Memory & Cognition
2002, 30 (3), 456-468

Alignable and nonalignable differences in causal explanations

ANN L. McGILL
University of Chicago, Chicago, Illinois

Prior research indicates that people may base their causal explanations on distinctive features between an event and a contrasting background instance in which the event did not occur. Research on similarity judgments suggests that there are two types of distinctive features: alignable differences, which are corresponding characteristics of a pair, and nonalignable differences, which are characteristics of one item for which there are no corresponding characteristics in the other. In three experiments, the hypothesis that people's evaluations of causal explanations vary as a function of feature alignment was examined. The results suggest that people will rate explanations differently on the basis of alignable or nonalignable differences, depending on the type of the event, and that alignability depends on the relational structure among the features of the event.

Prior research on causal reasoning suggests that people base their causal explanations on distinctive features between an event and a contrasting background instance in which the event did not occur (e.g., Einhorn & Hogarth, 1986; Hilton & Slugoski, 1986; Kahneman & Miller, 1986; McGill, 1989). For example, a school administrator may explain excellent standardized test scores at one school by comparing the target school with another school in the same area that performed worse. Possible explanations for the difference in performance would be based on other differences between the two schools—for example, the ratio of students to teachers, neighborhood type (e.g., urban vs. suburban), or the school's having been "adopted" by the school of education at a nearby university.

Research on judgments of similarity and analogical reasoning has identified two types of distinctive features: alignable differences, which are corresponding characteristics of a pair, and nonalignable differences, which are characteristics of one item for which there are no corresponding characteristics in the other (Goldstone, 1994; Markman & Gentner, 1993a, 1993b; Medin, Goldstone, & Gentner, 1993; see also Sanbonmatsu, Kardes, & Gibson, 1991, on types of difference features). Alignable differences are characteristics that can be "matched up" between two items. Alignable differences may occur along a continuous dimension. For example, the difference in the ratio of students to teachers between the two schools is an alignable difference. Alignable differences may also be binary characteristics that occupy corresponding roles in people's representations of the items (Markman & Gentner, 1996). For example, differences in neighborhood type are also alignable differences. Nonalignable differences are characteristics of one member of the pair that cannot be matched up with characteristics in the other. For example, one school's having been adopted by a university's school of education is a nonalignable difference.[1]

Previous research has demonstrated effects for feature alignment on similarity judgments and choice. The purpose of the present research was to evaluate whether people's causal judgments are also sensitive to feature alignment. In particular, in this research I examined whether causal explanations based on alignable differences would be rated higher than explanations based on nonalignable differences, as would be consistent with prior work on comparison processes, or whether relative ratings of alignable and nonalignable differences would depend on the nature of the event, as is suggested by work on causal judgment (Einhorn & Hogarth, 1986). In addition, in this research I tested the hypothesis that a feature may be perceived as alignable or nonalignable, depending on the perceived relational structure among features of an event, and that, as a consequence, people may base their explanations on different features, depending on how they structure their representations.

Alignable and Nonalignable Differences as Bases for Causal Explanations

The literature suggests two competing predictions regarding the evaluation of alignable and nonalignable differences in causal explanations. Research on comparison judgments suggests that alignable differences are more salient in the comparison process than are nonalignable differences (Gentner & Markman, 1994; Markman & Gentner, 1993a). Supporting this view is evidence that participants give greater weight to alignable differences than to nonalignable differences in similarity judgments (Markman & Gentner, 1996) and in choice (Lindemann & Markman, 1996; Markman & Medin, 1995; Slovic &

Correspondence concerning this article should be addressed to A. L. McGill, Graduate School of Business, University of Chicago, 1101 E. 58th Street, Chicago, IL 60637 (e-mail: ann.mcgill@gsb.uchicago.edu).

—Accepted by previous editorial team

Macphillamy, 1974; Zhang & Markman, 1998), that alignable differences are listed more frequently in comparison judgments than are nonalignable differences (Gentner & Markman, 1994; Markman & Gentner, 1993a), and that alignable differences are more likely to enter into justifications of decisions (Markman & Medin, 1995).

Two explanations have been proposed for the greater weight given to alignable differences in comparison judgments (Lindemann & Markman, 1996). First, alignable differences may have an attentional advantage. The comparison process, particularly that involved in similarity judgments, appears to focus more on the ways that things are alike than on the ways that they are different. Because alignable differences have a common element, they are therefore given more attention than are nonalignable differences. Second, alignable differences may be more easily evaluated than nonalignable differences (Hsee, 1996; Hsee, Blount, Loewenstein, & Bazerman, 1999). Nonalignable differences must be evaluated independently, and as a consequence, their relative size or meaning may be difficult to interpret. By contrast, alignable differences, by definition, may be evaluated in reference to the other item in the pair, which makes it easier to interpret their relative size or meaning. The more easily evaluated alignable differences may, therefore, have an advantage in the comparison process. The implication of these findings is that people will be more likely to base their causal explanations on alignable differences than on nonalignable differences when asked to explain an event. For example, people would be more likely to base their explanations for superior test performance at one school, as compared with another, on differences in the ratio of students to teachers or on differences in neighborhood type (i.e., alignable differences) than on one school's having been adopted by the nearby university (i.e., a nonalignable difference).

Research on judgments of causation suggests, however, that basing explanations on alignable or nonalignable differences may depend on characteristics of the event. Einhorn and Hogarth (1986) identified perceived similarity between a proposed cause and the effect as a commonly used "cue to causality." That is, in judging probable cause, people consider the similarity of the cause and the effect. Similarity may be judged in terms of physical resemblance—for example, when people consider the degree of similarity between a proposed cure for an illness and the symptoms of the illness (Nisbett & Ross, 1980). Similarity may also be judged in terms of length and strength, called "congruity" by Einhorn and Hogarth (1986, p. 10)—for example, when people expect large effects (e.g., the assassination of a president) to have large causes (e.g., a broad-based conspiracy).

The concept of congruity in judgments of probable cause may be extended to the evaluation of alignable and nonalignable differences. Specifically, congruity, as identified by Einhorn and Hogarth (1986), implies that people compare the *degree* to which an outcome and a possible causal candidate differ from the comparison case. The present research posits that people may also consider *how* outcomes and candidates differ from the comparison case. That is, the present research posits that people consider whether changes represent the addition (or deletion) of features from the comparison case or whether they represent modification of existing features. People are proposed to perceive an outcome that represents an addition to the comparison case (i.e., an outcome that has "come out of the blue") to be the result of some other addition to the comparison case (i.e., a cause that has also "come out of the blue"). For example, the first-time award of a federal grant may appear to be the result of the school's having been adopted that same year by the local school of education. By contrast, an outcome that represents a change in one characteristic of the comparison case may seem better explained with reference to a change in another characteristic in the reference case. For example, improved test scores from one year to the next may seem better explained by a shift to smaller class sizes over the same period. Reliance on congruity as a cue to causality suggests therefore that people will rate nonalignable differences as better explanations for nonalignable events but alignable differences as better explanations for alignable events.

In sum, the literature on comparison judgments seems to suggest that people will base their causal explanations for events on alignable differences more than on nonalignable differences, regardless of the type of event. By contrast, research on judgments of probable cause suggests that people will base their causal explanations on alignable differences for alignable outcomes and on nonalignable differences for nonalignable outcomes. Experiment 1 was intended to evaluate these competing predictions regarding the role of alignable and nonalignable differences in causal judgment.

EXPERIMENT 1

Method

Participants. Ninety-three members of the University of Chicago community were recruited through signs posted on campus. The participants were offered $5 to participate in a half-hour experiment on "decision making."

Stimuli and Procedure. The participants were presented with two scenarios in a booklet of materials, of which the present study was the second of several unrelated experiments. An alignable outcome version and a nonalignable outcome version were created for each scenario. The alignable outcome version of each scenario was created by describing an event that differed from the comparison case as a matter of degree along a single dimension. The nonalignable version was created by describing an event in binary terms—that is, as having occurred in the target case, but not in the comparison case. The participants rated an explanation, which was presented as an alignable or a nonalignable difference. Thus, the alignable and nonalignable versions of the scenarios described different events, but the participants rated the same explanation, which was presented as an alignable or a nonalignable difference, depending on experimental condition.

For example, in the alignable outcome version of the *store scenario*, the participants were asked to rate explanations for why an employee at a CD store was so much more effective than another employee in handling customers (as shown in italics), whereas in the nonalignable outcome version, the participants rated why one employee had decided to buy into the business and become a co-owner whereas the other had not (as shown in bold print):

> Barb Thomas owns several independent music stores that sell new and used CDs, mostly used. An important element to success in retail businesses like Barb's is being able to find, train, and retain good employees. Barb was thinking about this issue with two particular employees in mind. She hired Mike and Carl at the same time and they seemed like very similar young men in temperament and background. *They also were very much alike in how they interacted with customers when they first started out. Mike, however, has turned out to be amazingly good at dealing with customers. He is consistently personable even when other things in his life make that hard and he is able to diffuse situations when customers become irate. Carl, by contrast, is never really rude but he can be highly intolerant of people and a bit surly. Barb wonders what made such a difference*./**They were also very much alike in their ambitions when they first started out. Mike, however, has recently approached Barb about buying into the business to be a co-owner of one of the stores. He has saved some money and has worked out a possible plan for buying in over time. Carl, by contrast, continues as an hourly employee. Barb wonders what made such a difference**.

Information on the two employees followed in two paragraphs, with information on the target employee always presented first. The event was reinforced in each paragraph with a clause following each employee's name, as is shown in the descriptions of the employees below.

Feature alignability was manipulated by following the procedure of Markman and Medin (1995; see also Zhang & Markman, 1998). Specifically, the target feature *having worked with the owner* was presented as an alignable difference by describing the different levels of this feature for the target and the background employees. This feature was presented as a nonalignable difference by omitting information about this feature—that is, by deleting the sentence that is shown below in italics for the background employee:

> Mike, who is very good with customers, is 23 years old. He finished high school and had no real interest in going to college. His primary interest is music and he has very broad tastes. Mike works a forty-hour week, mostly Monday through Friday. He mostly works with a part-time person as an assistant. For his first year, he worked alongside the owner much of the time. He schedules his own hours.

> Carl, who is not as good with customers, is 24 years old. He too finished high school and had no real interest in going to college. His primary interest is music and he has very specific tastes confined to one or two styles. Carl works a forty-hour week, some weekdays and some weekends. He mostly works with one or two part-time people as assistants. *For his first year, he worked occasionally with the owner but not frequently*. He schedules his own hours.

The participants were then asked to rate three explanations for the event on a scale of 1 (*not at all related*) to 7 (*strongly related*). The second explanation was based on the target feature, whereas the other explanations were presented as filler items and were not analyzed. For example, for the continuous event version, the participants were asked to rate "how likely do you think it is that Mike's being good with customers is related to his" (1) "having broad musical tastes," (2) "having worked with the owner," and (3) "working during the week."

The *college scenario* was similarly constructed. The alignable event version asked the participants to rate explanations for why one student at a junior college graduated with a much higher grade point average than did another. The nonalignable event version focused on why one student had joined a fraternity whereas the other had not. The target feature, being ambitious, was manipulated as alignable or nonalignable in the same way as in the store scenario.

Design. The experimental design therefore contained two between-subjects factors, event type (alignable, nonalignable) and difference type (alignable, nonalignable), which were crossed to created four (2 × 2) versions of each scenario. The participants were randomly assigned to one of these four conditions and were presented the same versions of the store and college scenarios. Order of presentation of the scenarios was counterbalanced.

Pilot study. A pilot study was conducted to examine perceptions of events as *alignable*, which was operationalized in this study as differing in degree, and *nonalignable*, which was operationalized as being of a binary nature. For the pilot study, 18 people were recruited from the University of Chicago community in the same manner as that for the main study. The participants in the pilot study were presented with the following instructions:

> A list of events is presented below. Please consider each event and rate whether it is the sort of event that happens by degrees (i.e., ranges from low to high or from one end of a continuum to another) or whether it is a binary event (i.e., is a yes-no/on-off type of event that either happens or it doesn't).

The pilot study participants then rated six features "for a student at a junior college" and six features "for an employee in a CD store." Ratings were provided on a scale of 1 to 7, with endpoints labeled *binary* and *by degrees*, respectively. Hence, higher ratings indicated that the event was perceived more as one that happens by degree than as one that is of a binary nature. As is consistent with manipulation of event type, the alignable outcome for the store scenario *being polite with customers* ($M = 5.28$) was rated significantly higher than was the nonalignable outcome *deciding to buy in to be co-owner* [$M = 2.06$; paired $t(17) = 7.46$, $p < .001$], and the alignable outcome for the college scenario *getting good grades* ($M = 5.17$) was rated significantly higher than was the nonalignable outcome *joining a fraternity* [$M = 2.11$; paired $t(17) = 8.65$, $p < .001$].

Results

The participants' ratings of the target features were analyzed with a repeated measures analysis of variance (ANOVA), with scenario (store, college) as a within-subjects factor and event type (alignable, nonalignable) and difference type (alignable, nonalignable) as between-subjects factors. The results indicated no significant main effects or interactions for scenario (Fs < 1). Ratings were therefore analyzed for the store and college scenarios combined (see Table 1, which shows ratings of the target features averaged over the two scenarios by experimental condition). Analysis of the between-subjects factors indicated no significant main effect for difference type ($F < 1$) and a significant main effect of event type [$F(1,89) = 6.18$, $p < .05$]. The participants rated explanations for the alignable event higher than explana-

Table 1
Ratings of Explanations Based on Alignable and Nonalignable Differences for Alignable and Nonalignable Events in Experiment 1

	Event Type	
Explanation Base	Alignable	Nonalignable
Alignable differences		
Mean	6.13	4.72
Standard deviation	0.76	0.88
Number of participants	24	22
Nonalignable differences		
Mean	5.29	5.57
Standard deviation	1.16	1.45
Number of participants	24	23

tions for the nonalignable event. In addition, an analysis revealed a significant event type × difference type interaction [$F(1,89) = 13.66$, $p < .001$], indicating that the participants' ratings of the alignable and nonalignable differences depended on the type of event.

As was predicted, specific contrasts indicated that the participants rated the alignable difference higher than the nonalignable difference for the alignable event [$F(1,89) = 7.02$, $p < .01$] but rated the nonalignable difference higher than the alignable difference for the nonalignable event [$F(1,89) = 6.65$, $p < .05$]. Ratings of the alignable difference were significantly higher for the alignable event than for the nonalignable event [$F(1,89) = 18.89$, $p < .001$], whereas ratings for the nonalignable difference did not differ for the alignable and the nonalignable events ($F < 1$).

Discussion

The participants in Experiment 1 provided different ratings for explanations, depending on whether the explanations were based on alignable differences or nonalignable differences, a result that is consistent with prior research on similarity judgments and choice that has shown people's judgments to be differentially sensitive to these types of differences. In contrast to prior work on similarity judgments and choice, however, the participants in the present study did not rate alignable differences higher than nonalignable differences in all conditions. Instead, ratings of alignable and nonalignable differences depended on the type of event, a result that is consistent with the hypothesis that people consider congruity between difference type and event type to be a cue to causality.

These findings suggest that alignable differences may not have the attentional advantage in causal judgments that has been posited for similarity judgments. Nonalignable differences may be more salient in people's causal judgments than in their comparison judgments because the focus of causal judgments is on differences, whereas in comparison judgments it is on commonalities. It is also possible that alignable differences may not be more easily evaluated than nonalignable differences in causal judgments. In similarity judgments and choice, people assess the size or meaning of differences, which may make alignable differences easier to evaluate than nonalignable differences. In causal judgments, by contrast, people evaluate possible explanations by comparing the type of difference with type of event in an effort to match alignable differences with alignable events and nonalignable differences with nonalignable events. Alignable differences do not appear to have an evaluability advantage in this matching process. Hence, the two mechanisms that have been posited to give alignable differences an advantage in other comparison judgments, attentional salience and ease of evaluation, may not have been relevant to the present causal judgments.

Although the findings of Experiment 1 support the hypothesis that people attend to feature alignment in judging probable causes, the results do not indicate what it is about alignment that makes people want to match alignable events with alignable causes and nonalignable events with nonalignable causes. One possibility that was raised in the introduction is that the congruity effect observed in Experiment 1 reflects how people judge and perceive differences in representations. Specifically, people may be sensitive to the way in which the representation of the comparison case would have to be modified to achieve a representation of the target episode. In this view, people prefer to explain outcomes that require the addition of a feature in their representation with a cause that also requires the addition of a feature, whereas they prefer to explain a change in level of a preexisting feature with a change in another preexisting feature, thereby achieving a sort of balanced-entry system of representation.

An alternative account for the findings in Experiment 1 is that the participants judged congruity of cause and effect in terms of continuity–discreteness, and not in terms of feature alignment per se. Feature alignment is a distinct construct relative to continuity–discreteness in theory, but in practice they are operationalized in very similar ways. For example, nonalignable differences were operationalized in the first experiment, as in prior research, by omitting information about the target feature in the comparison case (cf. Markman & Medin, 1995; Zhang & Markman, 1998). It was assumed that the omission of information about the target feature in the comparison case would produce a representational difference consistent with the definition of alignment (i.e., the target feature would be included in the representation of the event, but no corresponding "slot" for this feature would be created in the representation of the comparison case). On the other hand, the participants may have included a slot for the missing feature in their representation of the comparison and, as such, perceived the target feature as a binary difference between the event and the comparison, present in one but absent in the other. Hence, the congruity effect observed in Experiment 1 may not reflect representational differences between alignable and nonalignable differences but may, instead, derive from people's belief that the productive force of the cause (a continuous change vs. a binary change) should be matched by the nature of the outcome (a continuous change vs. a binary change). Experiment 2 was conducted to examine this alternative explanation for the findings of Experiment 1.

EXPERIMENT 2

Method

Participants, Stimuli, and Procedure. Ninety-one members of the University of Chicago community were recruited in the same way as in Experiment 1. As in the earlier study, the participants were presented with two scenarios as part of a larger packet of materials for several unrelated studies. A continuous outcome version and a binary outcome version were created for each scenario. In contrast to the stimuli for Experiment 1, the continuous and binary outcome versions described approximately the same event. For example, in the *store scenario*, which was adapted from Experiment 1, the outcome involved an employee's decision to buy into the stores. For the continuous version, the target employee bought in at a higher level than the comparison employee (shown in italics below), but for the binary version, the target employee opted to buy in, whereas the contrasting employee did not (shown in bold below):

> Barb Thomas owns several independent music stores that sell new and used CDs, mostly used. An important element of success in retail businesses like Barb's is being able to find, train, and retain good employees. Barb was thinking about this issue with two particular employees in mind. She hired Mike and Carl at the same time and they seemed like very similar young men when they first started out. *Mike, however, has recently taken great advantage of the employee ownership program, which allows employees to buy into the business as the number of stores expands. Mike has decided to buy into one of the stores at an extremely high level, making him a major co-owner. Carl, by contrast, has taken little advantage of the employee ownership program. He has bought in at a very low level, making him a very minor co-owner*. **Mike, however, has recently approached Barb about buying into the business to be a co-owner of one the stores. Carl, by contrast, continues as an hourly employee and has no plans to buy into the business**. Barb wonders what made such a difference.

The possible causal candidate was also described as a continuous or a binary difference. For example, for the binary difference version of the store scenario, the participants were told of the target employee that "for his first year, he worked with the owner, sharing the same work schedule," whereas for the comparison employee, "during his first year, he did not work with the owner, having different work schedules." For the continuous difference version, the information for the target employee was that "for his first year, he worked alongside the owner much of the time, averaging almost forty hours per week," but the information for the comparison was that "during his first year, he worked infrequently alongside the owner, averaging no more than four hours per week."

The *dog scenario* described two dogs, one of which experienced emotional distress when boarded in kennels. For the binary version, the participants were told that the target dog had developed "kennel anxiety," which prevented him from being boarded anywhere, but that the comparison dog did not have kennel anxiety. For the continuous outcome version, the target dog was described as extremely nervous in kennels, whereas the comparison dog was not very nervous. In the binary cause version, participants were told that the target dog had a yard to play in but that the comparison dog had no yard. In the continuous cause version, the target dog had a large yard, whereas the comparison dog had an extremely small yard.

As in Experiment 1, the participants were asked to rate three explanations, the target feature and two filler questions that were not analyzed, on a scale of 1 to 7. For both the binary and the continuous outcome versions of the store scenario, the participants were asked to rate explanations for "Mike's decision to buy into the business." The target feature for the binary cause version was "having worked with the owner," and that for the continuous cause was "having worked frequently with the owner." For the dog scenario, the participants were asked to rate explanations for the dog's "having kennel anxiety" in the binary outcome version and for "being extremely nervous in kennels" for the continuous outcome version. The participants rated either the binary cause "having had a yard to play in" or the continuous cause "having had a large yard to play in."

Design. The experimental design contained two between-subjects factors, event type (continuous, binary) and cause type (continuous, binary), which were crossed to create four (2 × 2) versions of each scenario. The participants were randomly assigned to conditions and were presented the same version of the store and dog scenarios. Order of presentation of the scenarios was counterbalanced.

Results

The participants' ratings of the target feature were analyzed with a repeated measures ANOVA, with scenario (store, dog) as a within-subjects factor and event type (continuous, binary) and cause type (continuous, binary) as between-subjects factors. The results indicated a significant main effect of scenario, with the participants providing higher ratings for the dog scenario than for the store scenario [$F(1,87) = 4.60$, $p < .05$], but no significant interactions of scenario with the other variables (Fs < 1). The results were therefore analyzed for the store and dog scenarios combined (see Table 2, which displays ratings of the target feature averaged for the two scenarios by experimental condition). Analysis of the between-subjects factors indicated no significant main effect of cause type [$F(1,87) = 2.34$, $p < .15$] and a significant main effect of event type [$F(1,87) = 8.37$, $p < .01$]. The participants provided higher ratings of explanations for the binary event than for the continuous event. In addition, an analysis revealed a significant event type × cause type interaction [$F(1,87) = 11.88$, $p < .001$].

Specific contrast analysis revealed that the participants rated the continuous cause higher than the binary cause for the continuous event [$F(1,87) = 11.97$, $p < .01$]. Furthermore, the participants rated the binary cause higher than the continuous cause for the binary event, although this difference was not significant [$F(1,87) = 2.25$, $p < .15$]. Ratings of continuous causes did not differ between event types ($F < 1$), whereas ratings of binary causes were higher for binary events than for continuous events [$F(1,87) = 20.36$, $p < .001$].

Discussion

The results of Experiment 1 suggested that people use congruity as a cue to causation by comparing the nature of the outcome with the nature of the cause. The purpose of the second experiment was to evaluate whether the congruity effect observed in the first experiment might be due to feature continuity, and not necessarily to feature alignment. That is, the second experiment was intended to test the hypothesis that people attempt to match continuous (binary) events with continuous (binary) causes. Findings support this hypothesis, suggesting that people evaluate the congruity between the "motion" of the cause and that of the event. That is, people perceive a change in degree in an outcome variable as resulting from a similar change in a causal variable—for example, when an increase in temperature is caused by the turn of a ther-

Table 2
Ratings of Explanations Based on Continuous and Binary Differences for Continuous and Binary Events in Experiment 2

	Event Type	
Explanation Base	Continuous	Binary
Continuous differences		
Mean	5.88	5.77
Standard deviation	0.85	1.01
Number of participants	21	24
Binary differences		
Mean	4.89	6.15
Standard deviation	1.16	0.69
Number of participants	23	23

mostat dial. By contrast, a binary outcome appears to result from a binary cause—for example, when an explosion is caused by the push of a button.

The results of the second experiment suggest that the congruity cue operates around feature continuity. The results did not rule out the possibility, however, that other aspects of alignment—particularly, its representational component—may also be relevant in judgments of causation. The methodology employed in Experiments 1 and 2 is not, however, suited to isolating the effect of alignment per se from that of continuity. In these two experiments, the participants were asked to compare two cases that belong to the same nominal category—for example, two young employees, two dogs of the same breed, or two students. Furthermore, the cases were described in simple terms with few features. Under these conditions, alignment is readily confounded with continuity as the participants recast features that were intended to differ in representation (i.e., having a slot in one representation, but no slot in the other) as binary differences (i.e., present for one case but absent for another). I therefore conducted a third experiment to examine the effect of alignment on causal judgment in a different manner.

Structure Mapping and Alignability

One aspect of alignability relates to the presence or absence of corresponding "slots" in the representations of compared items. Recent research indicates that alignability is related not just to the presence or absence of feature information, but also to the relational structure of features within each pair. That is, items are perceived not just in terms of the features they possess, but also in terms of the relationship among the features, and comparison involves aligning objects in terms of the structured relations among the features.

More formally, in recent research on similarity judgments, how people represent objects and the process by which they compare them has been examined (see Markman & Gentner, 1996, for a more detailed discussion of these concepts). This research suggests that psychological representations may be characterized by structured, hierarchical systems, which include objects, attributes of objects, relations among attributes, and relations among relations (Gentner, 1983, 1989; Gentner & Markman, 1994; Markman & Gentner, 1993a, 1993b, 1996). For example, in Figure 1, which is adapted from Markman and Gentner (1993a), people may represent the scene in the first frame as containing the attributes *robot arm* and *car*, but also the relations *beside* and *repairs* between these two attributes.

In this view, comparison involves mapping structure from one object to another. As Markman and Gentner (1996) explain, "the main idea is that the comparison of complex structures involves not only the matching of features, but also the determination of which features and objects play the same role in the relational system" (p. 236). Hence, in comparing the first panel with the second in Figure 1, people may note that, in the first panel (i.e., Jack's repair shop), the robot repairs the car but that, in the second panel (i.e., Grant's repair shop), the man repairs the robot. This relational match would render the robot in panel 1 and the man in panel 2 alignable differences, because they are alignable elements in a common relational structure. They differ because one is mechanical, whereas the other is human, but both make repairs at their respective shops. Although the robot in panel 1 is physically similar to the robot in panel 2, these items would not be perceived as common features or even alignable differences, because they do not occupy the same role in the relational structure.

Hence, according to this structural alignment view, people do not represent objects as lists of features that they compare by noting features in one that are or are not present in the other (cf. Tversky, 1977). Instead, people represent objects as structured relations and compare objects along the lines of the shared structure. This structure-mapping process appears to be central to similarity judgment and is also at the core of analogical reasoning (Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1983, 1989; Hofstadter & Mitchell, 1994; Holyoak & Thagard, 1989; Keane, 1988; Medin et al., 1993). For example, the statement *an atom is like the solar system* hinges on understanding the common relations in the two pairs (e.g., smaller objects rotating about a larger object) and so is based not on simple feature matching but, instead, on relation matching (Gentner, 1983).

The implication of this view of comparison judgments is that alignability depends on the relational structure by which objects are represented. Furthermore, a change in representation may alter whether a feature in one object corresponds to a feature in another and so changes whether the features are perceived to be commonalities, alignable differences, or nonalignable differences. Barsalou (1982) makes a similar point—namely, that the feature representation of an object may be context dependent. The literature on comparison judgments extends this idea by noting that the relational structure of items may also be context dependent (see also Goldstone, Medin, & Gentner, 1991). For example, an insurance adjuster may represent the scenes in Figure 1, not according to the nature of the work being performed, which would align the robot at Jack's with the man at Grant's, but according to the costs that the insurance company might incur to replace, repair, or make well the machines and people in the repair shop. This perspective might, therefore, align the objects by physical characteristics, so that the two robot arms are placed in correspondence, whereas the man, the only element potentially requiring medical care, is perceived to be a nonalignable difference.

Prior research indicates, therefore, that feature alignability depends on relational structure. If people's causal explanations are sensitive to feature alignment, this finding implies that people's causal explanations will also de-

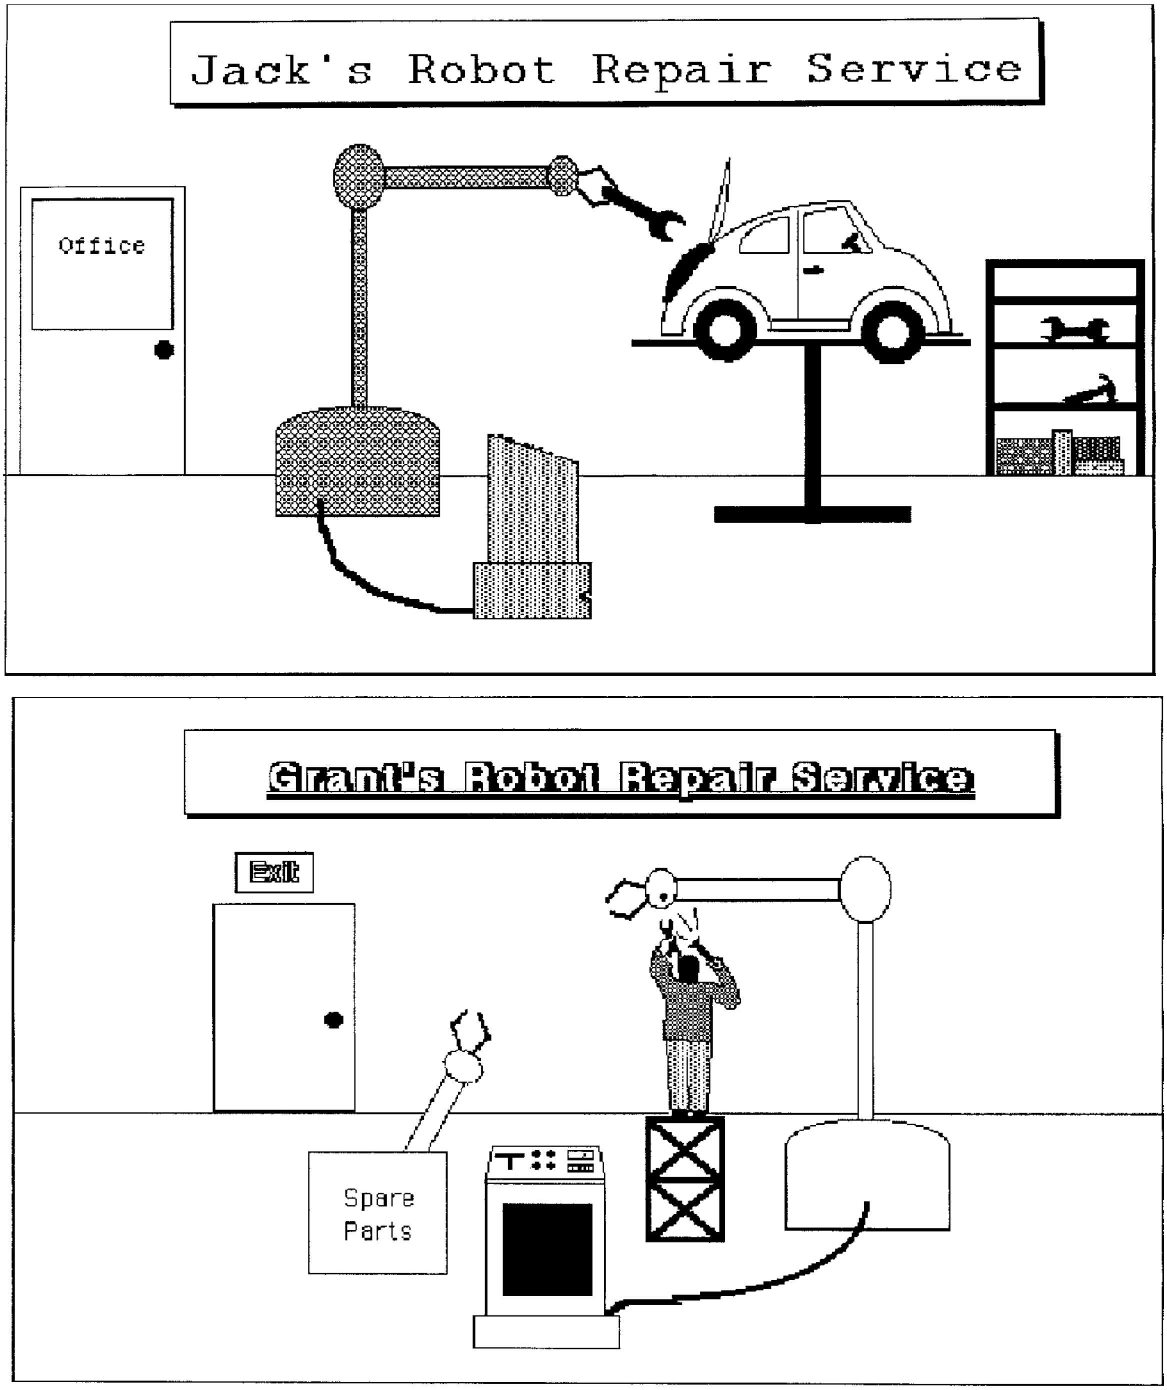

Figure 1. Drawings adapted from Markman and Gentner (1993a). From "Splitting the Difference: A Structural Alignment View of Similarity," by A. B. Markman and D. Gentner, 1993, *Journal of Memory & Language*, *32*, p. 519. Copyright 1993 by Academic Press. Adapted with permission.

pend on the relational structure by which they represent events. Experiment 3 was designed to test this hypothesis. In the third experiment, I manipulated the participants' representations of events by changing the context of the causal question. I then asked the participants to rate causal explanations that were based on differences that were either alignable or nonalignable, depending on the relational structure adopted. In this experiment, all events were presented as corresponding to an outcome in the comparison case. In accordance with the findings for Experiment 1, it was therefore predicted that the participants would rate explanations based on alignable differences higher than those based on nonalignable differences. Alignability, however, was expected to differ across conditions even though the event and the comparison case did not.

For example, the participants in Experiment 3 were presented with the pictures of the two repair shops shown in Figure 1. The participants in one condition were told that the repair being performed at Jack's shop was more expensive than the repair being performed at Grant's shop. The participants in this condition were expected to align objects according to the roles they occupied in the repair. Specifically, the participants in the *repair* condition of the robot event were expected to align the robot doing the repair at Jack's with the man doing the repair at Grant's and the car being repaired at Jack's with the robot being repaired at Grant's. By contrast, the participants in the *insurance* condition were told that the repair depicted at Jack's was more expensive to insure than the repair depicted at Grant's.[2] The participants in the insurance condition were expected to align objects according to physical characteristics. Specifically, they were expected to align the robot at Jack's with the robot at Grant's. It was expected that the car being repaired at Jack's and the man doing the repair at Grant's could not be aligned in either the repair or the insurance version. Hence, the robot scenario was constructed so that pairs of objects could be aligned by role (called *role-alignable* differences in the design), by physical characteristics (*form-alignable* differences), or by neither role nor physical form (*mismatches*; see Table 3, which shows feature alignability by relational structure for the robot event).

EXPERIMENT 3

Method

Participants. Eighty people were recruited for the third experiment in the same manner as the participants for Experiments 1 and 2.

Stimuli and Procedure. The participants were asked to rate explanations for two events, the *robot event* and the *household event*. For the robot event, the participants were presented with the two panels depicting Jack's and Grant's robot repair shops, as shown in Figure 1. In the *repair* version, the participants were told that "the repair that is depicted at Jack's is much more expensive than the repair that is depicted at Grant's. In the *insurance* version, the participants were told that "the repair that is depicted at Jack's is much more expensive *to insure* than the one depicted at Grant's" (emphasis added).

The participants were then asked to rate four possible explanations for the event on a scale of 1 (*very poor explanation*) to 7 (*excellent explanation*). Specifically, the participants were asked, "Please rate the following as possible explanations for the greater expense at Jack's." Explanations paired features of the panels and suggested that differences in sophistication caused the greater expense. Specifically, the participants rated (1) "the *robot* in Jack's is more sophisticated than the *man* at Grant's," (2) "the *car* in Jack's is more sophisticated than the *robot* at Grant's," (3) "the *robot* in Jack's is more sophisticated than the *robot* at Grant's," and (4) "the *car* in Jack's is more sophisticated than the *man* at Grant's" (italics in the original stimuli).

As was described in the overview of the experiment, these four explanations for the robot event were based on three types of differences. Specifically, Explanation 1 (robot–man) and Explanation 2 (car–robot) were based on role-alignable differences because the robot at Jack's and the man at Grant's, and the car at Jack's and the robot at Grant's occupied the same role in the repair operation. Explanation 3 (robot–robot) was based on form-alignable differences, because the robot at Jack's and the robot at Grant's were physically similar. Explanation 4 (car–man) was based on a *mismatch*, because these factors neither occupied the same role in the repair nor were they physically similar.

It was expected that the repair version, because of its focus on the work being performed, would cause the participants to structure their representations of the repair shops according to role in the repair operation. By contrast, the insurance version, because of its focus on replacement costs, would cause the participants to structure their representations according to physical form. As a consequence, it was expected that ratings of the role-alignable differences would be higher in the repair version than in the insurance version, ratings of the form-alignable difference would be higher in the insurance version than in the repair version, and ratings of the mismatch would not differ across versions.

For the household event, the participants were presented with the following information about two households:

> Consider the following two households. In Household A, Bill Jones works full time outside the home and is the primary source of income and Mary Jones works inside the home and provides care to Carl Jones who is 2 years old. In Household B, Linda Marks works full time outside the home and is the primary source of income and her sister Chris-

Table 3
Feature Alignability Depending on Relational Structure for the Robot Event in Experiment 3

	Relational Structure		
Difference	Repair (Aligned by Role)	Insurance (Aligned by Form)	Difference Type
Jack's robot–Grant's man (Explanation 1)	alignable	nonalignable	role alignable
Jack's car–Grant's robot (Explanation 2)	alignable	nonalignable	role alignable
Jack's robot–Grant's robot (Explanation 3)	nonalignable	alignable	form alignable
Jack's car–Grant's man (Explanation 4)	nonalignable	nonalignable	mismatch

tine Wilson works inside the home and provides care to Katie Marks who is 3 years old.

For the *happiness* version, the participants were told that Household A was happier, when considered in terms of family and social issues, than Household B. Specifically, the participants were provided with the following information:

Both Households A and B were part of a sample for a national research survey conducted on family and social issues. Answers on several questions indicated that Household A is significantly happier than Household B.

For the *health* version, the participants were told that Household A was healthier when considered in medical terms:

Both Households A and B were part of a sample for a national research survey conducted on medicine and health. Answers on several questions indicated that Household A is significantly healthier than Household B.

The participants then rated five explanations for the greater happiness or health of Household A. Specifically, the participants rated (1) "Bill Jones feels better than Linda Marks," (2) "Mary Jones feels better than Christine Wilson," (3) "Bill Jones feels better than Christine Wilson," (4) "Mary Jones feels better than Linda Marks," and (5) "Carl Jones feels better than Katie Marks."

It was expected that the happiness version, because of its focus on social and family issues, would cause the participants to structure the households in terms of traditional breadwinner–caregiver distinctions. By contrast, it was expected that the physical health version, because of its focus on physical well-being, would cause the participants to structure households according to physical similarity of members, particularly along the lines of gender, ob/gyn concerns, and age. Hence, explanations for the household event were also based on role-alignable differences, form-alignable differences, or a mismatch. A fourth type of difference, a *matching* difference, was also included for the household event. Specifically, Explanation 1 (Bill Jones–Linda Marks) and Explanation 2 (Mary Jones–Christine Wilson) were based on role-alignable differences, because Bill Jones and Linda Marks occupied the same role as breadwinner, whereas Mary Jones and Christine Wilson occupied the same role as caregiver. Although Mary Jones and Christine Wilson were both women, only one had had a child, which was expected to lessen their perceived physical similarity. This assumption was supported by the results of the pilot study on perceived similarity described below. Explanation 4 (Mary Jones–Linda Marks) was based on a form-alignable difference, because these household members were physically similar to the extent that both were women who had had a child. Explanation 3 (Bill Jones–Christine Wilson) was based on a mismatch, because the two people neither occupied the same role nor were physically similar. Explanation 5 (Carl Jones–Katie Marks) was based on a match, because the two children occupied the same role in the family and, on the basis of age (both were toddlers), were physically similar (see Table 3, which shows feature alignability by relational structure for the household event).

Relative ratings of the four types of explanations for the household event were expected to depend on relational structure implied by the version of the scenario. Specifically, ratings of the role-alignable differences were expected to be higher in the happiness version than in the health version, ratings of the form-alignable differences were expected to be higher in the health version than in the happiness version, and ratings of the mismatch and match were not expected to differ across versions.

Pilot study. A pilot study was conducted to examine the perceived similarity of items in the robot and household events. For the pilot, 20 students were recruited in the commons of the University of Chicago campus and offered $1 to fill out a questionnaire on "perceived similarity." The participants provided two sets of similarity ratings for both the robot event and the household event. For the robot event, the participants were presented with the pictures of the two repair shops and were asked to "please rate the similarity of the items below in terms of their *physical similarity*" and, again, to "please rate the similarity of the items below in terms of the *roles they fill in the repair shop*" (emphases in the original stimuli). For the household event, the participants were presented the description of the two households from the main experiment (but not the description of the greater happiness or greater health event) and were asked to "please rate the household members listed below in terms of *physical and health-related characteristics*" and to "please rate the similarity of the household members listed below in terms of *the family and social roles that they fill in their respective households*" (emphases in the original stimuli). Ratings were provided on scales of 1 to 7, with endpoints labeled *very different* and *very similar*, respectively. Order of presentation of the two similarity questions and of the robot and household events was counterbalanced.

The results of the pilot study supported the categorization of differences as role alignable, form alignable, and mismatched, as shown in Table 3 for the robot event. Specifically, the participants rated the role-alignable differences more similar in terms of role than of physical form [for Explanation 1, $M = 4.4$ and 2.5, respectively; paired $t(19) = 3.65$, $p < .01$; for Explanation 2, $M = 4.1$ and 2.9, respectively; paired $t(19) = 1.89$, $p < .09$]. By contrast, the participants rated the form-alignable difference (Explanation 3) less

Table 4
Feature Alignability Depending on Relational Structure for the Household Event in Experiment 3

Difference	Relational Structure: Happiness (Aligned by Role)	Relational Structure: Health (Aligned by Form)	Difference Type
Bill Jones–Linda Marks (Explanation 1)	alignable	nonalignable	role alignable
Mary Jones–Christine Wilson (Explanation 2)	alignable	nonalignable	role alignable
Bill Jones–Christine Wilson (Explanation 3)	nonalignable	nonalignable	mismatch
Mary Jones–Linda Marks (Explanation 4)	nonalignable	alignable	form alignable
Carl Jones–Katie Marks (Explanation 5)	alignable	alignable	match

Table 5
Ratings of Explanations Depending on Relational Structure for the Robot Event in Experiment 3

	Relational Structure	
Explanation Base	Repair	Insurance
Role-alignable difference		
Mean	4.26	3.11
Standard deviation	1.66	1.40
Form-alignable difference		
Mean	3.90	5.35
Standard deviation	2.07	1.88
Mismatch		
Mean	1.85	2.03
Standard deviation	1.42	1.37
Number of participants	40	40

similar in terms of role than of physical form [$M = 2.8$ and 5.0, respectively; paired $t(19) = -4.11$, $p < .001$]. The participants did not provide significantly different similarity ratings for the mismatched pair (Explanation 4) in terms of physical form and role [$M = 2.7$ and 2.4, respectively; paired $t(19) = 0.61$, $p > .55$].

Similarity ratings also supported the categorization of differences as role alignable, form alignable, mismatched, and matched, as shown in Table 4 for the household event. Specifically, the participants rated the role-alignable differences more similar in terms of role than of physical form [for Explanation 1, $M = 5.6$ and 3.9, respectively; paired $t(19) = 3.38$, $p < .01$; for Explanation 2, $M = 5.2$ and 4.5, respectively; paired $t(19) = 1.75$, $p < .10$]. By contrast, the participants rated the form-alignable difference (Explanation 3) less similar in terms of role than of physical form [$M = 3.3$ and 4.4, respectively; paired $t(19) = -2.10$, $p < .05$]. The mismatched pair (Explanation 4) did not differ in perceived similarity of role or physical form [$M = 3.1$ and 3.6, respectively; paired $t(19) = -1.00$, $p > .33$], nor did the matched pair differ [Explanation 5, $M = 4.2$ and 4.9, respectively; paired $t(19) = -1.51$, $p < .15$].

Design. The participants were presented one version of the robot scenario and one of the household scenario. The *repair* version of the robot scenario was paired with the *happiness* version of the household scenario, and the *insurance* version of the robot scenario was paired with the *health* version of the household scenario. Thus, the participants responded to the repair and happiness versions or to the insurance and health versions. Order of presentation of the scenarios was counterbalanced across scenarios.

Experimental design differed slightly for the robot and the household scenarios, which were therefore analyzed separately. The robot scenario involved a 2 × 3 mixed factorial design, with relational structure version (repair, insurance) as a between-subjects factor and explanation type (role alignable, form alignable, mismatch) as a within-subjects factor. The household scenario offered one additional explanation, so it involved a 2 × 4 mixed factorial design, with relational structure version (happiness, health) as a between-subjects factor and explanation type (role alignable, form alignable, mismatch, match) as a within-subjects factor.

Results

The results were analyzed separately for each scenario (see Tables 5 and 6, which show ratings of the explanations for each version of the repair and household scenarios, respectively). The participants' ratings for the robot scenario were analyzed with a repeated measures ANOVA, with relational structure (repair, insurance) as a between-subjects factor and difference type (role alignable, form alignable, mismatch) as a within-subjects factor. Ratings of the first two explanations, robot–man and car–robot, were combined to create the role-alignable measure. Results of the analysis did not differ when these explanations were analyzed separately. The ANOVA revealed no significant main effect of relational structure ($F < 1$) but a significant main effect of difference type [$F(2,77) = 63.97$, $p < .001$]. An additional analysis of this main effect indicated that the participants rated the form-alignable difference higher than the role-alignable difference [$F(1,78) = 9.29$, $p < .01$] and the role-alignable difference higher than the mismatch [$F(1,78) = 74.16$, $p < .001$].

As was predicted, however, the main effect of difference type was mediated by a significant relational structure × difference type interaction [$F(2,77) = 14.53$, $p < .001$]. Specific contrasts indicated that the participants' ratings of the role-alignable difference were higher in the repair version than in the insurance version [$F(1,78) = 11.17$, $p < .001$], whereas ratings of the form-alignable difference explanations were higher in the insurance version than in the repair version [$F(1,78) = 10.76$, $p < .01$]. Ratings of the mismatch did not differ across versions ($F < 1$).

The participants' ratings also differed within versions of the robot scenario [$F(1,78) = 29.09$, $p < .001$, for the repair version; $F(1,78) = 10.08$, $p < .001$, for the insurance version]. Consistent with the predicted effects of alignability, the form-alignable difference was rated higher than the role-alignable difference for the insurance version [$F(1,78) = 42.67$, $p < .001$], whereas the role-alignable difference was rated higher than the form-alignable difference for the repair version, although this latter difference was not significant [$F(1,78) = 1.13$, $p < .30$]. The mismatch received the lowest ratings in both versions of the robot scenario [$F(1,78) = 36.12$, $p < .001$, for the repair version; $F(1,78) = 42.67$, $p < .001$, for the insurance version].

Table 6
Ratings of Explanations Depending on Relational Structure for the Household Event in Experiment 3

	Relational Structure	
Explanation Base	Happiness	Health
Role-alignable difference		
Mean	4.62	3.86
Standard deviation	1.63	1.42
Form-alignable difference		
Mean	4.00	4.70
Standard deviation	1.78	1.68
Mismatch		
Mean	2.98	3.23
Standard deviation	1.61	1.80
Match		
Mean	4.13	4.83
Standard deviation	2.27	2.01
Number of participants	40	40

The participants' ratings for the household scenario were also analyzed with a repeated measure ANOVA, with relational structure (happiness, health) as a between-subjects factor and difference type (role alignable, form alignable, mismatch, match) as a within-subjects factor. As in the robot scenario, ratings for the first two explanations, Bill Jones–Linda Marks and Mary Jones–Christine Wilson, were combined to create the role-alignable differences. An analysis revealed no differences in the pattern of results when these explanations were analyzed separately, except as noted below.

An analysis of the household scenario revealed a pattern of results similar to that for the robot scenario. Specifically, the analysis revealed no significant main effect of relational structure ($F < 1$) but a significant main effect of difference type [$F(3,76) = 13.28$, $p < .001$]. Further analysis of this main effect indicated that ratings for the role-alignable, form-alignable, and matching differences did not differ ($F < 1$) but that all three of these explanations were rated higher than the mismatch [$F(1,78) = 34.56$, $p < .001$].

As in the robot scenario, the main effect of difference type was mediated by a significant relational structure × difference type interaction [$F(3,76) = 4.06$, $p < .01$]. Specific contrasts indicated that the role-alignable difference was rated higher in the happiness version than in the health version [$F(1,78) = 4.97$, $p < .05$] and that the form-alignable difference was rated higher in the health version than in the happiness version, although this difference in ratings was only marginally significant [$F(1,78) = 3.26$, $p < .08$]. Ratings of the match and mismatch did not differ across versions [$F(1,78) = 2.13$, $p < .15$, and $F < 1$, respectively]. A separate analysis for the two explanations that were combined to form the role-alignable difference revealed that ratings of Explanation 1 (Bill Jones–Linda Marks) and Explanation 2 (Mary Jones–Christine Wilson) were higher in the happiness version than in the health version but that this difference was significant for Explanation 1 [$F(1,78) = 14.45$, $p < .05$], whereas it was only marginally significant for Explanation 2 [$F(1,78) = 2.99$, $p < .09$].

Ratings also differed within versions of the household scenario [$F(1,78) = 10.93$, $p < .001$, for the happiness version; $F(1,78) = 21.70$, $p < .001$, for the health version]. Specifically, the role-alignable difference was rated higher than the form-alignable difference in the happiness version [$F(1,78) = 4.45$, $p < .05$], whereas the form-alignable difference was rated higher than the role-alignable difference in the health version [$F(1,78) = 8.00$, $p < .01$]. The mismatch received the lowest ratings in both versions [$F(1,78) = 11.98$, $p < .001$, for the happiness version; $F(1,78) = 4.63$, $p < .05$ for the health version]. Ratings of the matching difference did not differ from the other alignable differences for each version. That is, ratings of the matching feature did not differ from ratings of the role-alignable difference for the happiness version [$F(1,78) = 2.85$, $p < .15$], and ratings of the matching feature did not differ from ratings of the form-alignable difference in the health version ($F < 1$).

GENERAL DISCUSSION

The participants in Experiment 1 provided higher ratings for explanations that were based on the same type of difference as the event. That is, they rated alignable differences as being better explanations for alignable events than were nonalignable differences, and they rated nonalignable differences as being better explanations for nonalignable events than were alignable differences. The manner in which alignment was manipulated in Experiment 1 suggested a possible confound with perceived continuity. Experiment 2 was therefore intended to evaluate whether the participants would provide similar results for continuous and binary changes. The results for the second experiment supported this hypothesis. The participants rated explanations based on continuous causes as being better explanations for continuous events than were binary causes, but they rated binary causes as being better explanations for binary events than were continuous causes, although this latter difference was not reliable. The findings for Experiment 2 indicate, therefore, that people judge the degree of congruity between an outcome and a possible casual candidate in terms of continuity, and not necessarily in terms of alignability.

Experiment 3 was therefore intended to evaluate an effect of alignment on causal judgments by using a different methodology that would not confound alignment with continuity. In this study, alignment was manipulated by changing the perceived relational structure of events. In particular, the events in Experiment 3 could be represented in terms of task or in terms of physical form. The relative salience of these representations was manipulated by creating different versions of the scenarios. As was predicted, the participants provided higher ratings of explanations depending on whether the explanation was alignable or nonalignable according to the relational structure of the representation. Of particular relevance to the congruity hypothesis is the finding that the participants in Experiment 3 again provided higher ratings for explanations that were based on alignable differences for alignable events.

In addition, the participants' ratings of explanations also differed within versions of the scenarios as predicted by alignability. Explanations based on alignable differences were rated higher than those based on nonalignable differences within each version. This finding is important because across-scenario differences in ratings of the explanations might be attributed to differences in the versions other than relational structure and alignability. A within-scenario effect for alignability therefore bolsters support for hypotheses based on the effects of relational structure and factor alignability.

The results of Experiments 1–3 are, therefore, consistent with the hypothesis that people evaluate explanations by using congruity as a cue to causality. Prior research has conceptualized congruity in terms of magnitude—that is, in terms of length and strength of the candidate explanation and the event (Einhorn & Hogarth, 1986). Here, the concept of congruity is expanded to include the nature of the difference between the target and the comparison case, taking into account whether the candidate explanation and the event match in terms of alignment or in terms of continuity. It is important to note, at this point, that congruity is just one of several cues to causation and its impact on causal judgment may be limited, depending on the direction and size of other cues—in particular, perceived covariation (Einhorn & Hogarth, 1986). Nevertheless, these results point up the richness by which people may interpret this similarity-based cue to causation.

CONCLUSION

Prior research in causal reasoning has demonstrated that explanations may be based on distinctive features between an event and a comparison case. The present research extends these findings by showing that people attend to types of distinctive features. Specifically, evaluations of candidate explanations appeared to depend on whether features were perceived as alignable differences or nonalignable differences and whether they were perceived as binary differences or differences in degree. These findings, therefore, add to the literature on context effects in causal judgments. Previous studies have shown that explanations depend on the selection of a comparison case (e.g., Einhorn & Hogarth, 1986; McGill, 1989, 1993) and on the direction of comparison between the target and the comparison (e.g., Miller, Taylor, & Buck, 1991). The present research indicates that explanations also depend on the relational structure by which people represent the event and the comparison.

Evidence for the use of congruity as a cue to causality has implications for people's ability to understand their environment. In particular, it suggests that people may find it hard to identify "tipping point" or threshold effects in which high or low levels of a continuous variable can cause discontinuous outcomes. For example, it may be hard for people to understand or believe that an excess of sunlight may cause cancers. The idea that the slow, steady input of one variable can cause a sudden change in an output variable is not unheard of—for example, people may believe that someone may suddenly "boil over" in anger after having endured years of annoyance—but people who rely on congruity as a cue to causality may be slow to generate such hypotheses. For example, people may be more likely to blame an angry outburst on a recent one-time event, such as a change in job status or the death of a friend. Similarly, the finding that people think that continuous outcomes are produced by continuous inputs may cause them not to consider the effects of discontinuous inputs until other, more congruent explanations have been examined. Future research may examine conditions in which people are more likely to consider incongruent explanations for events.

REFERENCES

Barsalou, L. W. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, **10**, 82-93.

Einhorn, H. J., & Hogarth, R. M. (1986). Judging probable cause. *Psychological Bulletin*, **99**, 3-19.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, **41**, 1-63.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, **7**, 155-170.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity, analogy, and thought* (pp. 199-241). Cambridge: Cambridge University Press.

Gentner, D., & Markman, A. B. (1994). Structural alignment in comparison: No difference without similarity. *Psychological Science*, **5**, 152-158.

Goldstone, R. L. (1994). The role of similarity in categorization: Providing groundwork. *Cognition*, **52**, 178-200.

Goldstone, R. L., Medin, D. L., & Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognitive Psychology*, **23**, 222-262.

Hilton, D. J., & Slugoski, B. R. (1986). Knowledge-based causal attribution: The abnormal conditions focus model. *Psychological Review*, **93**, 75-78.

Hofstadter, D. R., & Mitchell, M. (1994). An overview of the Copycat project. In K. J. Holyoak & J. A. Barnden (Eds.), *Advances in connectionist and neural computation theory: Vol. 2. Analogical connections* (pp. 31-112). Norwood, NJ: Ablex.

Holyoak, K. J., & Thagard, P. (1989). Analogical mapping by constraint satisfaction. *Cognitive Science*, **13**, 295-355.

Hsee, C. K. (1996). The evaluability hypothesis: An explanation of preference reversals between separate and joint evaluation of alternatives. *Organizational Behavior & Human Decision Processes*, **46**, 247-257.

Hsee, C. K., Blount, S., Loewenstein, G. F., & Bazerman, M. H. (1999). Preference reversals between joint and separate evaluations of options: A review and theoretical analysis. *Psychological Bulletin*, **125**, 576-590.

Kahneman, D., & Miller, D. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, **93**, 136-153.

Keane, M. T. (1988). Analogical mechanisms. *Artificial Intelligence*, **2**, 19-23.

Lindemann, P. G., & Markman, A. B. (1996). Alignability and attribute importance in choice. In G. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Meeting of the Cognitive Science Society* (pp. 358-363). Mahwah, NJ: Erlbaum.

Markman, A. B., & Gentner, D. (1993a). Splitting the difference: A structural alignment view of similarity. *Journal of Memory & Language*, **32**, 517-535.

Markman, A. B., & Gentner, D. (1993b). Structural alignment during similarity comparisons. *Cognitive Psychology*, **25**, 431-467.

Markman, A. B., & Gentner, D. (1996). Commonalities and differences in similarity comparisons. *Memory & Cognition*, **24**, 235-249.

Markman, A. B., & Medin, D. (1995). Similarity and alignment in choice. *Organizational Behavior & Human Decision Processes*, **63**, 117-130.

McGill, A. L. (1989). Context effects in judgments of causation. *Journal of Personality & Social Psychology*, **57**, 189-200.

McGill, A. L. (1993). Selection of a causal background: Role of expectation versus feature mutability. *Journal of Personality & Social Psychology*, **67**, 701-707.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review*, **100**, 254-278.

Miller, D. T., Taylor, B., & Buck, M. (1991). Gender gaps: Who needs to be explained? *Journal of Personality & Social Psychology*, **61**, 5-12.

Nisbett, R. E., & Ross, L. D. (1980). *Human inference: Strategies and shortcomings in causal judgments.* Englewood Cliffs, NJ: Prentice-Hall.

Sanbonmatsu, D. M., Kardes, F. M., & Gibson, B. D. (1991). The role of attribute knowledge and overall evaluation in comparative judgment. *Organizational Behavior & Human Decision Processes*, **48**, 131-146.

Slovic, P., & Macphillamy, D. (1974). Dimensional commensurability and cue utilization in comparative judgment. *Organizational Behavior & Human Performance*, **11**, 172-194.

Tversky, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.

Zhang, S., & Markman, A. B. (1998). Overcoming the early entrant advantage: The role of alignable and nonalignable differences. *Journal of Marketing Research*, **35**, 413-426.

NOTES

1. The distinction between nonalignable and alignable differences is obviated if one assumes that people assign *characteristic not present* features to balance their object representations. For example, one school's having been adopted by the local university is a nonalignable difference if there is no corresponding feature for the other school, but it may be recast as an alignable difference by assigning the feature *not adopted* to the second school. Although conceptually possible, prior research suggests that people do not engage in across-the-board assignment of *characteristic not present* features to balance their representations of compared items. Were people to do so, it would imply that they represent all items to be compared with the same number of features, which is at odds with the finding that people's similarity judgments are directionally asymmetric, depending on the relative familiarity (i.e., number of features in the representation) of the two objects to be compared (Tversky, 1977) and with research that has shown the distinction between alignable and nonalignable differences to have empirical relevance.

2. Pretesting of the stimuli indicated that the participants would interpret the greater insurance expense as resulting from differences in the costs to repair or replace mechanical items in the shop and to provide medical treatment to people in the shop—for example, if there were a fire—and not as resulting from differences in the cost to warranty the quality of the repair.

(Manuscript received August 17, 2000;
revision accepted for publication December 11, 2001.)