.
- Harkness, Allan R., Kenneth G. DeBono, and Eugene Borgida (1985), "Personal Involvement and Strategies for Making Contingency Judgments: A Stake in the Dating Game Makes a Difference," *Journal of Personality and Social Psychology*, 49 (July), 22-32.
- Hilton, James L. and Steven Fein (1989), "The Role of Typical Diagnosticity in Stereotype-Based Judgments," *Journal of Personality and Social Psychology*, 57 (August), 201-211.
- Hutchinson, J. Wesley and Joseph Alba (1991), "Ignoring Irrelevant Information: Situational Determinants of Consumer Learning," *Journal of Consumer Research*, 18 (December), 325-345.
- Katz, Elihu and Paul F. Lazarsfeld (1955), *Personal Influence*, New York: Free Press.
- Kiel, Geoffrey C. and Roger A. Layton (1981), "Dimensions of Consumer Information Seeking Behavior," *Journal of Marketing Research*, 18 (May), 233-239.
- Klayman, Joshua and Kaye Brown (1993), "Debias the Environment Instead of the Judge: An Alternative Approach to Reducing Error in Diagnostic (and Other) Judgment," *Cognition*, 49 (October-November), 97-122.
- Kleiter, Gernot D. (1994), "Natural Sampling: Rationality without Base Rates," in *Contributions to Mathematical Psychology, Psychometrics, and Methodology*, ed. Gerhard H. Fischer and Donald Laming, New York: Springer, 375-388.
- Kleiter, Gernot D., Marianne Krebs, Michael E. Doherty, Hugh Garavan, Randall Chadwick, and Gregory Brake (1997), "Do Subjects Understand Base Rates?" *Organizational Behavior and Human Decision Processes*, 72 (October), 25-61.
- Macchi, Laura (1995), "Pragmatic Aspects of the Base Rate Fallacy," *Quarterly Journal of Experimental Psychology*, 48A (February), 188-207.
- Mann, D. H. (1972), "A Matrix Technique for Finite Bayesian Decision Problems," *Decision Sciences*, 3 (October), 129-136.
- McFall, Richard M. and Teresa A. Treat (1999), "Quantifying the Information Value of Clinical Assessments with Signal Detection Theory," *Annual Review of Psychology*, 50, 215-241.
- McKenzie, Craig R. M. (1994), "The Accuracy of Intuitive Judgment Strategies: Covariation Assessment and Bayesian Inference," *Cognitive Psychology*, 26 (June), 209-239.
- Meehl, Paul E. and Albert Rosen (1955), "Antecedent Probability

- and the Efficiency of Psychometric Signs, Patterns, and Cutting Scores," *Psychological Bulletin*, 52 (March), 194–216.
- Menon, Anil and P. Rajan Varadarajan (1992), "A Model of Marketing Knowledge Use within Firms," *Journal of Marketing*, 56 (October), 53–71.
- Nelson, Mark W., Robert Libby, and Sarah E. Bonner (1995), "Knowledge Structure and the Estimation of Conditional Probabilities in Audit Planning," *Accounting Review*, 70 (January), 27–47.
- Nisbett, Richard E. (1993), *Rules for Reasoning*, Hillsdale, NJ: Erlbaum.
- Nisbett, Richard E. and Lee Ross (1980), *Human Inference: Strategies and Short Comings of Social Judgment*, Englewood Cliffs, NJ: Prentice-Hall.
- Olshavsky, Richard W. and Donald H. Granbois (1979), "Consumer Decision Making—Fact or Fiction?" *Journal of Consumer Research*, 6 (September) 93–100.
- Pollatsek, Alexander, Arnold D. Well, Clifford Konold, Pamela Hardiman, and George Cobb (1987), "Understanding Conditional Probabilities," *Organizational Behavior and Human Decision Processes*, 40 (October), 255–269.
- Price, Linda L., Lawrence F. Feick, and Robin A. Higie (1989), "Preference Heterogeneity and Coorientation as Determinants of Perceived Informational Influence," *Journal of Business Research*, 19 (November), 227–242.
- Reingen, Peter H. and Jerome B. Kernan (1986), "Analysis of Referral Networks in Marketing: Methods of Illustration," *Journal of Marketing Research*, 23 (November), 370–378.
- Rogers, Everett M. (1983), *Diffusion of Innovations*, Glencoe, IL: Free Press.
- Rosch, Eleanor (1972), "Classification of Real-World Objects: Origins and Representations in Cognition," in *Thinking: Readings in Cognitive Science*, ed. Philip N. Johnson-Laird and Peter C. Wason, Cambridge: Cambridge University Press, 212–222.
- Rosch, Eleanor and Carolyn B. Mervis (1975), "Family Resemblances: Studies in the Internal Structure of Categories," *Cognitive Psychology*, 7 (October), 573–605.
- Rosen, Dennis L. and Richard W. Olshavsky (1987), "A Protocol Analysis of Brand Choice Strategies Involving Recommendations," *Journal of Consumer Research*, 14 (December), 440–444.
- Shaklee, Harriet and Michael Mims (1982), "Sources of Error in Judging Event Covariations: Effects of Memory Demands," *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8 (May), 208–224.
- Shaklee, Harriet and Diane Tucker (1980), "A Rule Analysis of Judgments of Covariation between Events," *Memory and Cognition*, 8 (September), 459–467.
- Sherman, Steven J., Matthew N. McMullen, and Igor Gavanski (1992), "Natural Sample Spaces and the Inversion of Conditional Judgments," *Journal of Experimental Social Psychology*, 28 (September), 401–421.
- Skov, Richard B. and Steven J. Sherman (1986), "Information-Gathering Processes: Diagnosticity, Hypothesis-Confirmatory Strategies, and Perceived Hypothesis Confirmation," *Journal of Experimental Psychology*, 22 (March), 93–121.
- Slowiaczek, Louisa M., Joshua Klayman, Steven J. Sherman, and Richard B. Skov (1992), "Information Selection and Use in Hypothesis Testing: What Is a Good Question, and What Is a Good Answer?" *Memory and Cognition*, 20 (July), 392–405.
- Solomon, Michael R. (1986), "The Missing Link: Surrogate Consumers in the Marketing Chain," *Journal of Marketing*, 50 (October), 208–218.
- (1987), "The Wardrobe Consultant: Exploring the Role of a New Retailing Partner," *Journal of Retailing*, 63 (Summer), 110–128.
- (1999), *Consumer Behavior: Buying, Having, and Being*, 4th ed., Upper Saddle River, NJ: Prentice-Hall.
- Stanovich, Keith E. and Richard F. West (1998), "Who Uses Base Rates and $P(D|H)$? An Analysis of Individual Differences," *Memory and Cognition*, 26 (January) 161–179.
- Trolier, Tina K. and David L. Hamilton (1986), "Variables Influencing Judgments of Correlational Relations," *Journal of Personality and Social Psychology*, 50 (May) 879–888.
- Urbany, Joel E., Peter R. Dickson, and William L. Wilkie (1989), "Buyer Uncertainty and Information Search," *Journal of Consumer Research*, 16 (September), 208–215.
- Weiner, Bernard (2000), "Attributional Thoughts about Consumer Behavior," *Journal of Consumer Research*, 27 (December), 382–387.
- West, Patricia M. (1996), "Predicting Preferences: An Examination of Agent Learning," *Journal of Consumer Research*, 23 (June), 68–80.
- West, Patricia M. and Susan M. Broniarczyk (1998), "Integrating Multiple Opinions: The Role of Aspiration Level on Consumer Response to Critic Consensus," *Journal of Consumer Research*, 25 (June), 38–51.
- Woodside, Arch G. and J. William Davenport, Jr. (1974), "The Effect of Salesman Similarity and Expertise on Consumer Purchasing Behavior," *Journal of Marketing Research*, 11 (May), 198–202.
- Wyer, Robert S. (1977), "The Role of Logical and Nonlogical Factors in Making Inferences about Category Membership," *Journal of Experimental Social Psychology*, 13 (November), 577–595.