

When Words Speak Louder Than Actions: Another's Evaluations Can Appear More Diagnostic Than Their Decisions

RICHARD GONZALEZ

University of Washington

When making decisions we often infer the value of the alternatives from the behavior of others. In some domains, however, we fail to make such inferences. Stock trading is a case in point—investors appear to undervalue the significance of the behavior of those who are willing to trade. A simulated market was developed to test this claim. Subjects made a series of portfolio reallocations on the basis of reports they read. Subjects in the experimental conditions also saw the responses of another person who based their responses either on the same information as the subject (high redundant information) or on different but equally predictive information (low redundant information). In fact, the responses subjects observed were identical in both conditions. A normative argument suggests that another's decision is more diagnostic when it is based on information having low redundancy rather than high redundancy. The dependent measure was an index of the degree of conformity to the other's decision. Results show that people become more sensitive to the diagnosticity of another's responses when observing how the person evaluated their information rather than how the person acted on their information. A salience hypothesis is involved to account for the pattern of results. Additional studies were conducted to test various alternative explanations. Implications for the measurement of conformity and influence are discussed. © 1994 Academic Press, Inc.

The behavior of others can be informative. In the context of decision making we often can infer the value of the alternatives from the behavior of others. For example, we may infer that a new compact disk player is of high quality if several of our music-loving friends recently bought it. Another example comes from Groucho Marx in his famous quip, "I wouldn't want to belong to any club that would accept me as a member."

This research was conducted as part of Richard Gonzalez' dissertation at Stanford University. I thank Herb Clark, Mark Lepper, Jim March, Lee Ross, and Amos Tversky for their helpful comments and suggestions. Ugo Colella and Josh Hon served as experimenters throughout the project. Address reprint requests to Richard Gonzalez, Department of Psychology, University of Washington, Guthrie Hall, NI-25, Seattle, WA 98195.

Groucho implies that he interprets a club's decision to grant him membership as information about the quality of the club. Such inferences regarding the quality of the alternatives may also occur when the behavior of others goes against expectations. We may suspect that a restaurant is less than adequate if we notice that the dining room is empty on what should be a busy Friday night. In short, it appears that people can infer the quality, or attractiveness, of the alternatives from the decisions of others.

Social psychologists have examined this behavior-of-others-as-information hypothesis within a more general context. One of the original concerns of attribution theorists (e.g., Heider, 1958; Kelley, 1967) was how the behavior of others can serve as a basis for making attributions about the environment. Kelley (1967) called this entity attributions.¹ For example, consider the judgment of whether a movie is humorous. "Am I to take my enjoyment of a movie as a basis for an attribution to the movie (that it is intrinsically enjoyable) or for an attribution to myself . . . ?" (Kelley, 1967, p. 194). Whether a movie is judged humorous partly depends on the behavior of others. Am I the only one in the movie theater laughing or are others in the audience laughing too? See Kassin (1979) for a review of the role of consensus in attribution research.

Economists have also been concerned about the information implied by the actions of others. This has been studied in, among other things, the used car market, the calculation of insurance premiums, and labor strikes (e.g., Akerlof, 1970; Grossman, 1981; Hayes, 1984; Spence, 1974; Varian, 1984). In the used car market a person's decision to sell a car may be interpreted as a "signal" about the quality of the car (Akerlof, 1970; Spence, 1974; Varian, 1984). That is, the buyer may infer that there is something wrong with the car (e.g., mechanical trouble) from the fact that the owner is selling. In actual transactions the buyer guards against the purchase of a "lemon" by investigating into the reasons the seller is selling. The buyer may decide to purchase the car only when the seller provides "satisfactory" reasons (e.g., liquidating an estate, moving to another country, etc.). If the seller is a used car dealer, the buyer might inquire if there is any information on the previous owner. The buyer may become suspicious if the dealer tells her that the sporty red car was previously owned by an elderly widow who only drove 3 miles a week to church.

Another example where the behavior of others serves as a "signal" is

¹ Some attribution theorists have followed a different path, also set by Heider (1958), where the object of inference is not the environment but the actor (e.g., Jones & Davis, 1965; Jones, Davis, & Gergen, 1961).

seen in a practice insurance companies use to set premiums. Insurance companies sometimes use the frequency of policy sales in a particular region as a cue to potential risks. Varian (1984) points out that

the purchasers of insurance may well have a better idea of the relevant risks than does the insurance company. For example, suppose an insurance company wants to offer a policy that insures bicycles against theft. They may examine the theft statistics for bicycles and set a rate that would allow them to break even on the policy. But consider what happens when they offer the policy to the public. The people who live in areas with very high theft rates will buy the insurance while those who live in less dangerous areas will not. Thus, the insurance company will have to revise its rates upward in order to break even.

Thus, an insurance company can interpret a “run” on their policies as an indication that clients know something about the risks that the company does not.

In some domains, however, people do not seem to make such inferences. Stock trading is a case in point. Unlike the used car situation, investors appear to undervalue the significance of the behavior of those who are willing to trade. Consider an investor who has formed a positive evaluation of a stock on the basis of some promising news. If the news is public and the price of the stock does not increase substantially, then the investor’s enthusiasm should be tempered—others may know something about the company the investor does not (Arrow, 1986; Plott & Sunder, 1988; Sharpe, 1985). Economic theory makes a strong prediction. Because the offer to buy or sell a stock is itself informative, the offer should produce a transfer of information that *inhibits* the transaction. Thus, economic theory is consistent with the intuition captured in Groucho’s quip. However, in contrast to what would be expected under this “signaling” model, trading in actual stock exchanges is exceptionally high (e.g., approximately 150 million shares exchange hands daily in the New York Stock Exchange alone). Even when taking into account inside information, excessive trading is difficult to explain by standard economic theory, especially if one considers the nontrivial transaction costs such as brokerage commissions and bid–ask spreads (Arrow, 1986). We propose that insensitivity to the diagnosticity of other people’s behavior may play a role in excessive trading.

Why do investors appear relatively insensitive to the information that can be inferred from the behavior of others? One simple explanation for the failure to make inferences from another’s transactions may be that investors do not believe that others possess different information. We call this the *redundancy hypothesis*. Investors may doubt that others have independent, nonredundant information about the value of the stock. An investor may formulate a positive impression of a company from a glow-

ing *Wall Street Journal* article and may not consider that those willing to sell their holdings may be acting under different information. Stronger versions of the redundancy hypothesis include beliefs that others possess information of less quality and others do not know how to "read" the market as well. The redundancy hypothesis is independent of the phenomenon that investors may undervalue the fact that most information is public, e.g., millions of people read the same newspaper article evaluating a particular company.²

Clearly, in the case of asymmetric information the redundancy hypothesis does not apply. If an investor does not have any information about a stock, then the behavior of others will, in general, appear informative. If I know nothing about a stock and I discover that other people are buying it, then I may infer that the stock is a favorable investment. Thus, under conditions of asymmetric information people do infer information from the behavior of others (Jaffe, 1974; Plott & Sunder, 1988) and, as shown in a simulated market by Plott and Sunder, people can use this information to achieve the same profit levels as those who have direct access to the information.

There is a normative argument that, all other things being equal, another's decision is more diagnostic when it is based on unique information rather than shared information. To illustrate the normative difference between unique and shared information conditions consider the following two investors who differ in their information sources: one investor has the same information you have (high redundancy), the second investor has different but equally predictive information (low redundancy). Normative considerations suggest that, to the extent that you are influenced by other investors, you should be less influenced by the investor with shared information than the investor with unique information because the former does not contribute "new" information (i.e., you learn less about the object when observing the decisions of someone with shared information).

An analogy with multiple regression may make the normative argument clearer. Suppose you predict GPA with one variable using linear regression. If you add a second predictor that is not correlated with the first, then predictive power (say, in terms of R^2) will increase more than if the second predictor is correlated with the first predictor (assuming that all predictors are equally correlated with the criterion). Several models of combining dependent sources of information share a similar intuition that redundancy lowers diagnosticity (e.g., Clemen & Winkler, 1985; French,

² I do not mean to imply that "signaling" does not occur in all aspects of the stock market. There is empirical work suggesting that analysts interpret "stock splitting" as evidence that the firm has a positive forecast (e.g., Brennan & Hughes, 1990; Sharpe, 1985).

1980; Winkler, 1981; Zeckhauser, 1971). Philosophers of science also argue that independent information is more diagnostic. For example, more weight should be attached to a conceptual replication of an experiment than to an identical replication—the former adds more credibility and generality to the empirical finding (e.g., Horwich, 1982).

The present paper reports five studies examining when people attend to the diagnosticity of another's behavior. The redundancy hypothesis, people do not believe that others have different information, is tested by presenting subjects the decisions of another person (i.e., a "previous participant") and manipulating the overlap between the subject's information and the previous participant's information. The dependent measure is an index of how much the subjects' decisions were influenced by the previous participant. The redundancy hypothesis suggests that people will be insensitive to the redundancy of the information, whereas normative considerations would have the subject differentially weight the response of the two previous participants.

STUDY 1: OBSERVING DECISIONS

The first study investigates the redundancy hypothesis. The task involved a series of portfolio decisions on a computerized market game. Subjects read stock reports and, on the basis of these reports, made decisions about the percentage of a portfolio that should be allocated to each stock. They also saw the decisions of a previous participant. The previous participant was framed as someone with average performance—doing very well on some trials and not so well on other trials. Redundancy was operationalized in terms of the overlap between the subject's information and the previous participant's information. There were two experimental conditions. In one, the previous participant purportedly read the same report (high redundancy with the subject's information); in a second, the previous participant purportedly read a different but equally predictive report (low redundancy with the subject's information). A third condition, a control group where subjects did not observe the decisions of a previous participant, was also included.

Method

Subjects. One hundred and seven subjects participated in this experiment (36 in the low redundancy condition, 35 in the high-redundancy condition, and 36 control subjects). Each subject received \$5.00 for participating in the study. Subjects were recruited from undergraduate dormitories at Stanford University to participate in a study on "stock market decisions."

Materials. Actual stock reports were used to add realism to the study. The reports were compiled from *Value Line*, a weekly information service

that provides independent evaluations on over 2000 common stocks. *Value Line* lists financial information (e.g., trends of stock prices, betas, etc.) as well as a four- to five-paragraph report of the company's status (approximately 300 words).

Three to five sentences were abstracted from 90 different *Value Line* reports published between 1972 and 1980. The reports emphasized marketing and managerial information such as demand and supply, management-labor relations, and competition. The companies were sampled from various industries including retail, transportation, energy, and food service.

The 90 stock reports were arranged into 30 portfolios—each consisting of three stocks. The portfolios were constructed so that the three stocks had a particular pattern. Namely, in each portfolio one stock had a clear signal either to buy or to sell. This stock is referred to as the target stock. The remaining two reports in the portfolio were mixed, or neutral (i.e., having a "hold" signal). The *Value Line* reports for the three stocks comprising a given portfolio were taken from the same calendar year.

To assess the direction of the *Value Line* reports (i.e., signaling a buy, sell, or hold response) five independent raters judged the favorableness of the 90 reports. The raters were given the following instructions.

We would like you to imagine that you are a consultant for a trust fund which holds some stock. You have been hired to evaluate the current portfolio of the trust fund and suggest a course of action that will most likely increase the total asset position of the portfolio. The time frame you are given is one year. That is, your recommendation will be acted upon immediately and then evaluated at the end of the year. The holders of the trust fund expect your decisions to maximize their position at the end of the year. In this sense, your rating can be viewed as a prediction (based on the written evaluations) of how favorable the stock will be in one year.

Raters used a five-point scale ranging from definitely buy to definitely sell.

Twelve portfolios where the majority of raters agreed on both the direction and extremity of the evaluation were selected. Across the five judges and the 12 selected portfolios, 82% of the ratings for reports *a priori* classified as having a clear buy or sell signal were at the appropriate endpoints of the favorableness scale (highly favorable or highly unfavorable), and 58% of the ratings for the neutral stocks were at the midpoint of the scale ("mixed outlook"). Moreover, 37% of the ratings for the neutral stocks fell in the "leaning toward" category, though no clear patterns emerged since raters were just as likely to rate "leaning toward sell" as "leaning toward buy." Six of the 12 portfolios included a target stock with a buy signal; the remaining 6 portfolios included a target stock with a sell signal.

An example of a report having a sell signal is:

Pay 'n Save Corporation
Operates Drug Store Chain

Pay 'n Save's 23 drugstores located in northern California lost money, partially reflecting intense price competition from supermarkets. The drugstore chain, including the problem northern California group, had a very poor Christmas selling season. Store opening expenses will increase sharply.

An example of a report having a buy signal is:

Almaden Corporation
California Wine Producer

The California grape harvest is again large. Three back-to-back bumper yields resulted in a sharp decline of prices for grape and bulk wines, a favorable development for Almaden. With the larger yield, Almaden will be able to expand production of the higher-priced and more profitable varieties.

Procedure. Subjects participated in 12 rounds; each round involved a different portfolio consisting of three stocks and a T-Bill account (i.e., a savings account). The task required subjects to play the role of a securities manager, examine the reports, and suggest a reallocation of the portfolio that would maximize total asset position in 1 year.

The experiment was conducted on a personal computer. An interactive program presented the initial holdings distribution over the portfolio. For example, here is one initial distribution over a portfolio:

Pan Am: 30%
IBM: 20%
General Motors 30%
T-Bills: 20%

Initial distributions varied across the 12 portfolios.

Subjects received a booklet containing the *Value Line* reports and rated the stocks in each portfolio on a five-point favorableness scale. The instructions for the rating task read:

After reading the reports in the portfolio please indicate how favorable each stock appears to you. The scale ranges from 1 to 5 and is represented below:

- 1: dismal
- 2: disappointing
- 3: fair
- 4: promising
- 5: excellent

Subjects were told that the T-Bill rates fluctuated between 4 and 11% during the 9-year period from which the reports were abstracted.

Subjects in the two experimental conditions were then presented the decisions of a previous participant. (Hereafter, the previous participant is

referred to as the Other.) Subjects were told that the performance of the Other was about average—doing very well on some trials and not so well on others. The Other had purportedly been given the identical initial distribution as the subject but the nature of the Other's information was manipulated. In the high-redundancy condition the Other purportedly read the same report; in the low-redundancy condition the Other purportedly read a report different than the one given to the subject. Subjects in the low-redundancy condition were told that the Other's report consisted of financial information such as price/earnings ratios, dividends, and the company's total asset position; subjects were also told that the financial information was as equally predictive as the managerial reports that they would receive. In fact, the reallocation decisions of the previous participant were identical in both redundancy conditions. Redundancy was manipulated between subjects. Subjects in the control condition did not receive the decisions of a previous participant.

The phrase "different but equally predictive" was included in the low-redundancy condition to minimize beliefs that the different information was inherently less diagnostic. We wanted to create a situation analogous to a multiple regression with predictors that are equally correlated with the criterion.

An important variable in social influence is the consistency of another person's response with the subject's response (e.g., Goethals, 1972). Because consistency might mediate subject's sensitivity to redundancy we chose to manipulate it. A consistent response is operationalized as the Other moving in the same direction as the *Value Line* report on the target stock. For example, if the target stock signaled "buy" ("sell"), the Other would increase (decrease) holdings by 20 or 25% points. An inconsistent response is operationalized as the Other holding even though the *Value Line* report on the target stock had a clear signal to buy or to sell (i.e., the Other leaving the holdings of the target stock at the initial level, a change of 0%). For the neutral stocks the Other's reallocations consisted of either a 5 or a 10% change. The sign for each of the 5 and the 10% changes on the neutral stocks, indicating an increase or a decrease in the reallocation, was randomly assigned. Consistency was varied within-subjects. In half the rounds the action of the previous participant was consistent with the reports on the target stock and inconsistent in the remaining half. The consistency of the Other's behavior was counterbalanced over portfolios. Four different random orderings of the 12 portfolios were used in this study.

Finally, the program collected the subject's reallocation decisions (the subject's recommended distribution in light of the evidence) and a confidence judgment on the overall portfolio decision. Confidence was as-

sessed on a five-point scale ranging from low confidence to high confidence.

Recall that the stock descriptions were abstracted from *Value Line* reports published between 1972 and 1980. Because actual *Value Line* reports were used in this study it was possible to provide immediate performance feedback after each portfolio decision, i.e., how well the subject's reallocation would have fared in the actual market 1 year later. After reallocating a portfolio subjects received the actual return rate on their decisions.³

In summary, for a given portfolio subjects saw the initial distribution, rated the "favorableness" of the three stocks, made their own reallocation and confidence rating, and, finally, saw the return that their reallocation would have produced. Experimental subjects also saw the decisions of a previous participant immediately after making their favorableness ratings. The task lasted approximately 40 min.

Results

To test the influence of the previous participant the subject's reallocation is compared with the reallocation of the Other. One natural measure is to treat the reallocation of a portfolio as a four-dimensional vector and compute the distance between the subject's vector and the Other's vector. The distance between the two vectors measures the match between two sets of decisions. A second possible measure is the projection of the subject's reallocation on the Other's reallocation. The projection can be interpreted as measuring how much the subject "moved" in the direction of the Other's decision (see Berscheid, 1966, for an example of a one-dimensional version of this measure). Note that the reallocation vector is itself a difference between the initial distribution vector and the final distribution offered by the subject in light of the evidence (or, in the case of the Other, a difference between the initial distribution and the response of the Other).

The dependent measure in this study combines these two ideas to create a directional distance measure, denoted D . The directional distance D indexes the discrepancy between the responses of the subject and the responses of the Other. The direction of D is defined with respect to the Other's decisions and a high D represents high discrepancy (i.e., D is inversely related to the impact of the Other's decisions on the subject). When the match is perfect $D = 0$. For example, when the Other increases

³ The feedback on the returns contributed to making the task engaging for the subject. Many subjects kept track of their trial-by-trial return rates and computed their overall mean. This occurred without any prompting from the experimenter.

by 20 percentage points and the subject increases by 20 or more percentage points on stock i , then $D_i = 0$. When the Other increases the holding of stock i in a portfolio by 20 percentage points and the subject increases the holdings of stock i by 15 percentage points, then $D_i = 5$ (i.e., $P_i - S_i = 20 - 15$). Thus, D_i approaches zero as the subject's response approaches the response of the Other, but $D_i = 0$ whenever the subject matches or exceeds the change made by the Other. The directional distance D between the subject's reallocation vector and the Other's reallocation vector is defined as the sum of the four D_i 's comprising a given portfolio (three stocks and one T-Bill). The directional distance D can be given a geometric interpretation in terms of a city block distance metric between the subject's decision vector and a subspace (not a point because differences that exceed the response made by the Other are scored as zero) defined by the Other's decision vector.⁴

Recall that subjects in the control condition did not observe the decisions of a previous participant. It is possible, however, to compute the directional distance between the control subjects' reallocation vector and the reallocation vector corresponding to the previous participant presented in the two experimental conditions.

The mean directional distance (over subject and portfolio) for each of the three conditions is presented in Fig. 1. The obvious finding is a striking conformity effect—there is a smaller discrepancy between the subject and the Other in the two experimental conditions than in the control conditions.⁵ Subjects in both redundancy conditions are influenced by the decisions of the Other ($M = 40.86$) relative to subjects in the control condition who did not observe the Other's decisions ($M = 48.72$), $t(104) = 3.42$, $p < .001$, for the 1, 1, -2 contrast in a between-subjects ANOVA. This result suggests that, contrary to what would be expected under the redundancy hypothesis, people are influenced by the market decisions of another person.

The main hypothesis tested in this experiment is whether subjects would conform less to the decisions of the Other who had shared infor-

⁴ The directional distance measure reported throughout the paper is defined over the four elements of each portfolio. Clearly, because the distribution over the portfolio is constrained to sum to 100% there are only three degrees of freedom (i.e., the vectors can be represented in a four-dimensional simplex). We chose to define D over the four components because it provides a natural normalization of the portfolios across the consistency manipulation. Defining the distance measure over the three stocks (i.e., omitting the T-Bill category from the measure) or using a different metric, such as Euclidean distance or a projection that does include the times the subject exceeded the Other, yields a similar pattern of results.

⁵ In this paper the terms *conformity* and *influence* will be used interchangeably because the main dependent variable, D , has both distance and direction interpretations (which are analogous to conformity and influence, respectively).




FIG. 1. Subject observes another's decisions: Redundancy variable (error bars denote ± 1 standard error).

mation than to the Other who had unique information. As seen in Fig. 1, a difference between the high- and the low-redundancy conditions does not emerge (contrast n.s.). A similar pattern is observed when one examines the D_i 's for either the target stock or the neutral stocks separately (as opposed to D which is defined over the four components in each portfolio).

As one would expect, there is also an effect of the consistency between the responses of the previous participant and the *Value Line* reports. Across the two experimental conditions, subjects are more influenced by the Other when the decisions are consistent with the *Value Line* report on the target stock (mean discrepancy between subject and Other = 38.76) than inconsistent (mean discrepancy between subject and Other = 48.25), main effect $F(1,104) = 38.14$, $p < .0001$ (using a 2×2 mixed ANOVA with redundancy and consistency as factors). As seen in Fig. 2, consistency does not interact with the redundancy manipulation.

At the end of each round subjects indicated their overall confidence in their reallocation of the portfolio on a five-point scale (1 means low confidence and 5 means high confidence). There are no differences between the three conditions on confidence. The mean confidence ratings for the high-redundancy, low-redundancy, and control conditions are 3.27, 3.29,




FIG. 2. Subject observes decisions: Redundancy vs consistency variables (error bars denote ± 1 standard error).

and 3.37, respectively. There is also no difference in confidence between the trials with consistent behavior on the part of the Other ($M = 3.31$) and trials with inconsistent behavior ($M = 3.25$).

Discussion

The results of Study 1 show that subjects were influenced by the decisions of another person. That is, the decisions of subjects who did not see the responses of a previous participant differ from the decisions of subjects who did. An additional result observed during debriefing is that the majority of subjects in the experimental conditions claimed that they were not influenced by the Other (even though their actual responses differed from those recorded under the control condition). This discrepancy between self-report and actual behavior is a common observation in conformity research (e.g., Asch, 1952; Nisbett & Wilson, 1977; Sherif, 1936).

The main purpose of this first experiment was to test the redundancy hypothesis. Subjects were not more influenced by the decisions of someone from whom they could infer new information than from someone who had the same information they did (and, thus, from whom new

information about the stock could not be inferred). It should be noted that during debriefing all subjects in the two experimental conditions correctly identified whether the Other had read a different report or the same report. Thus, it is not the case that subjects did not remember or attend to the Other's information source. Also, the main effect for consistency suggests that subjects were sensitive to the direction of the Other's decisions.

Our result goes against the normative argument because, despite awareness of the Other's information and sensitivity to the direction of the Other's decision, subjects did not distinguish between the two levels of redundancy. Subjects' decisions imply that they did not give more weight to the previous participant with low-redundant information.

The next study shows that if people observe another's *evaluation* of the information, then they do behave normatively—subjects give more weight to the evaluations of someone having unique information than to someone having shared information. The positive result observed in Study 2 gets around the logical problem encountered in the present study that the redundancy hypothesis predicts in favor of the null hypothesis. We defer this issue until the discussion of Study 2.

STUDY 2: OBSERVING EVALUATIONS

The procedure in this study focuses the subject's attention on the information the Other has by presenting the Other's evaluations rather than decisions. Subjects were presented the ratings of the *Value Line* reports purportedly made by the previous participant. We suggest that the evaluations make salient the information on which the responses are based, thus highlighting any redundancy. Seeing someone who has the same report you have give positive information a rating of "excellent" carries less diagnostic impact than seeing someone with a different report rate her information as "excellent." The latter case can, in effect, be treated as a sample of two. We call this the *salience hypothesis*. When the Other's information is made salient subjects should be less influenced by the evaluations in the high-redundancy condition than in the low-redundancy condition.

A comparison of the first two studies (i.e., the subject observed either the decisions or the evaluations of the Other) provides a test for another explanation of excessive trading. Clearly, several factors can influence an individual's investment behavior. People change a portfolio not only in response to changes in the favorability of stocks but also because of changes in liquidity and risk preferences, changes in other investment markets, etc. (Sharpe, 1985). Thus, an observer does not know whether to attribute the cause of another's decision to the information she must have had, to her investment preferences, or to any number of other explana-

tions (cf. Kelley's, 1972a, 1972b, discounting principle and analysis of causal schemata). We call this the *alternative attributions hypothesis*.

An observer who sees another's decision must infer the evaluation from the decision (e.g., "He's buying. He must have evaluated his information positively."). The Other's evaluation is thus a step removed. In contrast, when observing how another rated a particular stock one has direct access to his evaluation. The alternative attribution hypothesis implies that the decisions of the Other will have less impact because decisions, unlike ratings, can be attributed to other factors. Thus, the alternative attribution hypothesis implies more conformity (i.e., lower directional distance) when subjects see evaluations than when they see decisions.

Except for a small change in the instructions the materials and procedure were the same as in Study 1. The instructions were reworded to reflect the change from decisions to ratings. The ratings of the previous participant involved a simple transformation of the reallocation decisions presented in the first study. Recall that in Study 1 the reallocation decisions made by the Other on the target stock were either buy (20 or 25% increase), sell (20 or 25% decrease), or hold (no change). Small reallocations (± 5 or $\pm 10\%$) were made on the two neutral stock. In the present study the decisions of the Other were recoded into a five-point rating scale assessing the favorableness of the stock description. The recordings are listed below:

Decision	Rating
- 20 or - 25% →	1 (dismal)
- 5 or - 10% →	2 (disappointing)
0% →	3 (fair)
5 or 10% →	4 (promising)
20 or 25% →	5 (excellent)

Sixty-four subjects participated in this experiment (32 in each of the two redundancy conditions). Each subject received \$5.00 for their participation. Data from the 36 subjects in the control condition of Study 1 serve as a comparison.⁶

Results

The goal of Study 2 is to make salient the Other's information by presenting subjects the favorableness ratings purportedly made by the previous participant. To facilitate comparison between these results and

⁶ A second group of control subjects ($N = 37$) was tested as a replication. The results were comparable to those of the original group of 36 control subjects.

those of Study 1, the Other's ratings were transformed back into the decision scale so that the directional distance is on the same scale.

The mean directional distances are presented in Fig. 3. Recall that the directional distance indexes the discrepancy between the responses of the subject and the responses of the Other. The hypothesis that subjects would conform less when observing the ratings of someone with shared information is supported. The directional distance between the subject and the Other is greater in the high-redundancy condition ($M = 45.51$) than in the low-redundancy condition ($M = 39.13$), $t(97) = 2.78$, $p < .01$, for the 1, -1, 0 contrast in a between-subjects ANOVA. In fact, the mean directional distance in the high-redundancy condition does not differ significantly from the control condition, $t(97) = 1.44$, $p = .15$. Similar results are observed when the analysis is conducted on the target stocks and the mixed stocks separately.

The main effect for consistency observed in Study 1 is replicated. Subjects conform more when the Other's evaluation is consistent with the report on the target stock ($M = 40.82$) than inconsistent ($M = 48.4$), $F(1,97) = 19.88$, $p < .0001$. As in Study 1, there is no evidence of an interaction with the redundancy variable.

The confidence measure does not differ across redundancy conditions.




FIG. 3. Subject observes another's ratings: Redundancy variable (error bars denote ± 1 standard error).

The mean rating on the five point confidence scale is 3.28 for the low redundancy condition and 3.17 for the high redundancy condition ($F < 1$). Recall that the mean confidence rating in the control condition was 3.37. The mean confidence ratings for the consistent and inconsistent trials are 3.22 and 3.19, respectively.

Discussion

The findings in Study 2 support the salience hypothesis: when directly observing another's evaluations people do consider redundancy in the other's information. Subjects were less influenced by the evaluation of someone who read the same report than by the evaluation of someone who read a different report. Note that this finding contrasts with the null result found when people had to infer such evaluations from the decisions (Study 1). Thus, taking the first two studies together, an interaction between the redundancy manipulation and the nature of the Other's response (decision or evaluation) is observed (depicted in Fig. 4). In the case of low redundancy there was no difference between the evaluation and decision conditions, but in the case of high redundancy subjects were more influenced by decisions than evaluations.

The interaction shown in Fig. 4 has been replicated using a paper-and-pencil task involving decisions on single stocks (Gonzalez & Tversky, 1990). Gonzalez and Tversky also offered subjects financial incentives contingent on performance.

The normative argument that people should give less weight to the

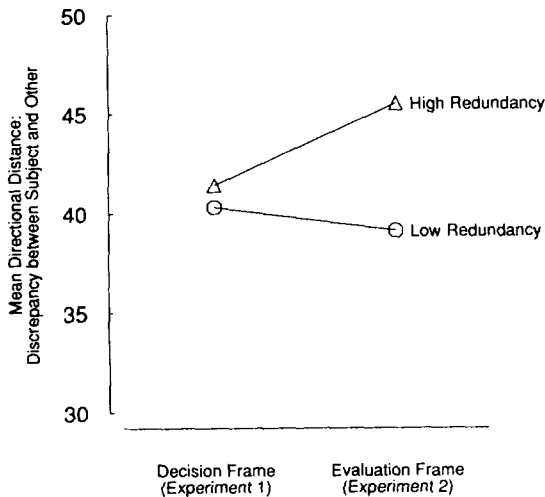


FIG. 4. Comparing Experiments 1 and 2.

responses of someone with redundant information implies that there should be a main effect for redundancy. However, there was no redundancy effect in the decision frame. From the standpoint of the normative model, the pattern in Fig. 4 (along with data from the control condition) suggests that subjects in the decision frame gave too much weight to the responses of the previous participant with shared information. That is, the simple normative model implies that the directional distance in the high-redundancy condition should not differ from that of the control condition (i.e., receiving nondiagnostic information should be equivalent to receiving no information at all, cf. Nisbett, Zukier, & Lemley, 1981). Because this result was observed in the evaluation frame (Fig. 3) but not in the decision frame (Fig. 1) one can argue that people in the high-redundancy condition gave too much weight to the Other's decisions. Further, while it can be argued that another's evaluations reveal more about the information they possess than do their decisions (along the lines of the alternative attributions hypothesis) it is not obvious that, in this context, decisions should be more revealing than evaluations. Thus, it is unlikely that subjects gave too little weight to the decisions of the Other with low redundant information.

How do the alternative attribution hypothesis and the redundancy hypothesis stand given these results? The interaction between the first two studies rule out the alternative attribution hypothesis because the discrepancy between the subject and Other was not lower in the evaluation frame than in the decision frame. That is, subjects were not more influenced in the evaluation frame. The redundancy hypothesis does not fare very well either. Even though subjects are not sensitive to redundancy when observing another's decisions, they do appreciate redundancy when observing another's evaluations. The salience hypothesis, however, cannot be rejected. It appears that subjects became sensitive to redundancy when the Other's information was made salient. People behaved normatively in the evaluation frame, but not in the decision frame.

The next three studies feature manipulations designed to encourage normative behavior while still presenting to subjects the *decisions* of the Other. In Study 3 the diagnosticity of the Other's decisions is increased. Subjects believed that they observed the average decision from a *group* of previously tested participants. Will increasing the diagnosticity of the decisions prompt subjects to be sensitive to the redundancy in the Other's information? In Study 4 the redundancy is made distinctive by the use of a within-subjects design where subjects observed the decisions of two previous participants—one had the same report as the subject, the other had a different report. In Study 5 attention is focused on the Other's information by having subjects estimate from the observed decision how the Other evaluated the information he or she read.

STUDY 3: A DECISION FROM A GROUP

In Study 3 we attempt to increase the diagnosticity of the previous participant by framing the decisions as an average from a “previous group of 80 participants.” Presumably, a cue coming from 80 people will be given more weight than a cue coming from a single person (assuming that all previous participants are drawn from a population of individuals with comparable expertise). Subjects in the high-redundancy condition were told the 80 previous participants had read the same report; subjects in the low-redundancy condition were told that the group of 80 previous participants had read a different but equally predictive report than the one they were given. By making the cue more diagnostic subjects may become more sensitive to the redundancy in the Other’s information.

Method

The reallocation decisions presented to subjects were identical to those in Study 1. The instructions were modified to reflect the change in framing from a single participant to the average reallocation from a group of previous participants.

Seventy-nine subjects participated in this experiment (40 in the low-redundancy condition and 39 in the high-redundancy condition). Each subject received \$5.00 for their participation. Subjects were recruited from undergraduate dormitories. Data from the 36 control subjects in Study 1 are also used.

Results

The dependent variable is the directional distance described in Study 1. The mean directional distances (over subject and portfolio) appear in Fig. 5. As observed in Study 1, subjects in both the high- and the low-redundancy conditions are influenced by the “group” response ($M = 32.2$) relative to subjects in the control condition ($M = 48.72$), $F(1,112) = 86.64$, $p < .0001$, for the 1, 1, -2 contrast. Again, a difference between the two redundancy conditions fails to emerge (contrast n.s.). Thus, even though subjects are influenced by the Other’s decisions they did not give differential weight in the two experimental conditions.

The finding that subjects conform more when the Others’ decisions and the reports on the target stock are consistent ($M = 20.53$) rather than inconsistent ($M = 31.69$) is replicated, main effect $F(1,100) = 117.51$, $p < .0001$. The consistency variable does not interact with the redundancy variable. Again, there is no difference between the confidence ratings in the low-redundancy ($M = 3.34$) and the high-redundancy ($M = 3.22$) conditions.

The effect of framing the decisions as coming from either a single “participant” (Study 1) or from a group of “participants” can also be as-




FIG. 5. Subject observes decisions of group average: Redundancy variable (error bars denote ± 1 standard error).

essed. A $2 \times 2 \times 2$ ANOVA with redundancy (low vs high), framing of the previous participant (individual vs group average), and consistency as factors reveals a main effect for framing the reallocations as decisions from a single person or a group. Subjects who observe the decisions corresponding to a group average are more influenced ($M = 32.02$) than subjects who believe the (identical) decisions come from a single previous participant ($M = 40.86$), $F(1,146) = 23.10$, $p < .0001$. This result serves as a manipulation check and confirms our assumption that subjects would give more weight to the group average than to the decisions of a single person. The only other effect to reach statistical significance in this three-way ANOVA is the main effect for consistency, $F(1,146) = 46.84$, $p < .0001$.

Discussion

The "group" frame was successful in increasing the diagnosticity of the Others' decisions. Subjects were more influenced by the group than by the single participant even though the numerical stimuli were identical in both cases. However, the increased diagnosticity did not influence sensitivity to redundancy: subjects did not exhibit differential levels of influence in the two experimental conditions. Despite the increase in diagnos-

ticity and conformity, subjects failed to appreciate the normative implications of the two levels of redundancy.

STUDY 4: THE DECISIONS OF TWO

In Study 4 we attempt to highlight redundancy through a within-subjects design. Subjects observed the decisions of two previous participants concurrently. One previous participant purportedly read the same report as the subject (high redundancy); the other previous participant purportedly read a different but equally predictive report (low redundancy). By presenting these two previous participants in the same trial the redundancy in the other's information becomes a distinctive feature.

This study also addresses alternative explanations involving the possibility that subjects were not attending to the Other's responses. While the equal weighting of high- and low-redundancy conditions observed in Studies 1 and 3 cannot be substantiated on normative grounds, a null finding is problematic and can always be challenged. For instance, one can question whether or not subjects attended to the source of the Other's information (i.e., same report or different report). Note that the 100% correct recall of the Other's information source observed in Study 1 and the significant difference found between the two redundancy conditions in the evaluation frame (Study 2) are inconsistent with this critique. The present study offers a more rigorous test that subjects, in this paradigm, do encode and have access to the nature of the report read by the previous participant[s].

This study examines an additional aspect of redundancy: the variability of the Other's responses. In the first three studies the responses of the previous participant[s] in both the high- and the low-redundancy conditions were identical. However, in the real world the responses of someone with unique information can be more variable than the responses of someone with shared information. One expects to observe extreme responses (such as the other person "buys" when the information you read signals "sell") more frequently when the other has access to unique information. There is a sense in which the first three studies sacrificed external validity by constraining the responses in the two redundancy conditions to be identical. This may have worked against us in a fundamental way: variability may have been eliminated as a cue for redundancy.

The variability of the Other's responses in the low-redundancy condition (i.e., the previous participant who had a different report) was manipulated. Two variability conditions were tested. In the *extreme variability condition* trials where the Other made extreme responses were included. That is, on some trials the previous participant with nonredundant information behaved in a manner opposite to the signal present in the subject's *Value Line* report. The case where the previous participant

with redundant information goes against the report on the target stock is not included because of possible discounting effects. Consider, for example, reading a report that gives a favorable forecast. You observe that a person who read the same report is "selling" their holdings of this promising stock. To the extent the person with the same information does not transact in an "obvious" manner you may discredit his investment judgment.⁷

In the *identical variability condition* the variability of the responses of both previous participants was identical. This second variability condition is included to allow comparisons with the previous three studies. In short, the two variability conditions differed only in the consistency between the decision of the Other with unique information and the subject's *Value Line* report.

Method

The materials and procedure were similar to those used in Study 1. The main difference was that subjects observed the decisions of two previous participants. One previous participant purportedly read the same report as the subject, and the other purportedly read a different but equally predictive report. Thus, unlike the three previous studies, redundancy was manipulated within subjects.

The variability of the Other's responses with the *Value Line* report were also manipulated. There were two levels of variability (a between-subjects factor). In the identical variability condition both previous participants exhibited the same level of variability relative to the *Value Line* reports. In the extreme variability condition the two previous participants exhibited different levels of variability. The previous participant with the different report sometimes exhibited the opposite decision than would be implied by the *Value Line* report (e.g., a "sell" response when the subject's *Value Line* reported signaled "buy"). We now discuss the responses of the previous participants in more detail.

Identical variability condition. First, consider the target stock in each portfolio. In two of the twelve portfolios the decisions of both previous participants went in similar directions (i.e., both previous participants either buy or sell when the *Value Line* report signals buy or sell, respectively). On the remaining ten portfolios one previous participant behaved in a consistent manner (buying or selling in accord with the *Value Line* report on the target stock) and the other previous participant behaved in an inconsistent manner ("holding," or not changing the initial holdings).

⁷ As pointed out by a reviewer, this assumption has interesting implications and should be explored in subsequent research. The discussion section to this study reports one preliminary study addressing this point.

The buy and sell reallocations were indexed by changes of ± 20 and ± 25 percentage points; the hold reallocation was indexed by a change of 0%. Consistency was counterbalanced across the two previous participants. For the two neutral stocks in each portfolio a reallocation from the set $\{+5\%, -5\%, +10\%, -10\%\}$ was randomly selected for each previous participant.

Extreme variability condition. In 4 of the 12 portfolios the behavior of the previous participant with highly redundant information was consistent with the target stock description (buy or sell) but the previous participant with low-redundant information did not change the reallocation (hold). Conversely, in 4 other portfolios the previous participants reversed the above roles. In the 4 remaining portfolios the behavior of the previous participant with the same information as the subject (high redundancy) was consistent with the *Value Line* report on the target stock and the behavior of the previous participant with different but equally predictive information (low redundancy) was the opposite; e.g., the previous participant with the same report buys and the previous participant with a different report sells. The buy, sell, and hold reallocations are indexed by 20, -20 , and 0%, respectively. The reallocations for the neutral stocks were $\pm 5\%$. The direction of the changes for the neutral stocks (an increase or a decrease) were fixed so the two previous participants were negatively correlated.

Both variability conditions. For each of the 12 portfolios the nature of the previous participant's report (i.e., the labels: "same" or "different") was displayed above the respective reallocation decisions on the computer monitor. Four random orders of the 12 portfolios were used. Also, the ordering of the two previous participants (which one was presented first to the subject) was counterbalanced.

Forty-seven subjects participated in this experiment (24 in the extreme variability condition and 23 in the identical variability condition). Subjects received \$5.00 for their participation (except for 9 subjects in the extreme variability condition who received partial credit for their introductory psychology course requirements).

Results

Two directional distance measures are computed for each portfolio. One is the directional distance between the subject's reallocation vector and the reallocation vector of the previous participant with the same report. The second is the directional distance between the subject's reallocation vector and the reallocation vector of the previous participant with a different report.

To facilitate comparisons with the previous three studies we first present the results for the subset of portfolios where one previous par-

participant acted in a consistent manner with the *Value Line* report on the target stock but the other did not change the allocation (hold). That is, we collapse over the variability condition but *exclude* (1) the two portfolios in the identical variability condition where both previous participants acted in a consistent manner with the *Value Line* report (e.g., both buy the target stock when the subject reads a favorable *Value Line* report) and (2) the four portfolios in the extreme variability condition where the previous participant with unique information acted in an opposite manner to the *Value Line* report on the target stock. This creates a 2×2 within-subjects ANOVA with redundancy and variability as factors.

The mean directional distances are presented in Fig. 6. Again, a significant difference between the high- and the low-redundancy conditions fails to emerge. Subjects are not influenced by the different degrees of redundancy ($F < 1$). Power analysis reveals that approximately 400 subjects are required to detect a significant difference between the low-redundancy ($M = 37.5$) and the high-redundancy ($M = 39.7$) conditions, setting the probability of a Type II error at 0.20 and using sample estimates in the calculation.

The main effect for consistency is replicated. When faced with the decisions of two previous participants subjects are more influenced by the




FIG. 6. Subject observes decisions of two "previous participants": Redundancy variable (collapsing over variability).

decisions of the participant who was consistent ($M = 35.30$) than inconsistent ($M = 41.95$), $F(1,45) = 12.82$, $p < .001$. As observed in the previous studies, the interaction between redundancy and consistency is not significant.

The main effect for consistency is qualified by an interaction with variability (extreme vs identical). The extreme variability condition shows a strong consistency effect, but the identical variability condition does not. The 8 critical portfolios in the extreme variability condition (i.e., one previous participant is consistent on the target stock, the other holds) shows a difference between the consistent ($M = 33.75$) and inconsistent ($M = 46.40$) conditions. However, on the 10 critical portfolios in the identical variability condition there is no difference between the consistent ($M = 36.8$) and the inconsistent ($M = 37.3$) conditions. A $2 \times 2 \times 2$ within-subjects ANOVA with variability, redundancy and consistency as factors reveals the interaction between variability and consistency is statistically significant, $F(1,45) = 10.92$, $p < .01$. Note that this interaction amounts to a comparison of the effects of the excluded trials on the critical portfolios. That is, the identical variability condition is defined by two trials where both previous participants agree with each other and with the subject's *Value Line* report; the extreme variability condition is defined by four trials where the previous participant with low-redundant information goes against the *Value Line* report. The results suggest that the four portfolios in the extreme variability condition served as strong cues for consistency and influenced the remaining eight trials.

We now turn to an analysis of the portfolios that define the two variability conditions. In the identical variability condition subjects are equally influenced by both the low- and the high-redundant previous participants (averaging over the two portfolios: $M = 18.8$ and $M = 18.9$, respectively). That is, when both participants are consistent with the *Value Line* report subjects are not influenced differentially by them. However, note that the directional distances are considerably lower than those observed for the remaining ten portfolios, reflecting a strong conformity effect on those two trials. In the extreme variability condition, on the other hand, subjects are more influenced by the participant who is consistent (and has high-redundant information) than the participant who acts in an opposite manner (and has low-redundant information). Averaging over the four extreme variability portfolios, directional distances between the subject and Other are $M = 37.92$ and $M = 78.39$, respectively. Interpreting the result on this subset of items is problematic because redundancy and consistency are confounded.

Again, there is no difference in confidence. The only comparison involving confidence that approaches statistical significance is a main effect

for variability. The mean confidence rating across subjects and portfolios is 3.27 in the extreme variability condition and 2.97 in the identical variability condition, $t(45) = 1.41$, $p = .17$.

Discussion

Study 4 shows that people did not appreciate the different degrees of diagnosticity across the two levels of redundancy even though the information source was made to be a distinctive feature. Contrasting the two sources of information was not sufficient for subjects to give more weight to the *decisions* of someone with shared information than to someone with unique information.

Recall that variability was manipulated to check the critique that our studies had eliminated an important cue to redundancy. The results suggest that subjects were less influenced by the Other who sometimes made inconsistent responses with respect to the subject's information. Thus, variability appears to have the opposite effect: Subjects give less weight to the more variable, but more diagnostic, cue (see Murrell, 1977, for a similar result using a categorization paradigm).

Similar conclusions regarding variability were made in a replication of Study 1 (Gonzalez, 1991). Subjects in the low-redundancy condition observed the previous participant choose in an incongruent manner with their own information in 2 of the 12 portfolios. The remaining 10 portfolios in the low-redundancy condition and all 12 portfolios in the high-redundancy condition were identical to those in Study 1. Redundancy was manipulated between subjects. The results were counternormative: subjects were more influenced by the Other who had read the same report than the Other who had read a different report and whose responses were more variable. It appears that including two trials where the Other with unique information provides extreme responses in the "wrong" direction serves to discredit him. As previously noted, these results must be interpreted with caution due to the confound between redundancy and consistency.

STUDY 5: INFERRING EVALUATIONS

The general pattern we observe, that people do not appreciate redundancy in another's information when observing *decisions* but do when observing *evaluations*, raised an interesting question. Why don't people infer the Other's evaluation from the decision? If you discover that someone is buying more shares of a particular stock, then you might infer that he or she interpreted the available information as promising or favorable. In the present study subjects were asked to infer the Other's evaluation from the decision the Other made. In effect, this is a conceptual replication of Study 2 except that subjects infer the evaluations rather than

observe them directly. Will inferred evaluations have the same impact as the evaluation frame presented in Study 2?

Method

The procedure and materials were identical to those in Study 1. The only addition was that subjects were asked to infer, or estimate, how the previous participant rated the information. Subjects were instructed to use the Other's reallocation to estimate the rating; subjects estimated the ratings immediately after observing the Other's decisions. They then made their final reallocations and confidence ratings. Subjects saw the decisions of one previous participant.

Eighty-one subjects participated in this experiment (41 in the high-redundancy condition and 40 in the low-redundancy condition). They were recruited from dormitories at Stanford University and were paid \$5.00. Data from the 36 control subjects tested in Study 1 will serve as a baseline.

Results

The pattern of results replicates Study 1—we find a strong conformity effect, but fail to observe sensitivity to the two levels of redundancy. The mean directional distances for each condition are presented in Fig. 7. Subjects who observe the decisions of the Other have lower discrepancy scores ($M = 37.81$) than subjects who do not ($M = 48.72$), $F(1,114) = 33.36$, $p < .0001$, as tested by a 1, 1, -2 contrast. As in Study 1, a difference between the high- and low-redundancy conditions fails to emerge (contrast n.s.). Subjects are not less influenced in the high-redundancy condition than in the low-redundancy condition. A similar pattern is observed in the target stocks and neutral stocks separately.

It appears that there is a trend in the direction of Study 2. Power analyses reveal that, using sample estimates, approximately 130 subjects in each of the two cells would be needed to achieve a Type II error rate of 20%. Thus, as can be seen in Fig. 7, the trend is quite weak.

The consistency variable again shows a main effect (2×2 ANOVA with consistency and redundancy as factors). Subjects are more influenced when the Other behaves in a manner consistent with the *Value Line* report on the target stock ($M = 35.89$) rather than inconsistent ($M = 46.44$), $F(1,114) = 68.78$, $p < .0001$. Consistency does not interact with the redundancy variable. As in the previous studies, there are no confidence effects.

A visual comparison between the present data and the data in Study 1 suggests that overall conformity increases when subjects are asked to infer the evaluation of the Other. A 2×2 between-subjects ANOVA with redundancy and whether or not subjects explicitly inferred the Other's




FIG. 7. Subject observes decisions and infers rating: Redundancy variable (error bars denote ± 1 standard error).

rating (Study 1 vs Study 5) as factors suggests a small, but significant effect. Subjects who are asked to infer the Other's rating have lower discrepancy scores ($M = 37.81$) than subjects who do not ($M = 40.86$), $F(1,148) = 2.76$, $p < .01$. Thus, subjects conform more when they inferred the Other's ratings.

The subjects' estimates of the Other's ratings were also analyzed. Table 1 shows the median estimates (across trials and subjects) conditional on the decisions of the Other. Note that the medians in Table 1 are similar to the mapping from decision to rating used in Study 2. Recall that Study 2 simply recoded the reallocation decisions into a five-point rating scale. When subjects in Study 2 observed the evaluations they gave less weight

TABLE 1
MEDIAN RATINGS INFERRED FROM THE DECISIONS OF "PREVIOUS PARTICIPANT"

Reallocation decisions by the other	Subjects' estimated rating (median)
-20 or -25%	→ 2
-5 or -10%	→ 2
0%	→ 3
5 or 10%	→ 3
20 or 25%	→ 5

to the responses of the Other who had the same information that they did. However, in the present experiment, even though subjects essentially inferred the mapping used in Study 2, they did not give less weight to the decisions of the Other who had the same information. One could argue that the inferred ratings are regressive relative to the mapping used in Study 2 and that this may diminish the impact of the ratings. However, an internal analysis suggests that regression effects cannot explain the result. Mean directional distances for the subset of subjects whose inferred estimates were identical to the mapping used in Study 2 show the same pattern—no difference between the two redundancy conditions.

Discussion

The findings in the present study are provocative. People can infer an evaluation from another's decision, but they do not give less weight to the decisions of someone with the same report than to someone with a different report. In other words, while subjects went "beyond the information given" and inferred the evaluation from the decision they did not discount according to redundancy. This contrasts with the finding in Study 2 that people who observed another's evaluation directly responded differentially to high and low redundancy. One explanation for the result that the inferred evaluation produced a different impact than the evaluation subjects observed directly is Slovic's (1972) concreteness hypothesis. Slovic stated

. . . that a judge or decision maker tends to use only the information that is explicitly displayed in the stimulus object and will use it only in the form in which it is displayed. Information that has to be stored in memory, inferred from the explicit display, or transformed tends to be discounted or ignored. (p. 14)

Thus, with regard to assessing the redundancy of another's evidence, subject-generated information carried less weight than information directly presented to the subject. However, asking subjects to infer the Other's evaluation did increase overall conformity relative to simply observing the Other's evaluation; i.e., subjects were more influenced in Study 5 than in Study 1.

GENERAL DISCUSSION

These studies were prompted by the hypothesis that there is "too much" trading in the stock market (e.g., Arrow, 1986). We initially offered one explanation: investors do not use the behavior of others to infer the quality of a stock. The present findings suggest, however, that people are influenced by the decisions of others and the pattern of influence is consistent with the salience hypothesis. We found that the responses of subjects who observed the decisions of another person differed from the

responses of subjects who did not. Yet, subjects did not seem to differentiate investors from whose decisions they could learn more (i.e., low redundancy) and investors from whose decisions they could learn less (i.e., high redundancy), except when directly observing the other's evaluations. Although the relevance of these results to excessive trading in the stock market is limited (subjects did not invest actual money, there were no transaction costs, subjects were undergraduates), the results are nonetheless interesting, consistent, and worthy of further investigation.

Alternative Explanations

What makes the decision frame different from the evaluation frame? These experiments do not offer a complete account, but some explanations can be ruled out. For example, the discussion of Study 2 mentions two explanations (the redundancy and alternative attributions hypotheses) that are not completely consistent with the data.

Even though an attributional analysis cannot completely account for the present results, there may be some benefit to casting the problem in a more general attributional framework. When observing another's decisions do we learn something about the person who made the decision or about the quality of the alternatives? Do we interpret another's stock decisions as "information-motivated" or "liquidity-motivated" (Sharpe, 1985)? What factors lead us to favor "person" inferences? What factors lead us to favor "decision alternatives" inferences? The attributional framework may be a natural and promising way to think about the inferences made from another's decisions.

Stillinger (1989) applied a general attributional analysis to understand how adversaries value options in a negotiation context. She proposed that "when an opponent offers a concession, the recipient is disposed to ask the attributional question, 'why is my adversary offering a concession, why *this* concession, and why *now*'" (p. 58). The attributional problem here is that the recipient may interpret an adversary's concession as a signal that the offer may be less than adequate.

A simple anchoring explanation can also be ruled out. One could imagine that the decision frame helped define the reallocation scale for the subject (i.e., how much to buy or sell). However, the evaluation frame, albeit informative, was not compatible with the decision scale the subject used when making the reallocation and, thus, could not serve as an anchor as easily (cf. Slovic, Griffin, & Tversky, 1990). While the decision frame probably served as an anchor, it seems unlikely that an anchoring effect accounts for the pattern in Fig. 4 because in the low-redundancy condition subjects were equally influenced by the previous participant in both the decision and the evaluation frames.

Measuring Conformity and Influence

A deeper, unresolved question deals with the concepts of conformity and influence. How do we define them? Are they different concepts? What are appropriate measures? As Abelson (1967, p. 11) pointed out, "assessing the 'strength of influence' of one individual upon another" is not a trivial task. See March (1955) for a set-theoretic approach to the measurement of influence.

For these studies we developed a dependent measure, the directional distance, that combines the intuitive ideas of conformity and influence. The directional distance served as an index of how much the subjects' decisions were affected by the responses of another person. Fortunately, conclusions made from conformity measures (e.g., Euclidean distance) and an influence measure (projection) yielded an identical pattern of results. But, in general, these two measures need not coincide. One can imagine situations where conformity is the key variable, e.g., when there are pressures toward uniformity such as in the Sherif paradigm (1936) and in the Asch paradigm (1952), or when there is a need to average such as in a group decision task (Allport, 1924). On the other hand, one can imagine situations where influence is the key variable such as in the group polarization effect (e.g., Kogan & Wallach, 1966; Myers & Lamm, 1976). The measurement of conformity and influence is both a theoretical and a methodological question that deserves further attention.

Extensions

The present paradigm provides a new way to assess conformity and influence. It was applied to test a specific hypothesis regarding the impact of another's redundant information. We are extending this paradigm to study (counternormative) cases when someone with redundant information is *more* influential than someone with nonredundant information. We find that if the Other's responses are construed as confirming a judgment, then more weight is given to the person with redundant information (Gonzalez, 1991). This extension fits nicely with predictions made within a social comparison framework (e.g., Goethals & Darley, 1977) and is also consistent with the "illusion of validity" effect (Kahneman & Tversky, 1973).

The issue of redundant information has been studied by social psychologists in other domains. Redundancy appears to influence the magnitude of the group polarization effect. Kaplan (1977) found that juries composed of members having nonredundant information polarized more than juries whose members had redundant information. The dependent variable was the rating of a defendant's guilt. Redundancy has also been examined in the context of impression formation (e.g., Dustin & Baldwin, 1966; Wyer,

1970). Modeling the effects of redundancy, however, is nontrivial (see, for example, Winkler, 1981, who discusses the difficulties in modeling the combination of correlated evidence into a single response).

REFERENCES

- Abelson, R. P. (1967). Mathematical models in social psychology. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 3, pp. 1-54). New York: Academic Press.
- Akerlof, G. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, *84*, 488-500.
- Allport, F. H. (1924). *Social psychology*. Boston: The Riverside Press.
- Arrow, K. (1986). Rationality of self and others in an economic system. *Journal of Business*, *59*, 385-399.
- Asch, S. E. (1952). *Social psychology*. New York: Prentice-Hall.
- Berscheid, E. (1966). Opinion change and communicator-communicatee similarity and dissimilarity. *Journal of Personality and Social Psychology*, *4*, 670-680.
- Brennan, M. J., & Hughes, P. J. (1990). *Stock prices and the supply of information*. Unpublished manuscript: University of California, Los Angeles and University of Southern California.
- Clemen, R. T., & Winkler, R. L. (1985). Limits for the precision and value of information from dependent sources. *Operations Research*, *33*, 427-442.
- Dustin, D. S., & Baldwin, P. M. (1966). Redundancy in impression formation. *Journal of Personality and Social Psychology*, *3*, 500-506.
- French, S. (1980). Updating of belief in the light of someone else's opinion. *Journal of the Royal Statistical Society (Series A)*, *143*, 43-48.
- Goethals, G. (1972). Consensus and modality in the attributional process. *Journal of Personality and Social Psychology*, *21*, 84-92.
- Goethals, G., & Darley, J. M. (1977). Social comparison theory: An attributional approach. In J. Suls & R. Miller (Eds.), *Social comparison processes* (pp. 259-278). Washington: Hemisphere Publishing.
- Gonzalez, R. (1991). *The impact of others' responses: Confirming an evaluation*. Unpublished manuscript: University of Washington.
- Gonzalez, R., & Tversky, A. (1990). The impact of others' choice on investment decisions. In K. Borcherding, O. Larichev, & D. Messick (Eds.), *Contemporary issues in decision making* (pp. 367-373). The Netherlands: North-Holland.
- Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *Journal of Law and Economics*, *24*, 461-483.
- Hayes (1984). Unions and strikes with asymmetric information. *Journal of Labor Economics*, *2*, 57-83.
- Heider, F. (1958). *The psychology of interpersonal relations*. Hillsdale, NJ: Erlbaum.
- Horwich, P. (1982). *Probability and evidence*. Cambridge: Cambridge University Press.
- Jaffe, J. (1974). Special information and insider trading. *Journal of Business*, *47*, 410-428.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 219-266). New York: Academic Press.
- Jones, E. E., Davis, K. E., & Gergen, K. J. (1961). Role-playing variations and their informational value for person perception. *Journal of Abnormal and Social Psychology*, *63*, 302-310.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*, 237-251.

- Kaplan, M. F. (1977). Discussion polarization effects in a modified jury decision paradigm: Informational influences. *Sociometry*, *40*, 262-271.
- Kassin, S. M. (1979). Consensus information, prediction, and causal attribution: A review of the literature and issues. *Journal of Personality and Social Psychology*, *37*, 1966-1981.
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 192-240). Lincoln, NE: University of Nebraska Press.
- Kelley, H. H. (1972a). Attribution in social interaction. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 1-26). New Jersey: General Learning Press.
- Kelley, H. H. (1972b). Causal schemata and the attribution process. In E. Jones, D. Kanouse, H. Kelley, R. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior*. Morristown, NJ: General Learning Press.
- Kogan, N., & Wallach, M. (1966). Modification of a judgmental style through group interaction. *Journal of Personality and Social Psychology*, *4*, 165-174.
- March, J. (1955). An introduction to the theory and measurement of influence. *American Political Science Review*, *49*, 431-451.
- Murrell, G. A. (1977). Combination of evidence in perceptual judgment. In M. Kaplan & S. Schwartz (Eds.), *Human judgment and decision processes in applied settings* (pp. 169-201). New York: Academic Press.
- Myers, D. G., & Lamm, H. (1976). The group polarization phenomenon. *Psychological Bulletin*, *83*, 602-627.
- Nisbett, R., & Wilson, T. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231-258.
- Nisbett, R., Zukier, H., & Lemley, R. (1981). The dilution effect: Nondiagnostic information weakens the implications of diagnostic information. *Cognitive Psychology*, *13*, 248-277.
- Plott, C. R., & Sunder, S. (1988). *Efficiency of experimental security markets with insider information: An application of rational expectations models*. Unpublished manuscript.
- Saxe, J. G. (1886). *Poems*. Boston: Ticknor and Fields.
- Sharpe, W. F. (1985). *Investments*. New Jersey: Prentice-Hall.
- Sherif, M. (1936). *The psychology of norms*. New York: Harper & Row.
- Slovic, P. (1972). *From Shakespeare to Simon: Speculations—and some evidence—about man's ability to process information* (Report No. 12(2)). Eugene, Oregon: Oregon Research Institute.
- Slovic, P., Griffin, D., & Tversky, A. (1990). Compatibility effects in judgment and choice. In R. M. Hogarth (Ed.), *Insights in decision making: Theory and applications*. (pp. 5-27). Chicago: University of Chicago Press.
- Spence, A. M. (1974). *Market signalling: Informational transfer in hiring and related screening processes*. Cambridge, MA: Harvard University Press.
- Stillinger, C. A. (1989). *Overcoming attributional barriers to the negotiated resolution of conflict*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.
- Varian, H. R. (1984). *Microeconomic analysis*. New York: Norton.
- Winkler, R. L. (1981). Combining probability distributions from dependent informing sources. *Management Science*, *27*, 479-488.
- Wyer, R. S. (1970). Information redundancy, inconsistency, and novelty and their role in impression formation. *Journal of Experimental Social Psychology*, *6*, 111-127.
- Zeckhauser, R. (1971). Combining overlapping information. *Journal of the American Statistical Association*, *66*, 91-92.

RECEIVED: January 24, 1992