
Typicality Can Create, Eliminate, and Reverse the Dilution Effect

Ellen Peters

Myron Rothbart

University of Oregon

Nisbett, Zukier, and Lemley presented evidence for a “dilution effect,” in which information nonpredictive of a stimulus person’s behavior “watered down” or diluted the predictive value of categorical information diagnostic of that behavior. Two experiments suggest, however, that nondiagnostic information influences prediction by altering the perceived goodness of fit between the stimulus person and the diagnostic category. The authors conclude that (a) Nisbett et al. found evidence for dilution because their nondiagnostic items decreased the goodness of fit between the stimulus person and the diagnostic category and (b) depending on the typicality of the nondiagnostic items, it is possible to dilute, enhance, or leave unchanged the predictability between category and behavior. The structure of social categories, and the importance of typicality in the stereotyping process, is discussed.

In 1981, Nisbett, Zukier, and Lemley presented an interesting idea about people’s use of diagnostic and nondiagnostic information in prediction. They showed that the predictive power of diagnostic information can be reduced by the addition of nondiagnostic information, which they called “the dilution effect.” For a given criterion behavior (e.g., “engages in child abuse”), diagnostic items predictive of the behavior (e.g., “has a drinking problem”) and nondiagnostic items irrelevant to the behavior (e.g., “he manages a hardware store”) were developed with pretest samples. Participants who were presented only with diagnostic information gave more extreme predictions than did participants presented with the same diagnostic information plus additional nondiagnostic information. The addition of nondiagnostic information, although independently judged to be of no value in the prediction, diluted the effects of otherwise diagnostic information. An individual with a drinking problem who managed a department store was

judged less likely to be a child abuser than an individual identified only as having a drinking problem.

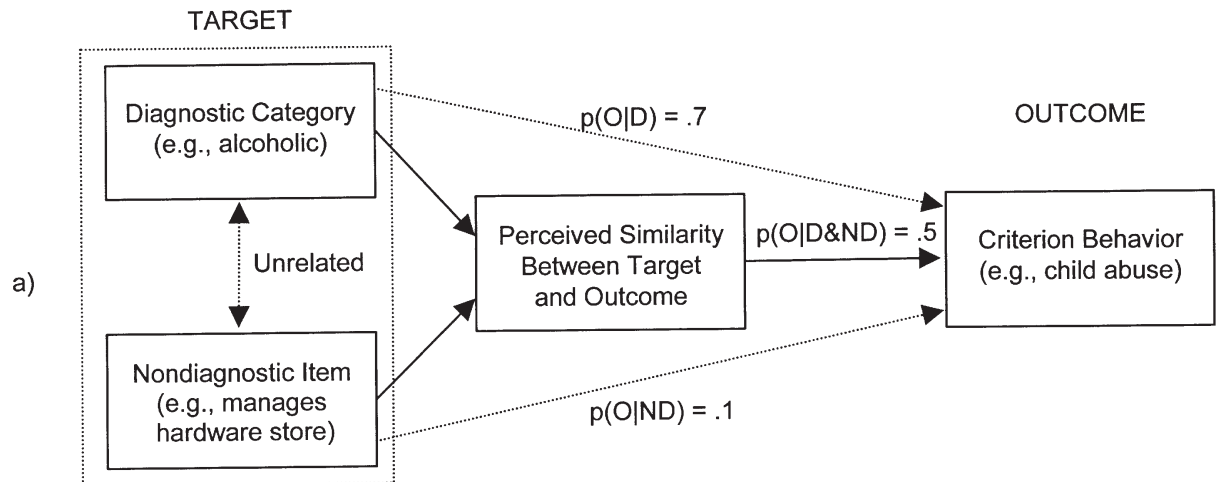
These findings are nonintuitive, nonnormative, and carry important implications for real-life judgments. Information judged to be of little value in predicting an outcome should not alter a prediction. A medical doctor’s diagnosis of her patient’s disease, for example, should not be influenced by patient information unrelated to the disease. A psychologist’s evaluation of the dangerousness of his client should not be reduced by exposure to the unrelated details of his client’s life.

Nisbett et al. (1981) explained the dilution effect in terms of Tversky’s (1977) features of similarity model, in which the presence of nondiagnostic information reduces the similarity between the stimulus person and the target behavior by increasing the number of distinctive elements. According to Tversky’s model, similarity is some function of the difference between common and distinctive features. In the Nisbett et al. paradigm, the overlap between the stimulus person and the target behavior represents the common features and thus the nondiagnostic items add to the number of distinctive features (present for the stimulus person and absent for the target behavior). As the number of nondiagnostic features increase, similarity between stimulus person and target decreases and the stimulus person is viewed as less likely to engage in the target behavior.

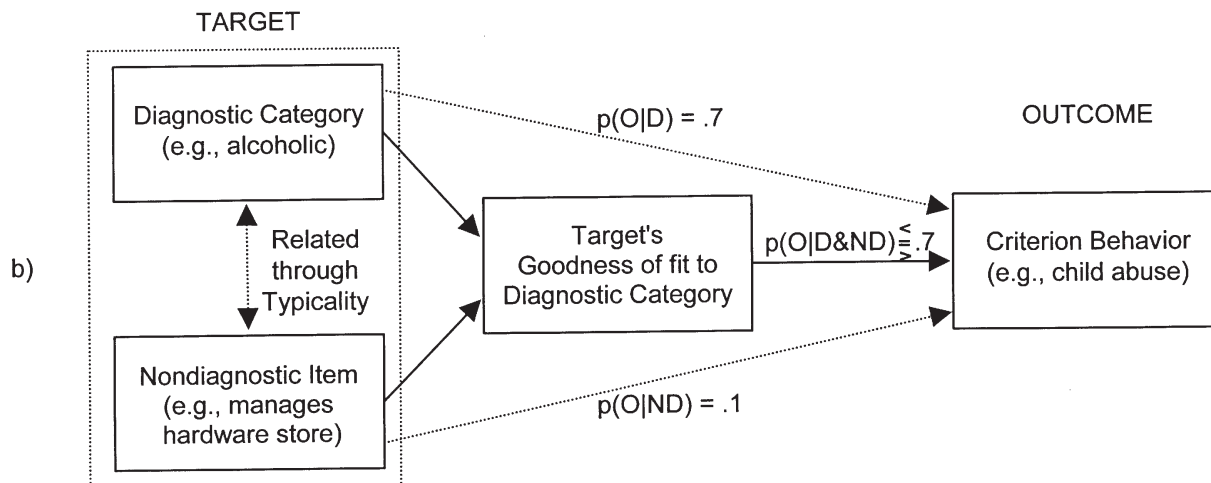
Authors’ Note: This research was supported in part by a National Institute of Mental Health (NIMH) grant (No. MH18935) to the first author and a NIMH grant (No. MH40662) to the second author. We gratefully acknowledge Lew Goldberg, Bertram Malle, Mary Rothbart, and two anonymous reviewers for their extensive comments on earlier drafts of this article. Reprints may be requested from Ellen Peters or Myron Rothbart, Ph.D., Department of Psychology, 1227 University of Oregon, Eugene, OR 97403-1227.

PSPB, Vol. 26 No. 2, February 2000 177-187

© 2000 by the Society for Personality and Social Psychology, Inc.



Nisbett et al.'s Similarity Model for Combining Diagnostic and Nondiagnostic Information



Alternative Model in Which Nondiagnostic Item Can Alter the Strength of the Diagnostic Category

Figure 1 Two models for describing the way in which diagnostic (D) and nondiagnostic (ND) combine to predict an outcome (O): (a) Nisbett, Zukier, & Lemely's (1981) model, in which the addition of ND information reduces the similarity between target and outcome, and (b) an alternative model in which the ND information alters the strength of the D information.

NOTE: $p(O|ND)$ is arbitrarily set at some low number (e.g., .1) to reflect the base rate of the outcome.

We believe that there is another plausible explanation for their findings. When developing their nondiagnostic items, Nisbett et al. (1981) considered only the relation between the nondiagnostic information and the predicted outcome (see Figure 1a). They did not consider, however, that the nondiagnostic information may influence the meaning or interpretation of the diagnostic information. This article considers whether nondiagnostic information may exert its influence on prediction by altering the strength of the diagnostic information. In the early literature on social perception, Asch (1946,

1948) argued persuasively that people will use context when interpreting the meaning of any given piece of information. Different pieces of information do not necessarily combine in a simple additive manner to yield a percept, but can interact in complex ways to influence perception. Although the nondiagnostic items, in isolation, may not predict the criterion, they may change the meaning of the diagnostic information when combined with it (see Figure 1b).

Consider the alcoholic-child abuse example cited earlier. Knowing that Joe has a drinking problem sug-

gests the potential for child abuse to many persons, although they do not know how serious the drinking problem is. Managing a hardware store does not have any obvious relation to child abuse and is therefore considered nondiagnostic. But a person who has a drinking problem and manages a hardware store may be considered less severely alcoholic by virtue of being reliably employed than a person who only has a drinking problem. A less severe alcoholic is less likely to engage in child abuse than a more severe alcoholic. Thus, although nondiagnostic information alone may be unrelated to the criterion, it may serve to weaken the strength of the diagnostic information, resulting in a weaker, or diluted, prediction.

In our view, the Nisbett et al. (1981) research does not constitute a strong test of the dilution hypothesis because the nondiagnostic information may have altered the meaning of the diagnostic information. The stronger test would include nondiagnostic information that is not only unrelated to the criterion, but also does not influence the interpretation of the diagnostic category. Nondiagnostic information unrelated to both the criterion and to the diagnostic category that is still able to weaken the predictive value of the diagnostic information would provide far stronger evidence for a real (and nonnormative) dilution effect.

More generally, the idea that nondiagnostic information may alter the predictive strength of the diagnostic information raises the possibility that nondiagnostic information could be created that would strengthen, rather than weaken, prediction of a criterion. If so, we should be able to demonstrate an enhancement effect—opposite to the dilution effect—using nondiagnostic information that strengthens, rather than weakens, the target's fit with the diagnostic category. Information that increases or decreases a stimulus person's goodness of fit to a category influences that person's typicality, an important determinant of stereotypic judgments (Maurer, Park, & Rothbart, 1995; Rothbart & Lewis, 1988; Rothbart, Sriram, & Davis-Stütt, 1996).

If our hypothesis is correct, it should be possible to show dilution, enhancement, or no effect depending on the typicality of the nondiagnostic items selected. We believe that Nisbett et al. (1981) inadvertently selected a relatively large proportion of atypical nondiagnostic items, making the stimulus person appear to be an atypical category member and thus creating a dilution effect. A different selection of items (i.e., typical items) might have resulted in an article about the "enhancement effect" instead.

Two studies are presented. The first study examined the within-subject partial correlations between an item's typicality (i.e., its tendency to weaken or strengthen the diagnostic category information) and the criterion

behavior, partialing out any direct relation between the item and the criterion. If typicality is important, then the partial correlation should be significantly different from zero. In the second study, participants received either one, three, or five nondiagnostic items that were either typical, irrelevant to, or atypical of the diagnostic category. The effects of increasing the number of nondiagnostic items should be different for the three types of items: (a) for typical items, predictions should increase (a positive slope); (b) for irrelevant items, there should be no change (a zero slope) and (c) for atypical items, predictions should decrease (a negative slope).

For both studies in this article, we chose the stereotype of fraternity members as the diagnostic category and the number of books a university student would read outside of class assignments as the criterion behavior (a typical fraternity member is perceived as reading fewer books than the average student). We then examined the effect of the typicality of added information on the prediction of number of books read, controlling for the diagnosticity of that information.

STUDY 1

Overview

In the first of two sessions, participants rated each of 46 items for typicality (how typical or atypical the item was of fraternity members) and diagnosticity of number of books read. In a second session 2 weeks later, each of the same items were associated with an individual described as a fraternity member, and the same participants were asked to estimate the number of books read by that individual. A partial correlation was computed for each participant across the 46 items, between an item's typicality and the estimated number of books read, partialing out the item's ability to directly predict the number of books read. Because strengthening the stereotype of the fraternity member should result in fewer books read, evidence for a typicality effect would be manifest as a significant negative correlation (tested against zero).

Method

Sixty-seven undergraduate students from introductory psychology courses received partial class credit for their participation. The study was completed in two sessions. In the first session, participants responded to two series of questions (i.e., typicality and diagnosticity) about behavioral items that could describe a fraternity member (see Figure 2 for a partial list of items). The order of presentation of the two series was counterbalanced, and participants completed unrelated tasks in between the two ratings. No significant order effect was found, and order will not be mentioned again.

In the typicality ratings, participants were shown a list of 46 items (e.g., “likes to work on his tan,” “has a pair of Birkenstock shoes”) and were asked to rate each item on a 7-point scale according to its goodness of fit with their image of fraternity members (1 = *very strongly disagrees with my image of a typical fraternity member* to 7 = *very strongly agrees with my image of a typical fraternity member*, with 4 = *unrelated to my image of a typical fraternity member*). In the diagnosticity series, participants rated the same 46 items on another 7-point scale according to its value in predicting the number of books a person would read outside of class compared to the average student (1 = *item predicts person would read very many fewer books* to 7 = *item predicts person would read very many more books*, with 4 = *item predicts person would read neither less nor more books than the average student*).

Participants returned approximately 2 weeks later for the second session, in which they predicted the number of books an individual fraternity member (identified by initials such as R.P.) would read in a year outside of assigned course material. Participants responded to 46 fraternity members described briefly with one of the behavioral items rated in Session 1 (diagnostic category plus one nondiagnostic item) and to 3 individuals described only as “fraternity members” without any nondiagnostic information (used as a baseline measure).

Results

Correlations were calculated within each subject, across 46 items, between the typicality ratings and the book reading estimate, partialing out the items’ diagnosticity of number of books read.^{1, 2} The partial correlations were converted to *z* scores and averaged. The average *z* transformed score was $M = -.25$ ($SD = .17$), which was highly significant when tested against a population value of zero, $t(66) = 12.05$, $p < .001$. Sixty-four out of 67 subjects (96%) showed the negative correlation between an item’s typicality and the estimated number of books read. Fraternity members described with more typical characteristics were predicted to read fewer books outside of class than those described with less typical traits. In this case, a dilution of the fraternity stereotype occurs when the number of books read increases to that of students in general. The correlational result indicates that as the items reduce the typicality of the fraternity member, the number of books read increases, diluting the effect of fraternity membership.

Although this result is consistent with the present hypothesis, the correlation alone does not specify whether the absolute level of books read is higher with atypical items than with no items at all. Figure 2 presents the aggregated data for estimated number of books read plotted against the typicality of the items. This figure includes a horizontal baseline depicting the estimated

number of books read by a fraternity member, with no other information present. Items above this horizontal line indicate dilution of the fraternity stereotype because the estimated number of books read increased toward that of students in general. According to the typicality model, the atypical items on the left side of the abscissa should cause dilution (and be above the horizontal baseline), and the typical items on the right side should lead to enhancement (and be below the baseline). The figure shows that most items had a diluting effect but at least two of the most typical items (e.g., “wears a Greek sweater,” “likes to drink beer”) had the opposite (enhancement) effect of decreasing the number of books read. Thus, although dilution appeared more likely, there was also some evidence of enhancement.

The analysis presented above examined the correlation between an item’s typicality and the criterion, controlling for diagnosticity. It also was possible to examine the correlation between an item’s diagnosticity and the criterion, controlling for typicality. This correlation also was highly significant, with $M = .48$ ($SD = .18$) and $t(66) = 21.20$, $p < .001$, when tested against a population value of zero. This positive correlation indicates that an item can increase or decrease the prediction of books read, independent of any effect of typicality. That is, a person who writes poetry is more likely to read books than a person who likes to drink beer, apart from these items’ stereotypic consistency with fraternity membership.

Discussion

Both the typicality model and the original dilution model received some support in the present study. The dilution model predicts that nondiagnostic information, whether typical or atypical of the diagnostic category, should dilute predictions based solely on category membership. About half of the items showed a dilution rather than an enhancement effect, and about a third of the items demonstrated little or no effect on predictions. However, the negative correlation between typicality and the criterion, partialing out diagnosticity, suggests that dilution may occur in part by decreasing the typicality of the category member. The more atypical (or less typical) a given item was of fraternity members, the more the predicted number of books read approximated that of the average student away from the typical fraternity member. The stronger argument, that highly typical items should show the opposite of dilution—enhancement—received only modest support, with only two items showing enhancement.

There may be two reasons for the failure to find more definitive support for either model, although clearly more dilution than enhancement was found in Study 1. First, there may be a limit as to how much dilution or




Figure 2 Estimated number of books read plotted against a residualized index of typicality (where typicality has been regressed against diagnosticity). NOTE: The negative slope represents the regression line and the horizontal line represents the estimation for a fraternity member with no individuating information.

enhancement is possible with only a single item. In Study 2, we vary the number of nondiagnostic items, similar to the method used by Nisbett et al. (1981). Second, because the average fraternity member was already expected to read fewer than four books in a year, it may have been hard to show the further reduction needed for enhancement due to a floor effect. In Study 2, the criterion variable also was modified slightly, to “number of books read in 4 years,” to increase the overall magnitude of participants’ estimates.

STUDY 2

Overview

Participants judged descriptions of 9 stimulus persons in a 3 (number of items: one, three, five) × 3 (type of items: typical, atypical, irrelevant) within-subjects design. All items were selected to be nondiagnostic of the crite-

riion. Each stimulus person was associated with one, three, or five items that were either all typical, atypical, or irrelevant to the stereotype of fraternity members. We predicted that (a) dilution would be obtained as the number of atypical items increased, (b) enhancement would occur as the number of typical items increased, and (c) no change would result from increasing the number of irrelevant items. Nisbett et al. (1981) would predict the same dilution effect regardless of the typicality of the nondiagnostic items selected.

Method

Psychology students ($N = 178$) completed this experiment in a large prescreening session along with a number of other unrelated studies. Each participant responded to 9 brief descriptions of fraternity members by predicting the number of nonrequired books each individual read during 4 years of college. Descriptions

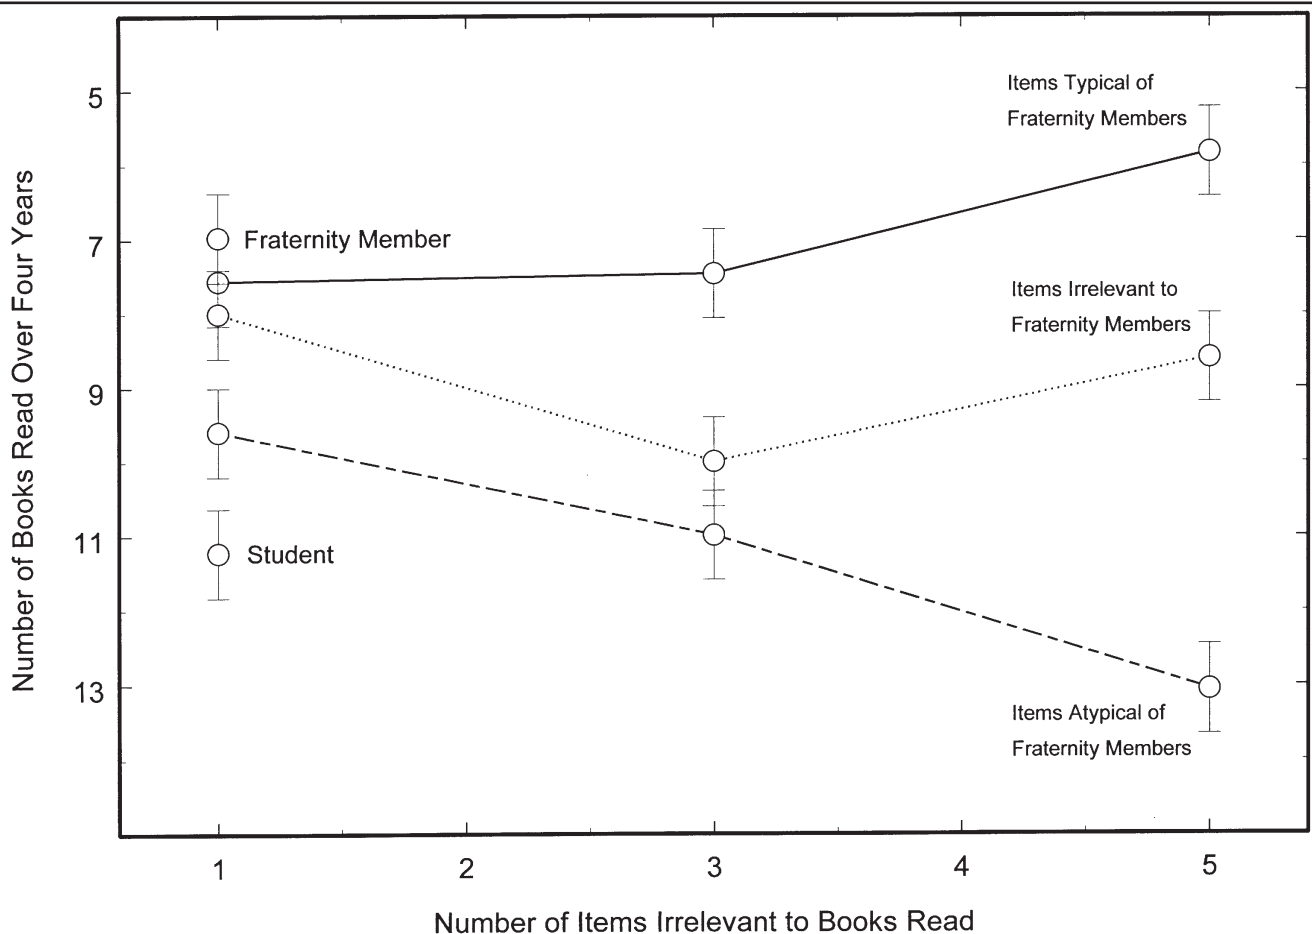


Figure 3 Estimated number of books read as a function of number and type of nondiagnostic items.

were constructed from (a) diagnostic information (i.e., fraternity membership) and (b) nondiagnostic information judged by pretest participants to be typical, atypical, or irrelevant to images of a typical fraternity member. The pretest data came from participants in Study 1 ($N = 67$), and from other undergraduates recruited from introductory psychology courses ($N = 11$).

The mean diagnosticity of all nondiagnostic information was between 3.5 and 4.5 (a rating of 4 on this 7-point scale indicated that the item was irrelevant to the prediction). The nondiagnostic items were selected to be typical, atypical, or irrelevant to pretest participants' image of a fraternity member (pretest participant means on the typicality scale were greater than 4.5 for typical items, less than 3.5 for atypical items, and between 3.5 and 4.5 for nondiagnostic items that were irrelevant to the stereotype).³ Fraternity members were described with one, three, or five of these characteristics in a 3 (typicality of item) \times 3 (number of nondiagnostic items) design. In addition, participants responded to three filler items

containing one, three, or five moderately diagnostic characteristics, as well as two fraternity members and two university students for whom no other information was presented.

Each participant received one of nine versions of the stimulus set. Across the nine forms, each nondiagnostic item appeared equally often within the one, three, and five item conditions. The items were rotated through the nine forms, but the order of question type remained constant (i.e., a fraternity member described with one nondiagnostic irrelevant item always appeared first, one described with five nondiagnostic typical items always appeared third, etc.).

Results

Figure 3 presents the predicted number of books read by a fraternity member as a function of typicality of the nondiagnostic item (typical, irrelevant, or atypical) and number of items (one, three, or five) associated with that fraternity member. Figure 3 also includes the estimated

number of books read by a fraternity member and an average university student with no additional associated information.

The fraternity student was predicted to read 7.0 books on average, whereas the average university student was predicted to read 11.2 books. Thus, dilution of the diagnostic information (fraternity member) is evidenced by an increase in the number of books read. To be consistent with figures from Nisbett et al. (1981), however, we wish to show a dilution effect as a decrease in the criterion variable; for this reason, we have reversed the scale of the ordinate so that a dilution effect (i.e., an increase in books read) appears as a decrease on the graph.

Within each subject, three slopes were computed across the number of items for (a) typical items, (b) irrelevant items, and (c) atypical items. It was predicted that as the number of typical items increased, the criterion estimates would be enhanced and the slope would be positive (recall that enhancement entails a reduction in number of books read, resulting in a positive slope given the reversed scale). As the number of irrelevant items increased, we expected no significant change in the criterion measure (a zero slope), and as the number of atypical items increased, we expected to find the standard dilution effect (a negative slope). The three slopes were subjected to a one-way, within-subjects ANOVA, with typicality of item as the single factor. The average slopes for typical, irrelevant, and atypical items were, respectively, .42, -.15, and -.86. Differences among the slopes were highly significant, $F(2, 354) = 18.62, p < .001$, with all three slopes being significantly different from one another (all t s > 2.94). In addition, the positive slope for the typical items was significantly different from zero, $t(177) = 4.35, p < .001$, the slope for the irrelevant items did not differ significantly from zero, $t(177) = 1.37, ns$, and the negative slope for the atypical items also differed significantly from zero, $t(177) = 3.92, p < .001$.

The descriptive statistics presented in Figure 3, and the inferential statistics presented above, both strongly support the predicted pattern of results: an enhancement effect for typical items, no effect for irrelevant items, and a dilution effect for atypical items.

Although the slopes are quite consistent with the premise that typicality mediates the effects of nondiagnostic information, an inspection of Figure 3 also shows that overall, additional information is more likely to dilute than enhance prediction from a pure category (i.e., a category with no individuating information). Table 1 formalizes this finding by subtracting the book-reading estimates for stimulus persons associated with one, three, or five items from estimates provided for the pure fraternity member. Note that of the nine values presented in Table 1, eight show a positive value (a positive

TABLE 1: The Average Difference Between Predictions for Member Only^a and Predictions for a Fraternity Member With Added Characteristics ($N = 178$)

Characteristics	1	3	5
Typical	0.62	0.52	-1.11*
Irrelevant	1.03	3.50*	1.64*
Atypical	2.67*	4.08*	6.11*

a. The “pure” fraternity member was predicted to read 7.0 nonrequired books in 4 years. All numbers in the table are differences calculated from this baseline. A positive number indicates a dilution effect.

* $p < .01$.

score is evidence for dilution), and among the six values that are significantly different from the pure category baseline, five are in the direction of dilution and one is in the direction of enhancement.

Discussion

At this point, it would be useful to summarize what the data in Figure 3 and Table 1 imply for the typicality model proposed here and for Nisbett et al.’s (1981) similarity model. First, the effects of increasing the amount of nondiagnostic information from one to three to five clearly supports the typicality model. Nisbett et al. found evidence for increasing dilution as the amount of nondiagnostic information increased, and we find evidence of that only for the items that are atypical of fraternity membership. For the irrelevant items, the slope does not show evidence of dilution and does not differ significantly from zero. And for the typical items, we actually obtain the opposite of dilution—enhancement—as evidenced by the significant positive slope. Second, the theoretically predicted enhancement effect, as evidenced by the addition of five typical items, is consistent with the typicality model and incompatible with the similarity model. That is, according to Tversky’s (1977) features of similarity model, adding uncommon, distinctive elements could not possibly increase the similarity between target and outcome, yet that is precisely what we find. Third, the typicality model is also supported by some of the data in Table 1, which shows that the amount of dilution (the magnitude of the positive scores) varies with the typicality of the items: The greatest dilution is found for the atypical items, less for the irrelevant items, and the least for the typical items. Fourth, more dilution than enhancement was found, however, when dilution was measured by the discrepancy between the pure category and category members associated with any individuating information. This last finding is exactly the kind of data that Nisbett et al. (1981) considered as evidence for the similarity model.

Although this last measure of dilution—the difference between the pure and individuated category

members—provides some evidence for the prevalence of dilution, we believe it is not convincing evidence for the similarity model adopted by Nisbett et al. (1981). There are a number of possible explanations for this finding. First, we are still dealing with a bounded scale that is more constrained on the low end than on the high end. The estimated number of books read by a “pure” fraternity member alone ($M = 7.00$) cannot go below zero but it can increase to a very large number. Second, an unadorned category label, such as fraternity member, implicitly may include many typical attributes, making it relatively easy to decrease the strength of the stereotype but difficult to strengthen the category beyond its already extreme level. The nature of category labels, and their implications, is a topic of importance and will be considered further below.

It should be noted, however, that despite the floor effect, the nature of the bounded scale, and the extremeness of pure category labels, we nonetheless find a significant enhancement effect above and beyond the category label—evidence that is clearly inconsistent with Nisbett et al.’s (1981) use of Tversky’s (1977) similarity model as an account of dilution.

GENERAL DISCUSSION

Nisbett et al. (1981) presented evidence that nondiagnostic information can dilute the predictive power of diagnostic information. They invoked Tversky’s (1977) analysis of similarity as the difference between common and distinctive elements to account for the dilution effect, with nondiagnostic information serving to increase the number of distinctive elements and therefore reduce similarity. The results of the present experiments seriously question both the nature and interpretation of the dilution effect. Social categories served as the basis for the diagnostic information, and the participants’ predictions based on these categories will be subject to the general rules of stereotypic inference. One of those rules is typicality or representativeness: The better the goodness of fit between a target person and his or her category, the more likely the category’s attributes will be ascribed to the category member (Hewstone, 1994; Hewstone & Lord, 1998; Maurer et al., 1995; Rips, 1975; Rothbart & Lewis, 1988; Rothbart et al., 1996; Werth & Lord, 1992). Although Nisbett et al. (1981) selected their nondiagnostic items to be irrelevant to the criterion, those same items may have altered the goodness of fit or typicality between the target person and his social category. Study 1 in the present article demonstrated such a correlation between typicality and the criterion, with diagnosticity partialled out. It is possible that the Nisbett et al. research disproportionately drew on nondiagnostic items that reduced the typicality of the target person, inadvertently leading to dilution. If true,

it ought to be possible to add diagnostic information that would increase the typicality of the target person, or even to leave the typicality unaltered. Study 2 succeeded in demonstrating all three possible effects: dilution, no effect, and enhancement. Both studies suggest that nondiagnostic information may alter the predictive power of diagnostic information, not by decreasing the overlap between the target person and the criterion but by altering the meaning and strength of the diagnostic category.

There are several implications that derive from the present findings. First, results from prior research on the dilution effect may be explained by the inadvertent preponderance of atypical nondiagnostic information. The Nisbett et al. (1981) participants, for example, were presented with a great deal of (nomothetically determined) nondiagnostic information, some of which may have been truly irrelevant to the category but at least some of which appear atypical of the category. On page 254 of Nisbett et al. (1981), for example, one stimulus person, John Varner, is described as an engineering major; a sophomore with a 3.1 grade average; from Detroit; whose father is a manager for a steel company; whose mother is a housewife; Catholic; had previously attended Western Michigan; likes to snow ski, play chess, and listen to music; reads mostly science fiction and philosophical literature; and whose last read book was *Cat’s Cradle* by Kurt Vonnegut. Given this complex description, it is difficult to know the exact proportion of nondiagnostic information that confirms, negates, or is irrelevant to the category of engineer. Perhaps reading science fiction and playing chess are consistent with the category, being from Detroit and being Catholic are irrelevant to the category, and reading philosophy and Kurt Vonnegut are incongruent with the category. What we cannot tell from this description of the stimulus person is the relative impact of typical, atypical, and irrelevant information on the participants’ perceptions of the target stimulus person. Simply counting the number of each type of item is not sufficient to predict its impact on perception and judgment—a point we will elaborate on shortly.

The confounding effects of typicality may be present in other research on the dilution effect as well. Zukier (1982), for example, examined the impact of adding nondiagnostic items perceived as either varying widely in his college student population (e.g., some students date a lot, whereas others do not date at all) or as being narrowly dispersed (e.g., all students have about the same number of close friends). Nondiagnostic items drawn from broadly dispersed dimensions (e.g., the number of dating partners in a given year) caused more dilution of a grade point average (GPA) prediction than did nondiagnostic items drawn from narrowly dispersed dimensions. It seems likely, however, that information

perceived as narrowly dispersed will be either irrelevant to or typical of a target's group membership. It is more likely that broadly dispersed information, particularly if drawn from more extreme ends of the distribution, will be atypical of a target's group membership and, as predicted by the present article's model, should produce dilution. Using Zukier's example, a student who scored high on the Scholastic Aptitude Test (SAT) who dated three people during the school year may be a less typical high-SAT student and, thus, may receive a lower GPA prediction. A similar analysis can be made of Zukier and Jennings's (1983-1984) article.

Fein and Hilton (1992) also argued that not all nondiagnostic information is equally diluting and distinguished between nondiagnostic information that was perceived as being broadly useful for most social judgments (pseudorelevant or high in typical diagnosticity) or not useful for most social judgments (irrelevant or low in typical diagnosticity). They found evidence of greater dilution when category members were associated with high, rather than low, typical diagnosticity. It should be noted that their use of the term *typical* differs from ours in a significant way. They are varying the prevalence of attributes that are typical of people in general, whereas we are looking specifically at information that modifies goodness of fit to the stimulus person's social category. Fein and Hilton did not vary the typicality of the information relevant to the category, and it is likely that their pseudorelevant information was atypical of the category. Indeed, they present evidence that individuals associated with pseudorelevant information were judged as less typical or representative of their category. It is possible that their results may have been due, at least in part, to the manipulation of goodness of fit to the category rather than to typical diagnosticity in the population at large.

A second general implication concerns the nature of social category labels, which may make it easier to dilute than to enhance predictions about a category member. Although we do not agree with Nisbett et al.'s (1981) conclusion that nondiagnostic information inevitably dilutes diagnostic information, we nonetheless agree that in general it is probably easier to weaken than to strengthen the image conveyed by a pure category label (cf. Table 1). Lakoff (1987) has argued that categories vary in the nature of their best-fitting exemplars. For some categories, the best-fitting exemplars may well be the average or the modal (most frequent) members, but other categories may be best represented by ideals or paragons. An idealized exemplar (idealized in either a favorable or unfavorable direction) is like a caricature (Mauro & Kubovy, 1992) in which critical features are accentuated. A paragon, in turn, may be an extreme form of an idealized exemplar, in which the best-fitting

member of the category is also the most extreme example of the category.

Research by Judd, Ryan, and Park (1991) has shown that even when thinking about mundane categories, such as engineering and business majors, respondents have an idealized or extreme view of the target group (which is stronger for out-group than for in-group members). Thus, there is reason to believe that for categories such as alcoholics or engineers (two of the examples present in the Nisbett et al. [1981] research), these simple category labels, unadorned with individuating information, may well include some of the more extreme aspects of the category. Additional individuating information, therefore, is more likely to weaken than to strengthen the image evoked by the category label because the pure category already includes many stereotypic elements. When we think of an alcoholic, we are probably less likely to think of a housewife who battles loneliness with vodka, yet manages to keep her children well-fed and clothed. We may be more likely to think of a derelict on the Bowery, with no job or family, sleeping off an alcoholic binge. The former may be statistically more frequent than the latter, but the latter may be closer to the (erroneous) cultural image. The more extreme the cultural image associated with a social category, the more likely it is that an individual's common behavior will decrease, rather than increase, the goodness of fit between himself or herself and the category.

Again, to sharpen the differences between our own approach and that of Nisbett et al. (1981), we interpret the preponderance of dilution over enhancement observed in Table 1 to be due to the relative ease in decreasing the typicality of an image composed largely of typical features rather than to their preferred explanation based on the addition of distinctive features. More generally, Nisbett et al. argued,

Stereotypes may play a somewhat less powerful and pervasive role in person perception than has been assumed by many social psychologists . . . [and] may operate most powerfully in the abstract, applying primarily to undifferentiated groups or individuals, and they may exert relatively little impact on judgments about concrete, individuated targets. (p. 272)

In contrast, our own view is that the process of stereotyping is directly implicated in the dilution effect and that it is possible to individuate targets in a way that enhances, as well as dilutes, stereotypic judgments.

A third implication is to repeat, once again, the importance of Asch's argument that meaning is derived from context: Participants will combine information in a way that maximizes apparent coherence and meaningfulness. Given participants' sensitivity to typicality information, and their tendency to combine information in a

way to achieve meaning, it is not surprising to find that irrelevant information can influence prediction by altering the meaning of the relevant information. Although Nisbett et al. (1981) portrayed their participants' reasoning as nonnormative, that description does not seem applicable to the present participants, who used the nondiagnostic information to interpret the meaning of the diagnostic information. Although there is error in

the extremeness of our cultural images, the use of a representative heuristic or goodness of fit principle is not in itself evidence of irrationality. More generally, although as experimenters we often create stimulus presentations by the mechanical combination of elementary components, we should not forget that participants will attempt to assess the meaningful relations among those components (Asch, 1948).

APPENDIX

Nondiagnostic and Diagnostic Items Used in Study 2

<i>Nondiagnostic Items</i>			
<i>Atypical</i>	<i>Irrelevant</i>	<i>Typical</i>	<i>Moderately Diagnostic Items</i>
Has a full beard	Enjoys Italian food	Likes rock music	Likes beer
Wears bellbottoms	Has brown hair	Is talkative	Is a physics major
Rides the bus	Likes to jog	Wears a baseball cap	Plays chess
Is celibate	Juggles	Goes downhill skiing	Rarely skips class
Works part-time in a flower shop	Has a pair of Birkenstocks	Plays pool	Reads the school newspaper
Enjoys country line dancing	Voted for Clinton in the 1996 election	Wears contacts instead of glasses	Is planning on joining the Peace Corps
Does pottery in his spare time	Is from Portland, Oregon	Has short hair	Often skips class
Does not like parties	Likes Pink Floyd	Dresses fashionably	Is religious
Is overweight	Prefers McDonald's to Burger King	Is extroverted	

NOTE: Only nondiagnostic items were analyzed for Study 2. Moderately diagnostic items were included only as fillers to provide context for the nondiagnostic items.

NOTES

1. Study 1 uses an idiographic rather than a nomothetic estimate of nondiagnosticity. In the original Nisbett, Zukier, and Lemley (1981) research, and in the present Study 2, the nondiagnosticity of an item was determined by aggregated ratings across pretest judges. Because unanimous agreement among judges is rare, aggregated ratings are less desirable than using each judge to standardize his or her own items. Both idiographic and nomothetic methods are used in the present research.

2. The average zero-order correlations among the indices were as follows: (a) $M = -.39$ for typicality and the criterion, (b) $M = .55$ for diagnosticity and the criterion, and (c) $M = -.35$ for typicality and diagnosticity. All means were significantly different from zero (for all means, $t > 10$ and $p < .001$).

3. See the appendix for the complete list of items used in Study 2.

REFERENCES

- Asch, S. E. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology, 41*, 258-290.
- Asch, S. E. (1948). The doctrine of suggestion, prestige, and imitation in social psychology. *Psychological Review, 55*, 250-276.
- Fein, S., & Hilton, J. L. (1992). Attitudes toward groups and behavioral intentions toward individual group members: The impact of nondiagnostic information. *Journal of Experimental Social Psychology, 28*, 101-124.
- Hewstone, M. (1994). Changing outgroup perceptions: A cognitive intergroup perspective. *International Review of Social Psychology, 7*, 45-62.
- Hewstone, M., & Lord, C. G. (1998). Changing intergroup cognitions and intergroup behavior: The role of typicality. In C. Sedikides, J. Schopler, & C. A. Insko (Eds.), *Intergroup cognition and intergroup behavior*. Mahwah, NJ: Lawrence Erlbaum.
- Judd, C. M., Ryan, C. S., & Park, B. (1991). Accuracy in the judgment of in-group and out-group variability. *Journal of Personality and Social Psychology, 61*, 366-379.
- Lakoff, G. (1987). Cognitive models and prototype theory. In U. Neisser (Ed.), *Concepts and conceptual development: ecological and intellectual factors in categorization* (pp. 63-100). Cambridge, MA: Cambridge University Press.
- Maurer, K. L., Park, B., & Rothbart, M. (1995). Subtyping versus subgrouping processes in stereotype representation. *Journal of Personality and Social Psychology, 69*, 812-824.
- Mauro, R., & Kubovy, M. (1992). Caricature and face recognition. *Memory and Cognition, 20*, 433-440.
- Nisbett, R. E., Zukier, H., & Lemley, R. E. (1981). The dilution effect: Nondiagnostic information weakens the implications of diagnostic information. *Cognitive Psychology, 13*, 248-277.
- Rips, L. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior, 14*, 665-681.
- Rothbart, M., & Lewis, S. (1988). Inferring category attributes from exemplar attributes: Geometric shapes and social categories. *Journal of Personality and Social Psychology, 55*, 861-872.
- Rothbart, M., Sriram, N., & Davis-Stütt, C. (1996). The retrieval of typical and atypical category members. *Journal of Experimental Social Psychology, 32*, 309-336.
- Tversky, A. (1977). Features of similarity. *Psychological Review, 84*, 327-352.

Werth, J. L., & Lord, C. G. (1992). Previous conceptions of the typical group member and the contact hypothesis. *Basic and Applied Social Psychology, 13*, 351-369.

Zukier, H. (1982). The dilution effect: The role of the correlation and the dispersion of predictor variables in the use of nondiagnostic information. *Journal of Personality and Social Psychology, 43*, 1163-1174.

Zukier, H., & Jennings, D. (1983-1984). Nondiagnosticity and typicality effects in prediction. *Social Cognition, 2*, 187-198.

Received June 12, 1998

Revision accepted October 27, 1998