Memory & Cognition
1995, 23 (3), 354-367

Judging a book by its cover: Interpretative effects of content on problem-solving transfer

MIRIAM BASSOK, LING-LING WU, and KAREN L. OLSETH
University of Chicago, Chicago, Illinois

We examine how cover stories of isomorphic problems affect transfer. Existing models posit that people retain content in problem representations and that similarities and differences between the "undeleted" cover stories might interfere with recognition of structural similarities. We propose that cover stories can affect transfer in another way—by inducing semantic knowledge that modifies problem structures. Two experiments examined how people represent and solve permutation problems dealing with random assignment of elements from one set to elements from another set. Although the problems were structurally isomorphic, cover stories involving different pairs of element sets led subjects to abstract different "interpreted structures." Problems involving objects and people (e.g., prizes and students) led subjects to abstract an asymmetric structure ("get") and problems involving similar sets of people (e.g., doctors and doctors) led subjects to abstract a symmetric structure ("pair"). Transfer was mediated by similarities and differences between the interpreted structures of the learned and the novel problems.

Structurally isomorphic or analogous problems can be solved in the same or in a similar way despite the fact that they are "dressed up" in different content covers. Thus, per definition, the problem's content is superficial to the solution. This simple logic guides mothers who attempt to teach their daughters via examples from their own "archaic" past, teachers who present students with structurally isomorphic word problems in math or in physics, and—of course—psychologists who study analogical problem solving. Yet, all too often this logic works only for the mothers, the teachers, or the experimenters—and not for the daughters, the students, or the subjects. In fact, the most prevalent finding in research on problem-solving transfer is that differences in the cover stories of analogous problems severely impair spontaneous retrieval and/or application of previously learned solutions. For example, many subjects fail to recognize that a convergence solution learned in the context of a military problem can be applied to an analogous medical problem (Gick & Holyoak, 1980); that a statistics principle learned from an example about weather forecasting can be applied to a problem about arrangements of pizza toppings (Ross, 1987); or that a physics equation learned for motion can be applied to an analogous problem dealing with bushels of potatoes (Bassok & Holyoak, 1989).

In addition to effects of such overall similarities and differences in content domains, transfer is affected by similarities and differences between specific entities that serve as arguments of the relational statements (Bassok, 1990; Gentner & Toupin, 1986; Holyoak & Koh, 1987; Ross, 1987, 1989). For example, subjects in Holyoak and Koh's (1987) study learned a convergence solution in the context of a story in which either laser beams or ultrasound waves were used to fuse a filament of a light bulb. Then they were presented with Duncker's (1945) radiation problem in which X-rays had to destroy a tumor. The laser version yielded a much higher level of spontaneous transfer to the X-rays problem than the ultrasound version, indicating that similarities between the specific forces affected retrieval of analogous solutions. Similarities and differences between corresponding entities in the learned (base) and the novel (target) problems also affect people's ability to apply the learned solutions. For example, children in Gentner and Toupin's (1986) study learned a story in which a squirrel was the protagonist and a frog was the villain. Then they had to enact an analogous scenario with a chipmunk and a toad. Children's performance was much better when they could map the chipmunk to the role of the squirrel (i.e., the protagonist) and the toad to the role of the frog (i.e., the villain) than when they had to reverse the structural roles of the chipmunk and the toad—to put the toad in the squirrel's role (the protagonist) and the chipmunk in the frog's role (the villain).

Preparation of this manuscript was supported in part by a grant from the University of Chicago School Mathematics Project. Portions of this research were presented at the 64th Annual Meeting of the Midwestern Psychological Association, May 1992, Chicago. The experiments formed the basis of the second author's masters thesis. We want to thank Cecily Cornelius and Denise Nelson for their help in collecting the data and Valerie Chase and Christina Honde for their editorial comments. Many thanks are extended to Dedre Gentner, Gerd Gigerenzer, Susan Goldin-Meadow, Bill Goldstein, Keith Holyoak, Susan Levine, Laura Novick, Steve Reed, and Lance Rips for their insightful comments on previous versions of this manuscript. The paper benefited greatly from discussions with Douglas Medin. Correspondence should be addressed to M. Bassok, Department of Psychology, University of Chicago, 5848 South University Ave., Chicago, IL 60637 (e-mail: mbmb@speech.uchicago.edu).

Two Types of Content Effects: Memory and Interpretation

Why do similarities and differences in content affect transfer? Such effects can reflect people's knowledge, or belief, that very often similarities in content indicate similarities in structure. For example, Hinsley, Hayes, and Simon (1977) have shown that after reading no more than a few words of an algebra word problem, subjects could correctly identify the problem's category (e.g., a work problem, a distance problem), thereby accessing information relevant to the problem's solution. In addition to such domain-specific expectations, people might be guided by a generalized assumption that content could provide them with valuable information about problems' structures. Such an assumption would be consistent with their tendency to infer essence from appearance (e.g., round things can roll; things with wings can fly)—a tendency that is often justified by experience (Medin & Ortony, 1989; Smith & Heise, 1992).

How does people's tendency to rely on similarities and differences in content affect transfer? The prevalent explanation for such effects is that people retain content in their representations and rely on similarities and differences in such "undeleted" aspects of content for retrieval and application of previously learned solutions. Specifically, existing models of transfer (e.g., Gentner, 1983, 1989; Gentner & Forbus, 1991; Holyoak & Thagard, 1989; Thagard, Holyoak, Nelson, & Gochfeld, 1990) posit feature-matching mechanisms that establish one-to-one matches between various aspects of the base problem and the target problem. The aspects on which the base and the target problems are matched are of two general types: aspects of structure (relational predicates, such as: "attack (forces, target)") and aspects of content (arguments of these predicates, such as: the forces are X-rays and the target is a tumor). It is assumed that matches and mismatches in aspects of content (object attributes) either support or compete with matches and mismatches in aspects of structure (relations between objects).

Within this view, failures of transfer occur when mismatches in superficial aspects of content win in their competition with matches in the solution-relevant aspects of structure. For example, subjects fail to recognize that the same structural relation "attack (forces, target)" is hidden behind different cover stories because the arguments of the "attack" relation are different (e.g., in one problem the forces are X-rays, whereas in the other problem the force is an army). In a similar way, mapping errors occur when matches in superficial aspects of content (e.g., a chipmunk is more similar to a squirrel than to a frog) win in their competition with mismatches in solution-relevant aspects of structure (e.g., the chipmunk is supposed to be the villain, but the squirrel was the protagonist). We will refer to this view as the "memory" explanation, because it posits that people *remember* (or retain) the specific content and that transfer is mediated by similarities between the specific *content instantiations of the base and the target problems.*

In the present paper we propose another way in which content mediates transfer. We will argue that in addition to a general belief that similar books are likely to have similar covers, people actually judge each book by its cover.[1] That is, we propose that people *interpret* the structure of the base and the target problems using semantic knowledge triggered by the situations described in these problems—by the particular objects and relations that appear in each cover story. Transfer between structurally isomorphic problems dressed up in different contents would depend on similarities between the *interpreted structures abstracted for the base and the target problems.* Because interpreted structures might not coincide with the objective structures of the problems, transfer might be impaired by differences between the interpreted structures of the base and the target problems.

The fact that different covers or setups affect problem representations is well documented. A famous example would be Duncker's (1945) work on "functional fixedness." Subjects had to use a box as a platform on which they could mount a candle. Presented with a box that contained candles, the functional role of the box as a container prevented them from noticing that the box could also serve as a platform, but when the box was empty subjects easily solved the problem. More recently, Kotovsky, Hayes, and Simon (1985) demonstrated that inferences triggered by object attributes can affect the relative difficulty of two isomorphs of the Tower of Hanoi problem. In both problems, subjects had to place disks of different sizes on top of each other according to a prespecified set of rules, but in one of the problems the disks supposedly represented acrobats. Subjects had little difficulty placing a big disk on top of a small disk, but they refrained from placing a big acrobat on the shoulders of a small acrobat. The relative difficulty of the "acrobats" version was explained by a structural constraint inferred from people's knowledge that a big acrobat cannot jump on the shoulders of a small acrobat without hurting him.

Although effects of semantic knowledge on problems' difficulty are well documented, the possibility that content affects transfer by inducing inferences about problems' structures generally has not been addressed, the work on isomorphs of the Tower of Hanoi problem being a notable exception (Hayes & Simon, 1977; Kotovsky, Hayes, & Simon, 1985). The present paper aims to demonstrate that interpretative effects of content affect the way in which people apply previously learned solutions to novel problems. In order to clarify the distinction between effects of memory for particular content instantiations of problems and interpretative effects of content, we contrast the predictions derived from these two explanations.

Object-Mapping Versus Mapping of Interpreted Structures

Suppose a child first learns a problem (the base) in which plates are divided equally among several trays, and the problem is solved using an equation in which the plates are in the numerator and the trays are in the denominator (plates/trays). Then, the child is asked to apply the learned solution to a new division problem (the target) dealing with cookies and plates. Would the child now write an equation in which the plates were in the numerator and the cookies in the denominator (plates/ cookies), or would the child put the cookies in the numerator and the plates in the denominator (cookies/plates)?

Existing models of transfer would predict that because both problems share the same structure ("divide") and because the structure does not dictate which elements should be the divided and the dividing set, the child's solution to the target problem would depend on similarities and differences between the objects that serve as arguments of the divide relation in the two problems. Specifically, the child should engage in *object-mapping*—she should prefer to place similar objects in similar structural roles (i.e., similar slots in the learned equation which corresponds to the "divide" relation). Because cookies are both unlike plates and unlike trays, the child should be indifferent about where she should put the cookies. However, because plates appear in both problems, and because they appear in the numerator of the base problem, the prediction would be that the child would put plates in the numerator and cookies in the denominator (plates/cookies).

It is important to note that feature-matching models that posit such object-mapping effects compute separate matches for each pair of objects and relations in the base and the target problems. For example, the similarity between cookies (food) in the target problem and plates (dishes) in the base problem would be determined irrespective of the fact that there were trays in the base problem and plates in the target problem. Moreover, because each match is established independently of other matches, identical objects (plates and plates) would be always more similar to each other than any pair of nonidentical objects. These independent matches and mismatches compete or support each other until the system achieves the best overall correspondence between all the relations and the objects in the two problems.

Note, however, that the specific object attributes and the specific relations in each problem are not independent from the context in which they appear (Medin, Goldstone, & Gentner, 1993; Ortony, 1979). For example, a problem involving plates and trays and a problem involving plates and cookies are likely to trigger semantic knowledge suggesting that the problems have a "container" structure. Unlike the mathematical "divide" structure, the "container" structure constrains the structural roles of the paired elements: In the base problem the plates are the contained elements and the trays are the containers, whereas in the target problem the cookies are the contained elements and the plates are the containers. That is, the plates in the base and the target problems are not simply identical dishes (match). Rather, the plates in the base problem have the salient attribute of being the contained elements, whereas the plates in the target problem have the salient attribute of being the containers (mismatch).

Given that the child abstracts an interpreted "container" structure for the base and the target problems, our prediction was that she would compute similarities between the interpreted objects of the base and the target problems and place the cookies and the plates in their respective structural roles in the interpreted structure ("container"). The equation plates/trays used for solving the base problem implies that the contained elements (plates) are in the numerator and the containers (trays) are in the denominator. Hence, we predicted that the child would engage in *mapping of the interpreted structure*—placing the contained elements (cookies) in the numerator and the containers (plates) in the denominator (cookies/plates).

In general, our claim is that when people are presented with a problem involving several entities, they reason about the situation described in the problem using knowledge about the way in which these entities typically interact with each other (e.g., lollipops can be put into jars, but jars cannot be put into lollipops; children eat cakes, but cakes do not eat children). As a result, they abstract interpreted structures that include the reasons why certain entities play certain structural roles. The interpreted structures also determine the relevant "respects" of similarity between the objects and the relations in the base and the target problems. That is, similarities between the objects and the relations in the base and the target problems are computed on attributes inferred from knowledge about the situations described in the cover stories.

To date, object-mapping effects have been documented in several studies. As mentioned earlier, Gentner and Toupin (1986) demonstrated such effects in their study on children's enactments of stories about animals (e.g., mapping a toad rather than a chipmunk into the role of a frog). Ross (1987, 1989) obtained similar results with adults solving probability problems. For example, subjects in the 1989 study first learned a permutation problem in which cars were randomly assigned to mechanics; then they had to apply the learned solution to another permutation problem in which scientists were randomly assigned to computers. The subjects received the previously learned equation and were asked to instantiate it with the appropriate values (i.e., with the number of scientists and the number of computers). Ross found that, although scientists had to be placed in the structural role of the cars (because both served as the randomly assigned set) and computers in the structural role of the mechanics (because both served as assignees), subjects erroneously solved the target problem by placing scientists in the role of the mechanics and computers in the role of the cars.

To understand these findings, let us examine the structure of such permutation problems. The "assign" structure involves random assignment of elements from one set to elements from another set, and the solution depends on the direction of assignment—which elements are the randomly assigned set and which are the set of assignees. When a person (e.g., a manager) performs the assignment, this person can decide whether to assign cars to mechanics or mechanics to cars, or whether to assign scientists to computers or computers to scientists. Thus, the "assign" structure does not constrain the structural roles of the elements—each set of elements can be either the assigned set or the set of assignees. Under the assumption that the problems share the "assign" structure and that the objects and their structural roles are independent, the object-mapping hypothesis posits that subjects determine the structural roles of the objects by computing direct similarities between the objects in the base and the target problems. Specifically, according to Ross (1989), subjects remember that "one variable was instantiated by animate objects and the other by inanimate objects" (p. 465) and therefore place scientists in the role of the mechanics (animate) and computers in the role of the cars (inanimate).

The object-mapping hypothesis accounts quite well for this pattern of transfer results. However, the preference to place scientists in the role of the mechanics and computers in the role of the cars can be also explained by our interpreted-structure hypothesis. Note that regardless of the direction of assignment (e.g., cars to mechanics or mechanics to cars; computers to scientists or scientists to computers), the *outcome* of such assignments is that mechanics *get* cars and scientists *get* computers. That is, in the "get" outcome of assignment, the structural roles of the elements (receivers and givens) are constrained: Inanimate elements cannot be the receivers of animate elements, they can only be the givens. The subjects could have assumed that because the cars and the computers were the givens they were also the assigned set, and because the mechanics and the scientists were the receivers they were also the assignees. That is, the subjects could have abstracted an interpreted "get" structure rather than the objective "assign" structure. If so, mapping between animate elements, they were actually mapping the receivers, and mapping between inanimate elements, they were actually mapping the givens.

Notice that the object-mapping hypothesis and the interpreted-structure hypothesis make the same prediction about transfer performance on these problems. Unlike the previous example of division problems, where direct similarities between objects (plates and plates) were separated from similarities in the interpreted structural roles of these objects (contained elements and containers), in the permutation problems used by Ross (1989) similarities between objects (animate vs. inanimate) coincided with similarities in the interpreted structural roles of these objects (receivers vs. givens). Whenever the default similarities in object attributes coincide with the interpreted attributes of these objects, it is difficult to distinguish between object mapping (e.g., animate/ inanimate) into the uninterpreted structures (e.g., "assign"), mapping of the interpreted objects (e.g., receivers/ givens) into the interpreted structure (e.g., "get"), or some combination of both types of mapping.

Our experiments were designed to examine the hypothesis that subjects abstract interpreted structures for the base and the target problems and place the entities in their interpreted structural roles. In Experiment 1 we used experimental materials and an experimental design similar to those used by Ross (1987, 1989),[2] but we modified the materials and the design to demonstrate the existence of interpretative effects of content that cannot be explained by mapping. In Experiment 2 we contrasted predictions derived from the object-mapping hypothesis with predictions derived from our interpreted-structure hypothesis using a design analogous to that of the previous example involving division problems with cookies, plates, and trays.

EXPERIMENT 1

The main purpose of Experiment 1 was to examine whether different content covers lead to abstraction of different interpreted structures and thereby affect transfer. All permutation problems in Experiment 1 had the same mathematical structure and used the same relational predicate: "assign (assigner, assigned set, set of assignees)." That is, in all problems, someone (e.g., a teacher) randomly assigned elements from one set (e.g., prizes) to elements from another set (e.g., students). We used two types of cover stories for the training and transfer problems: problems involving two similar sets of people (e.g., doctors from Chicago and doctors from Minnesota) and problems involving a set of people and a set of objects (e.g., students and prizes). The paired elements and the cover stories were chosen such that subjects would be likely to abstract different interpreted structures. Specifically, problems involving two similar sets of people described situations in which the elements from the two sets played symmetric roles in the outcome of assignment (e.g., doctors from two hospitals work together)—a *"pair (set A, set B)"* interpreted structure. Problems involving a set of objects and a set of people described situations in which the elements from the two sets played asymmetric roles in the outcome of assignment (e.g., students get prizes)—a *"get (receivers, givens)"* interpreted structure.

To test whether subjects interpret the assignment structure by including structural constraints implied by the outcome of assignment, we asked subjects to solve the permutation problems without receiving training from us. We examined the equations that subjects spontaneously constructed for the "pair" and the "get" permutation problems. We predicted that the "pair" problems would lead subjects to construct equations in which both sets of people played symmetric roles and the "get" problems would lead them to construct equations in which people and objects played asymmetric roles.

Experiment 1 also examined how people solve these two types of problems after receiving a worked-out solution of a problem involving a different pair of element sets. In particular, we examined how subjects solve "get" problems following training on "pair" problems. As we will elaborate later, this condition does not permit mapping to example. We predicted that, without the "benefits" of mapping, subjects would operate under the default assumption that the direction of assignment corresponded to the structural roles of the two sets of elements in the outcome of assignment (objects assigned to people rather than vice versa). We refer to this tendency as an *interpretative bias*. We also examined how subjects solve "get" problems following training on "get" problems. We predicted that in this case subjects would map people to people and objects to objects. Such results would replicate prior findings documenting object mapping, but they would be also consistent with the possibility that subjects *map interpreted structures* across problems (mapping receivers to receivers and givens to givens). Moreover, we predicted an interaction between the interpretative bias and mapping: more mapping when the training problems are consistent with the interpretative bias (objects assigned to people) and less when they are inconsistent with the interpretative bias (people assigned to objects).

Method

Subjects

The subjects were 198 college students, 137 from the University of Chicago and 61 from Northwestern University. They were recruited by advertisements posted on campus asking for people to participate in a study piloting new instructional materials in various domains (e.g., physics, economics, probability) and saying that subjects would be asked to study one or more topics from these domains. The subjects varied in their knowledge of probability, but were screened for knowledge of the target material. Specifically, 56 subjects constructed a correct equation to a permutation problem before receiving training from us and were excluded from the study. We report results obtained from the remaining 142 subjects, 64 females and 78 males, who initially constructed incorrect equations to the training permutation problems. The subjects were paid for participating in the study.

Materials

A short chapter (two pages) that covered the topic of permutations was written for the study and was used in training. The chapter first presented the necessary background information (a definition of probability as the relative frequency of the target outcome, a definition of permutations as particular arrangements of r elements out of n possible elements, etc.). The background information was illustrated by random drawings of balls from a bag. Then, the chapter presented the equation that should be used for the solution of permutation problems, $1/n(n-1)(n-2) \ldots (n-r+1)$, and a worked-out solution to one of the permutation problems described below. Note that only one set of elements—the randomly assigned set (n)—is relevant to the solution of such problems and appears in the equation (the r term stands for the number of random drawings).

Four isomorphic permutation word problems that differed in their cover stories were constructed for the study and served either as training problems or as transfer problems. Two problems paired people and objects: One problem served as a training problem (secretaries and computers), and the other served as a transfer problem (students and prizes). The other two problems paired elements from two similar sets of people: one problem served as a training problem (doctors from two hospitals) and the other served as a transfer problem (children from two nursery schools).[3] All problems involved two sets of ordered elements, m and n, that differed in size. For example, one problem involved a set of $m = 25$ secretaries listed in order of their work experience and a set of $n = 21$ computers labeled with serial numbers 10075 through 10095. In all problems, someone (e.g., a teacher, a manager) randomly assigned three elements from the n set (e.g., computers) to three ordered elements from the m set (e.g., secretaries) and subjects had to find the probability that a specific subset of elements from the two sets would end up being paired with each other (e.g., the first three students in alphabetical order receive the first three prizes, respectively). The outcome question in all problems was consistent with the hypothesized interpreted structure. The question in problems with two similar sets of people asked for the probability that certain pairs of people would end up working together, and the outcome question in problems with objects and people asked for the probability that certain people would end up receiving certain objects.

Each of the four different cover stories had two versions that reversed the direction of assignment. For example, in one version, computers (n) were randomly assigned to secretaries (m), and in the other version, secretaries (n) were randomly assigned to computers (m). To control for possible effects of set size, each of the eight different problems had two versions that reversed the number of elements in the two sets: In one version, m was bigger than n, and in the other version m was smaller than n. Thus, each of the four cover stories had four different versions, and altogether there were 16 different permutation problems. One version of a problem for each of the four cover stories appears in Appendix A.

Note that our problems were quite difficult to understand, leaving sufficient room for interpretative effects of content. Because the solution to the permutation problems depends on the size of the randomly assigned set, in problems involving two sets of elements subjects have to identify the direction of assignment to determine which of the two sets is the randomly assigned set. Hence, two-set permutation problems are more difficult to understand than permutation problems dealing with random drawings of balls from a bag where there is only one set of elements (the balls).

Our problems were also more difficult than the two-set permutation problems used by Ross (1987, 1989). Because Ross examined effects of memory for content rather than the interpretative effects of content, when he reversed the structural roles of objects and people he changed the phrasing of the problems to avoid unreasonable scenarios—he interpreted the problems for the subjects to clarify the direction of assignment. For example, some pairs of reversed problems differed in their relational predicates: mechanics chose cars (a two-place *"choose"* predicate), whereas cars were assigned to mechanics by a manager (a three-place *"assign"* predicate). Other pairs of reversed problems did not involve the same element sets: In one problem, mechanics (people) chose *cars* (objects), whereas in the other, *car owners* (people) chose mechanics (people). Unlike Ross, we wanted to examine the effects of inferences triggered by people's knowledge that objects cannot either choose or get people. Hence, we did not interpret the problems for the subjects: all our problems had the same relational predicate ("assign"), the reversed problems used the same element sets, and the paired sets in each problem were described in a similar way (as ordered sets). This standard phrasing eliminated various auxiliary cues that typically help subjects to decide what the direction of assignment is. Thus, although the direction of assignment was explicitly stated in every problem (e.g., "The manager of the company randomly assigns computers to secretaries"), it was less transparent than in the problems used by Ross.

Procedure

Each subject received one training problem and one transfer problem that differed in content. Using the two directions of assignment of the "secretaries and computers" problem, we created two training conditions: objects assigned to people (OP) and people assigned to objects (PO). The "doctors and doctors" problem served as the third training condition—people assigned to people (PP); the direction of assignment in this problem was randomized across subjects. Using the two directions of assignment of the "students and prizes" problem, we created two transfer conditions: OP and PO. The "children and children" problem served as the third transfer condition (PP); the direction of assignment in this problem was randomized across subjects. Thus, our design involved nine experimental conditions: three training (OP, PO, PP) by three transfer (OP, PO, PP). Subjects were assigned randomly to one of the nine conditions according to the combination of the training and transfer problems they had received. The number of subjects in these nine conditions ranged between 15 and 18.

The subjects were tested individually in one session that lasted between 20 and 30 min. They first received the training permutation problem and were asked to solve it as best they could. The subjects then received the training chapter, which included a worked-out solution of the problem that they initially had solved on their own. After notifying the experimenter that they were comfortable with the learned material, the subjects received the transfer problem together with the learned equation. The equation was already adjusted for the fact that there were only three random drawings ($1/n(n-1)(n-2)$), and the subjects were asked to instantiate the equation with the appropriate value (n). Forty-seven subjects were asked to talk aloud while solving the training and transfer problems, and their solutions were tape-recorded for later analysis.[4]

Results

Equations Constructed for the Training Problems Before Instruction

As mentioned earlier, 56 subjects constructed correct equations, and their results were excluded. We examined the solutions of the remaining 142 subjects who solved either problems with objects and people (OP and PO, $N = 95$) or problems with two sets of people (PP, $N = 47$). Forty-eight out of the 142 subjects (34%) failed to construct an equation. There was no difference in the frequency of failure to construct an equation in subjects who received problems with objects and people (35%) and those who received problems with two sets of people (33%) [$\chi^2(1) = 0.15$, n.s.]. A total of 94 subjects (66%) constructed incorrect equations. These incorrect equations varied in complexity, but most of them (93%) took the form of a fraction that resulted in a value smaller than 1. This suggests that the vast majority of subjects who constructed an incorrect equation for the training problem had some understanding of probability but did not know how to solve the particular permutation problems.

We classified the incorrect fraction equations into two categories, asymmetric and symmetric, according to the roles given to the two sets of elements in the equation each subject constructed. Asymmetric equations were fractions that either involved only one set of elements, m or n (e.g., $1/n$, $1/m!$), or placed the two sets of elements into different roles (e.g., m/n, n/m^3). Symmetric equations were fractions that placed the two sets of elements in similar roles (e.g., $1/m \cdot n$, $1/(m+n)^3$). As predicted, problems involving objects and people (both PO and OP) and problems involving two similar sets of people (PP) led subjects to construct different types of equations. Out of the 62 subjects in the PO and OP conditions who constructed incorrect equations, 87% constructed asymmetric equations and only 13% constructed symmetric equations. In contrast, of the 32 subjects who constructed incorrect equations to problems with two sets of people (PP), 22% constructed asymmetric equations and 78% constructed symmetric equations. The difference in frequency of symmetric and asymmetric equations generated for these two types of problems was highly significant [$\chi^2(1) = 39.41, p < .001$].

These results provide strong evidence that permutation problems with cover stories involving objects and people were perceived as different in type from permutation problems with cover stories involving two sets of people. Even though the problems were structurally isomorphic and the characteristics of the elements in the two sets were superficial to the solution of the problems (i.e., to the direction of assignment), problems that differed in their element sets and in their outcomes led subjects to interpret the problems in very different ways. In particular, permutation problems involving two sets of people and a "pair" outcome led subjects to abstract a "pair" structure: "The first doctor is *paired* with the first doctor in the two lists. Now for the second doctor in the two lists ... *because it has to work for both groups*" (Subject 29). In contrast, problems involving objects and people and a "get" outcome led subjects to abstract a "get" structure: "Each secretary would have a 21 in 25 chance of *getting* a computer" (Subject 15).

Instantiation of the Learned Equations

Solutions were scored as correct if subjects instantiated the given equation (n) with the number of elements in the randomly assigned set and as incorrect if subjects instantiated n with the number of elements in the set of assignees. Solutions were also scored as incorrect for 9 subjects (6% of our subject population) who refused to instantiate the given equation with either one of the two sets. Hence, for 94% of our subjects, an incorrect solution indicates that they chose to instantiate the equation with the alternative set of elements. This methodology and scoring procedure, similar to that used by Ross (1987, 1989), enabled us to measure subjects' preference for one of the two sets of elements to serve in the role of the assigned set. Table 1 presents the percentage of correct solutions for the nine different combinations of training and transfer problems.

Subjects' performance was not affected by differences in the size of the assigned and the receiving sets [$\chi^2(1) =$

Table 1
Percentage of Correct Solutions on Transfer Problems in the Nine Different Combinations of Training and Transfer in Experiment 1

	Transfer					
	OP		PO		PP	
Training	% Corr.	n	% Corr.	n	% Corr.	n
OP	89	18	0	16	60	15
PO	33	15	53	15	50	16
PP	67	15	31	16	50	16
Total	65	48	28	47	53	47

Note—OP denotes objects assigned to people, PO denotes people assigned to objects, and PP denotes people assigned to people.

0.08, n.s.], so we combined the results of subjects who had received problems with reversed sizes of the two sets. The data were analyzed using the logit model based on maximum likelihood estimation, with training (3) and transfer (3) serving as factors. The analysis yielded a significant effect for the difference between the three types of transfer problems [$\chi^2(2) = 11.17, p < .004$]. Overall, subjects' performance was best (65% correct) on problems in which objects were assigned to people (OP) and worst (28% correct) on problems in which people were assigned to objects (PO). The level of performance on problems that involved two similar sets of people (PP) was intermediate (53% correct). In addition, there was a strong interaction between the training and the transfer conditions [$\chi^2(4) = 15.44, p < .004$]. In what follows, we interpret the transfer results and relate them to the erroneous equations constructed by the subjects for the training problems. We first discuss transfer to problems with two sets of people and then discuss transfer to problems with people and objects.

Transfer to problems with two sets of people. Overall, performance on PP transfer problems was at chance level (53% correct), and it was similar for all three training conditions [$\chi^2(2) = 0.31$, n.s.]. Because subjects had to choose only which of the two sets of elements to instantiate as n in the equation, transfer performance of about 50% indicates lack of preference for either set. However, a closer analysis of subjects' performance on the PP transfer problems revealed that the three training conditions did, in fact, have different effects on transfer. Subjects trained on PP problems ($N = 16$) were equally likely to instantiate the equation with either one of the two sets of people. Because the outcome of PP problems is that the elements from the two sets are paired with each other in a symmetric way, subjects had difficulty deciding which of the two sets of people should play the role of the assigned set and which should play the role of the assignees: "Okay, then, so my problem is here in that I know that it's going to be *either* 16, 15, 14 *or* 20, 19, 18" (Subject 20). Although these subjects learned that only one set of elements should appear in the equation, they failed to learn from the symmetric PP training problems which of the two similar sets of people should appear in the equation.

Transfer performance of subjects in the asymmetric OP and PO training conditions was quite different. Out of the 31 subjects in these two conditions, only 22 instantiated the equation with one set of elements (77% of them were correct). The error made by the remaining 9 subjects was not that they chose the incorrect set of people—they refused to instantiate the equation with only one set (either m or n). These subjects crossed out the equation that appeared with the test problem and constructed a symmetric equation that included both sets. For example, they combined the two sets of people and sampled pairs of people from the combined set: "What's that have to do with 16 kids from 20 kids? ... So you have to take all the kids from one *and* all the kids from the other ... *Sum of elements* is 36, ... 36, 34, 32 ... *Every time you do this you take out 2*" (Subject 49). These subjects understood that problems involving two symmetric sets of elements differed in structure from the learned problems involving two asymmetric sets of elements. Just as the subjects who initially constructed symmetric equations for the PP training problems, these subjects constructed symmetric equations for the transfer problems.

Transfer to problems with a set of objects and a set of people. As can be seen from the bottom row of Table 1, subjects were, overall, much more successful on transfer problems that assigned objects to people (65% correct) than on problems that assigned people to objects (28% correct) [$\chi^2(1) = 10.90, p < .001$]. Because there were only two sets of elements to choose from, this differential success reflects the fact that in both conditions the majority of subjects chose the set of objects rather than the set of people to serve as the assigned set (this choice is represented by the 65% correct and 72% incorrect solutions in the OP and PO conditions, respectively). The bias to assign objects to people in the transfer problems is consistent with the functional asymmetry between people and objects in the outcome of assignment: regardless of the direction of assignment, people end up getting objects—objects do not get people. These results support our claim about the existence of an interpretative bias.

In addition to the overall interpretative bias there was a very strong interaction between the training and the transfer problems [$\chi^2(1) = 14.71, p < .001$]. The interpretative bias was clearly demonstrated in the PP training condition. In this condition, subjects could not engage in object mapping because the elements in the transfer problem were either equally similar (students) or equally dissimilar (prizes) to the two sets of doctors in the training problem. Moreover, the subjects could not engage in mapping of the interpreted elements because in the training problem the elements were arguments in a "pair" relation (doctors paired with doctors), whereas in the transfer problem the elements were arguments in the "get" relation (students get prizes). Nonetheless, 67% of the subjects in the PP training condition were correct on the OP transfer problems but only 31% were correct on the PO transfer problems. These results establish the existence of a baseline interpretative bias that is unrelated to either object mapping or mapping of the interpreted structure.

In the OP and PO training conditions, when both the example and the transfer problems involved people and objects (first two rows × first two columns in Table 1), a total of 73% of the subjects were correct when objects and people played similar roles in the training and transfer problems (an average of the cells OP→OP & PO→PO), but only 16% were correct when the roles of objects and people were reversed (an average of the cells OP→PO & PO→OP [$\chi^2(1) = 14.71, p < .001$]. This pattern of results replicates the "cross-mapping" results obtained by Ross (1987, 1989) and by Gentner and Toupin (1986), showing that when subjects apply a learned solution to a novel problem they tend to place similar elements in similar structural roles. However, as explained in the introduction, the present design cannot determine whether these results reflect object mapping, mapping of the interpreted "get" structure, or some combination of both types of mapping. The protocols obtained from several subjects in these two training conditions indicate that at least some of the subjects engaged in mapping of the interpreted "get" structure. Subjects trained on OP problems understood that n represents the set of givens: "It's like in the previous question ... the best secretaries are going *to get* the first ones, and the best students are theoretically going *to get* the top prizes" (Subject 11). Subjects trained on PO problems understood that n represented the set of receivers: "Because on secretary problem computers were being *given out* and students are given prizes, so I choose students to find the probability of *what they are getting*" (Subject 52).

As predicted, our design revealed both mapping effects and interpretative effects of content that are unrelated to mapping. These two effects interacted with each other. When the direction of assignment was consistent with the interpretative bias (OP training), 94% of the subjects instantiated the transfer equation with objects, solving correctly 89% of the matching OP transfer problems and none of the nonmatching PO transfer problems (0%). In contrast, when the direction of assignment was inconsistent with the interpretative bias (PO training), only 60% of the subjects instantiated the transfer equation with people, solving correctly 53% of the matching PO transfer problems and 33% of the nonmatching OP transfer problems. The difference in the magnitude of mapping in these two conditions was highly significant [$\chi^2(1) = 7.39, p < .001$]. Although there was a large difference between these two conditions in the absolute frequency of responses that matched the solutions of the training problems, the two conditions did not differ in the magnitude of mapping relative to the baseline interpretative bias established in the PP training condition: The OP training increased subjects' tendency to instantiate the equation with objects by 26% (94% vs. 68% object instantiations in the OP vs. PP training conditions, respectively), and the PO training decreased subjects' tendency to instantiate the equation with objects by 28% (40% vs. 68% object instantiations in the PO vs. PP training conditions, respectively).

Discussion

As predicted, permutation problems with symmetric sets of elements and permutation problems with asymmetric sets of elements were perceived by subjects as different types of problems. The different cover stories (objects and outcome) led subjects to construct either symmetric or asymmetric equations for the training problems and affected the way in which subjects applied the learned solutions to isomorphic problems "dressed up" in different contents. Our results show a baseline interpretative bias (preference to place objects rather than people in the role of the assigned set), mapping to example (placing similar elements in similar structural roles), and an interaction between the interpretative bias and mapping.

It is important to note that interpretative effects of content might either impair or enhance subjects' understanding of the correct structure of the problem. In our permutation problems, the two sets of elements had to play asymmetric roles because the problems specified that elements from one set were assigned to elements from the other set. Problems with two similar sets of people impaired subjects' ability to notice the asymmetry of the assignment relation, whereas problems with people and objects led subjects to an interpretation that was consistent with the asymmetry inherent in the assignment relation. However, the asymmetric outcome of assignment in "get" situations biased subjects to assign objects to people regardless of the direction of assignment stated in the problems.

The high level of interpretative bias in the OP and PO transfer conditions and the chance level of performance in the PP transfer condition suggest that our subjects had little understanding of the correct mathematical structure of the permutation problems, possibly because our problems made it difficult to decide which elements were the randomly assigned set. One might wonder whether interpretative effects of content would occur when subjects have a better understanding of the mathematical structure and/or when the direction of assignment is more transparent. Note that an extreme case of reliance on interpretative effects of content implies nonarbitrary guessing on OP and PO problems (choosing objects as the assigned set) and arbitrary guessing on PP problems (choosing one of the two sets of people as the assigned set).

The scope of interpretative effects of content and the extent to which such effects depend on expertise and on the clarity of the mathematical structure will have to be established in future work. We have some evidence indicating that interpretative effects occur even when the subjects are not forced into guessing. We analyzed the performance of the 56 subjects who constructed correct equations for the training problems and whose results were not included in the main analysis. These subjects showed a strong interpretative bias on the training problems: They solved correctly 80% of the OP problems but only 50% of the PO problems. Moreover, all of them engaged in mapping on the OP and PO transfer problems even though they solved correctly 85% of the PP transfer problems. Thus, even subjects who knew the correct equation before receiving training from us and could correctly determine which set of elements was the assigned set in the symmetric PP problems used the asymmetric outcome to determine the direction of assignment in the OP and PO problems.

EXPERIMENT 2

Experiment 2 was designed to unconfound predictions derived from the object-mapping hypothesis from predictions derived from the interpreted structure hypothesis. Specifically, this experiment used a hierarchically ordered triplet of element sets: carts, caddies, and golfers, with caddies serving as a variable in both the training and the transfer problems. All subjects initially learned a solution to a permutation problem in which caddies were assigned to golfers. In the solution of this problem, the number of caddies appeared in the equation (n = caddies). One group of subjects received a transfer problem in which caddies were assigned to carts. As in the training problem, the number of caddies in this transfer problem should appear in the equation (n = caddies). Another group of subjects received a transfer problem in which carts were assigned to caddies. Because in this problem the carts were the randomly assigned set, contrary to the training solution, the number of carts should now appear in the equation (n = carts).

According to the object-mapping hypothesis, subjects establish independent matches between each pair of objects in the base and the target problems. Clearly, caddies are more similar to caddies than to golfers. Using the animate/inanimate distinction, carts (objects) are as different from caddies (people) as they are from golfers (people), but they are more similar to caddies than to golfers because both carts and caddies are known to carry equipment. Given that identity matches (caddies to caddies) are stronger than matches in a subset of shared attributes (carts to caddies), the caddies in the transfer problem should win this competition. Note that caddies will also win the competition against carts as being more similar to golfers (people and people vs. objects and people). Although it is not obviously clear which of these two "relative wins" is stronger, identities should win over partial matches even in such relative decisions. Hence, object mapping should lead subjects to map caddies to caddies. Because in the training problem the caddies were the assigned set, subjects should perform better on the "caddies assigned to carts" than on the "carts assigned to caddies" transfer problem.

According to the interpreted structure hypothesis, however, subjects should perform better on the "carts assigned to caddies" transfer problem than on the "caddies assigned to carts" problem. First, due to a baseline interpretative bias that is unrelated to mapping, they should perform better on this problem, expecting that the carts (objects) should be the assigned set because the carts are the givens. Moreover, they should perform better on this problem because they could map the interpreted "get" structure: in the training problem, the givens appeared in the equation, and in the transfer problems, the givens were the carts.

In order to estimate the relative effects of object mapping and mapping of the interpreted structure, two additional groups of subjects received training on the "caddies assigned to golfers" problem and were tested on transfer problems dealing with students and prizes. Performance on problems with prizes and students cannot be explained by object mappings, because students are as similar to golfers as they are to caddies and prizes are as different from golfers as they are from caddies. Thus, these problems served to establish a baseline interpretative bias that is unrelated to mapping—a preference to assign objects to people (prizes to students) rather than vice versa (students to prizes). If subjects engage in object mapping on the caddies and carts transfer problem, their "interpretative" bias to assign carts to caddies (i.e., objects to people) should be smaller than in the case of prizes and students.

Method

Subjects

Eighty-one subjects from a population similar to that used in Experiment 1 participated in this experiment. As in Experiment 1, we report results obtained only from the subjects who failed to construct correct equations for the caddies and golfers training problem ($N = 69$, 36 females and 33 males). The subjects were paid for their participation in the study.

Materials

The training chapter was identical to the chapter used in Experiment 1. Three isomorphic permutation word problems that differed in their cover stories were constructed for the study. One problem (caddies assigned to golfers) served as the training problem in all conditions. The other two served as transfer problems: one involved caddies and carts, and the other, students and prizes. Each of the two transfer problems had two versions that reversed the direction of assignment. Also, each of these problems (one training and four transfer) had two versions that reversed the relative sizes of the two sets (n and m). Thus, there were 10 different permutation problems altogether. The training problem and two problems that reversed the direction of assignment of the caddies and carts transfer problem appear in Appendix B. The transfer problem involving students and prizes was used in Experiment 1; it appears in Appendix A.

Procedure

All subjects received the same training. They were randomly assigned to one of four experimental conditions according to the transfer problem they had to solve. The size of the two sets of elements was randomized across subjects. The number of subjects in each of the four transfer conditions ranged between 12 and 21. The large differences in the numbers of subjects per cell is due to uneven distribution of subjects who correctly solved the training problem (a total of 12 subjects) and therefore were excluded from the study. The experimental procedure was identical to that used in Experiment 1. Fifty-six subjects talked aloud while solving the training and the transfer problems (see again Footnote 4).

Results and Discussion

Equations Constructed for the Training Problem Before Instruction

Of 81 subjects, 12 constructed correct equations, 22 did not construct an equation, and 47 constructed erroneous equations. As in Experiment 1, we classified the erroneous equations as either symmetric or asymmetric. Of the 47 subjects who constructed erroneous equations, 51% constructed equations in which caddies and golfers played asymmetric roles (e.g., m/n) and 49% constructed equations in which caddies and golfers played symmetric roles (e.g., $1/(m \cdot n)$). The frequency of asymmetric equations constructed for golfers and caddies is lower than the frequency of asymmetric equations constructed for problems involving objects and people in Experiment 1 (51% vs. 87%) but is higher than the frequency of asymmetric equations constructed for problems involving symmetric sets of people in Experiment 1 (51% vs. 22%). This pattern of results is probably reflective of people's knowledge that, on the one hand, caddies and golfers play asymmetric roles and, on the other hand, both sets of elements are people and therefore could play symmetric roles.

Instantiation of the Learned Equation

Subjects' performance was not affected by differences in the size of the assigned and the receiving sets [$\chi^2(1) = 0.04$, n.s.], so we combined the results of subjects who received problems with reversed sizes of the two sets. Table 2 presents the percentage of correct solutions for the two directions of assignment on the "caddies and carts" and on the "students and prizes" transfer problems (i.e., problems with objects and people).

The data were analyzed using the logit model, with cover story (2) and direction of assignment (2) serving as factors. As predicted, there was a strong interpretative bias: Subjects correctly solved 94% of the problems in which objects were assigned to people, but were correct on only 21% of the problems in which people were assigned to objects [$\chi^2(1) = 21.06, p < .001$].

In addition to the overall interpretative bias to assign objects to people, subjects' performance on the two transfer problems that differed in their cover stories was virtually the same [$\chi^2(1) = 0.09$, n.s.], and there was no interaction between the problems' cover story and the direction of assignment [$\chi^2(1) = 0.07$, n.s.]. Because the training problem involved two sets of people, the biased performance on the "students and prizes" problem cannot be explained by object mapping, and the similarity in subjects' performance on the two transfer problems implies that subjects did not engage in object mapping on the "caddies and carts" problem either.

In fact, the most impressive result of Experiment 2 was the magnitude of the interpretative bias on the "carts and caddies" transfer problem. As predicted, subjects were highly successful when carts were assigned to caddies (94% correct), but their performance was very poor when caddies were assigned to carts (24% correct) [$\chi^2(1) = 19.67, p < .001$]. This bias occurred despite the fact that the caddies were the assigned set in the training problem. Thus, in sharp contrast to previous results showing that subjects apply the learned solution by mapping similar elements to similar variable roles (e.g., Gentner & Toupin, 1986; Ross, 1987, 1989), subjects in the present experiment did not attempt to map caddies into their role in the training problem. Instead, they were guided by what they knew about carts and caddies in general—that caddies get carts but carts do not get caddies—and instantiated the equation with the given set. Thus, when the experimental design did not confound similarities between the elements in the base and the target problems (e.g., being inanimate) and similarities in the interpreted roles of the elements (e.g., being the givens), matches in object attributes (caddies and caddies) did not win the "competition" against the interpreted structural roles of these objects.

It is interesting to note that some subjects did not merely assume that the direction of assignment was compatible with the "get" outcome; they actually replaced the "assign (manager, caddies, golfers)" three-place predicate with a "select (golfers, caddies)" two-place predicate: "It was like, the caddies are selecting the carts, the golfers select the caddies" (Subject 124). We do not know how many subjects replaced the "assign" relation with a "select" relation, or how many subjects merely assumed that the direction of assignment was consistent with the "get" outcome. Regardless of the particular nature of such interpreted structures, subjects were highly successful when the direction of assignment in the target problems was consistent with the natural roles of the elements with respect to each other and to erroneous performance when the direction of assignment reversed such natural roles. Guided by similarities in the interpreted structural roles of the elements in the training and the transfer problems, our subjects ignored direct matches in object attributes—they did not engage in object mapping.

Table 2
Percentage of Correct Solutions on the Four Transfer Problems in Experiment 2

Direction of Assignment	Cover Story: Prizes and Students % Corr.	*n*	Carts and Caddies % Corr.	*n*	Total % Corr.	*n*
Objects to People	94	18	94	18	94	36
People to Objects	17	12	24	21	21	33
Total	63	30	56	39	59	69

GENERAL DISCUSSION

Our results demonstrate that people use the particular content instantiations of problems to interpret the structures of problems. Specifically, our subjects interpreted the structure of permutation problems dealing with random assignment using knowledge about the structural roles that the entities play with respect to each other in the outcome of assignment. Knowing that the structural roles of doctors from two hospitals were symmetric whereas the structural roles of computers and secretaries were asymmetric, the subjects constructed equations that reflected the expected symmetry and asymmetry in the structural roles of these element sets. Because our subjects abstracted interpreted structures that differed from the structure that actually determined the problems' solutions, they were "successful" on problems in which the solution to the interpreted problems ("get") happened to coincide with the solution dictated by the mathematical structure (i.e., objects assigned to people), but arrived at erroneous solutions when the solution to the interpreted problems reversed the solution dictated by the mathematical structure (i.e., people assigned to objects). Moreover, they performed at chance level when the interpreted structure ("pair") was symmetric (people assigned to people) and therefore incompatible with the mathematical structure of the permutation problems.

Our results also show that when similarities between objects (animate vs. inanimate) were correlated with the structural roles of these elements in the interpreted "get" structure (receivers vs. givens) subjects applied previously learned solutions to isomorphic problems by placing similar entities in similar structural roles. However, when similarities between the objects (caddies and caddies) conflicted with the structural roles of these objects in the interpreted "get" structure (givens vs. receivers), subjects ignored direct similarities between the objects in the base and the target problems.

Object-Based Inferences

It is by now well established that people draw inferences on the basis of knowledge about nonarbitrary constraining dependencies between objects and relations. Such inferences are a natural byproduct of language comprehension (e.g., Anderson & Ortony, 1975; Gentner, 1981; Gentner & France, 1988; Ortony, 1979). In particular, people are aware of the constraining structural roles that various objects can, or tend to, play with respect to each other. For example, Strohner and Nelson (1974) found that even 2-year-olds know that cats chase mice rather than vice versa, misinterpreting sentences that switch such familiar roles. In many situations we consider such inferences as evidence for intelligent behavior. For instance, in an "analogies" test, we would credit people with intelligence points for choosing "apples: baskets" rather than "baskets:apples" or "apples:oranges" as the most appropriate analogy to "books:shelves."

Our results document that such inferences are drawn naturally during problem solving, and that people expect the mathematical structures of problems to correspond to the likely relations between the objects that serve as arguments in these structures. Note that in our problems the interpretative effects were not "purely" object based because the outcome question always supported such inferences. It remains to be established to what extent similar effects will occur when they are not supported by an outcome question. Recently we found evidence for such object-based inferences in a problem generation study. We presented college students with various pairs of element sets (e.g., flowers and vases) and asked them to construct simple word problems dealing with division (Bassok, 1993). As was the case for the probability problems used in the present experiments, elements known to play different roles (e.g., cars and mechanics) induced asymmetric scenarios that constrained the roles of elements from the two sets into the divided (numerator) and the dividing (denominator) roles. For example, the majority of subjects divided cars among mechanics rather than mechanics among cars. Moreover, when subjects were presented with two similar sets of elements (e.g., apples and oranges), they often created a common superset (e.g., fruit) and then asked about the relative proportion of one set in the combined set (e.g., oranges/fruit). Thus, although the subjects could rely only on semantic knowledge that was implied by the paired elements, spontaneous scenarios constructed by subjects for division problems with symmetric and asymmetric sets of elements were similar to the interpreted scenarios inferred for our permutation problems.

It is important to note that the specific direction of asymmetry between objects and people obtained in our experiments was implied by the particular sets of elements chosen by us for the assignment situations. Of course, other sets of elements (e.g., taxi cabs and passengers) imply scenarios that reverse the direction of asymmetry (people assigned to objects). In addition, some pairs of element sets induce more than one likely scenario. For example, in our division study, the direction of asymmetry between books and students depended on the inferred type of books: subjects divided textbooks among students but divided students among encyclopedias. Our claim is not that subjects always assign objects to people. Rather, we argue that it is important to consider the type of scenarios, and therefore structural constraints, that are activated by people's semantic knowledge about typical relations between the entities that appear in the problem. Clearly, much work is needed to establish which scenarios are triggered by various objects and how these scenarios interact with problem texts and with the "objective" problem structures.

Other than the work of Kotovsky et al. (1985), to which we referred in the introduction, we are not aware of any other line of research that has examined the effects of object-based inferences on transfer. Nonetheless, we believe that structural inferences triggered by object attributes are quite common in problem solving, and that such inferences are likely to affect transfer by affecting the representations of the base and the target problems. For example, our results are consistent with "erroneous" reversals of object roles by subjects solving "impossible" algebra word problems (Paige & Simon, 1966), where subjects switched the roles of quarters and dimes in a given problem to match their knowledge about the monetary values of these coins. Our results are also consistent with on-line inferences made by students while studying worked-out solutions to physics problems (Chi, Bassok, Lewis, Reimann, & Glaser, 1989). For instance, a student presented with a figure depicting a block hanging from three strings tied by a knot reads the following statement: "Consider the *knot* at the junction of the three strings to be the *body*." This statement does not fit her expectations about the functional roles of knots and blocks: "Why is the knot the body? And I thought that W (the block) was the body," leading her to repeated attempts to understand this "unreasonable" statement.

It appears that people's default reasoning about problems, regardless of whether they attempt to understand problem texts, figures, or physical displays, is affected by structural inferences implied by attributes of the specific objects that serve as arguments in such problems.

As argued by Greeno, Moore, and Smith (1993), people reason about problem situations in terms of "affordances" of the materials (object-attributes) and in terms of the "abilities" of the agents acting on these materials (object attributes). Such affordances and abilities might induce structural inferences and affect transfer.

Interdependence Between Content and Structure: Implications for Models of Transfer

People can rely on direct similarities in problem contents because they know that similar contents are likely to indicate similar structures. Such effects can occur even when it is obvious to the most naive subject that the specific content instantiation of the problem is superficial to the problem's solution. For example, Ross (1984) found that subjects were reminded of particular editing commands by the topic of the edited text that served as the training context. However, people know that very often content and structure are not merely correlated by chance and that such correlations might be well justified. Hence, as we have shown in the present study, people actually assume that the specific content instantiations of problems can assist them in interpreting the problems' structures. According to Medin and Ortony (1989), people's belief about the existence of nonarbitrary dependencies between appearance and essence serves two functions: "It enables surface similarity to serve as a good heuristic for where to look for deeper properties, *and* it functions as a constraint on the predicates that compose our mental representations" (p. 182). Both interpretative effects of content and reliance on similarities in "undeleted" aspects of content can serve the first heuristic function. Our experiments demonstrate the second function—that content constrains the type of predicates that compose our mental representations of problems.

People tend to use heuristic methods either when they do not have more valid rules or when such methods can serve as useful shortcuts that can speed up the processing relative to reliance on more certain rules or algorithms. Indeed, many studies document that people with better understanding of the conditions that determine problem solutions are more likely to ignore similarities and differences between content covers of analogous problems (e.g., Bassok & Holyoak, 1989, 1993; Brown, Kane, & Echols, 1986; Catrambone & Holyoak, 1989; Gentner & Toupin, 1986; Gick & Holyoak, 1983; Novick, 1988; Reed, 1987). Also, results from many studies support the notion that reliance on similarities in content of the base and the target problems serves as a useful shortcut for speeding up the retrieval and/or mapping process. Even people who have good understanding of the problem's structure rely on similarities in content during initial retrieval or mapping but spontaneously recover from erroneous analogies when they attempt to apply the retrieved solutions in a more mindful way (e.g., Bassok, 1990; Gentner & Ratterman, 1991; Gentner & Toupin, 1986; Gick & Holyoak, 1980; Novick, 1988).

Although research on analogical transfer provides ample evidence for the heuristic value of reliance on similarities in content, the possibility that content can impose constraints on our mental representations of problems received very little attention. In principle, interpretative effects of content could be incorporated into existing models of analogical transfer, because proponents of such models explicitly acknowledge that the feature-matching mechanisms should operate on relations and attributes that comprise the "psychologically relevant" representations of the base and target problems. Nonetheless, the validity of models that account for retrieval and mapping of analogous solutions is established on "uninterpreted" representations—on relational predicates that are assumed by the experimenters to be hidden behind "superficial" cover stories and objects that serve as arguments in such "uninterpreted" structures. Because the experimenters know that the problems share the same structure, they might be overlooking various interpretative effects of content.

We believe that lack of research on interpretative effects of content is not a mere error of omission. Rather, it appears that research on analogical transfer is guided by an implicit assumption that effects of content and effects of structure can be studied independently of each other. This assumption is implied by the dichotomy and the competition between content and structure in feature-matching models of transfer, and by the common methodology that independently varies similarities in aspects of content and structure to estimate their relative impact on retrieval and application (e.g., Gentner & Landers, 1985; Gentner & Rattermann, 1991; Holyoak & Koh, 1987; Ross, 1987, 1989).

In fact, it appears that the common methodology in transfer research attempts to "undo" and/or average across existing correlations between content and structure. For example, to ensure a clean separation between content and structure, Gentner and Toupin (1986) arbitrarily placed a variety of animals (e.g., a seal, a cat, a giraffe) into their structural roles in animated stories and subsequently averaged across any interpretative effects of content that might have occurred in each of the specific scenarios. Similarly, object-mapping effects in the studies of Ross (1987, 1989) were averaged across the two directions of assignment (people to objects and objects to people) because, in the "uninterpreted" mathematical structure, the objects and their structural roles were independent. Although such studies demonstrate effects of memory for specific aspects of content, they ignore interpretative effects that might be quite strong in natural situations. That is, in most cases people do not have to reason about animated stories in which mice can chase cats and foxes can fly, and in most cases people are justified in believing that those who present them with problem examples attempt to be informative rather than misleading (e.g., Grice, 1975) and therefore do not "undo" the correlations between content and structure.[5]

Our results are very robust, but our work should be seen mainly as a proof of existence for interpretative effects of content on transfer. We have demonstrated such effects using a handful of examples and a single problem structure, leaving many important questions unanswered. For example, it remains unclear how people induce scenarios from various combinations of objects, whether objects can induce interpretative effects without the "assistance" of explicit relations stated in the text (e.g., the outcome questions), to what extent interpretative effects of content will occur when the problem's text provides sufficient cues to disambiguate the problem's structure, or whether interpretative effects of content can affect access as well as mapping. Of course, much more work is needed to determine how memory and interpretation interact with each other. Such questions will be answered by studies that explicitly attempt to discover the rules by which people exploit nonarbitrary dependencies between semantic and structural aspects of problems.

REFERENCES

ANDERSON, R. C., & ORTONY, A. (1975). On putting apples into bottles—a problem of polysemy. *Cognitive Psychology*, **7**, 167-180.

BASSOK, M. (1990). Transfer of domain-specific problem solving procedures. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **16**, 522-533.

BASSOK, M. (1993, November). *On putting baskets into apples: Effects of world knowledge on problem solving*. Paper presented at the meeting of the Psychonomic Society, Washington, DC.

BASSOK, M., & HOLYOAK, K. J. (1989). Interdomain transfer between isomorphic topics in algebra and physics. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 153-166.

BASSOK, M., & HOLYOAK, K. J. (1993). Pragmatic knowledge and conceptual structure: Determinants of transfer between quantitative domains. In D. K. Detterman & R. J. Sternberg (Eds.), *Transfer on trial: Intelligence, cognition, and instruction* (pp. 68-98). Norwood, NJ: Ablex.

BROWN, A. L., KANE, M. J., & ECHOLS, C. H. (1986). Young children's mental models determine analogical transfer across problems with a common goal structure. *Cognitive Development*, **1**, 103-121.

CATRAMBONE, R., & HOLYOAK, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 1147-1156.

CHI, M. T. H., BASSOK, M., LEWIS, M. W., REIMANN, P., & GLASER, R. (1989). Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, **13**, 145-182.

DUNCKER, K. (1945). On problem solving. *Psychological Monographs*, **58**(Whole No. 270).

GENTNER, D. (1981). Some interesting differences between verbs and nouns. *Cognition & Brain Theory*, **4**, 161-178.

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

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

GENTNER, D., & FORBUS, L. K. (1991). MAC/FAC: A model of similarity-based access and mapping. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 504-509). Hillsdale, NJ: Erlbaum.

GENTNER, D., & FRANCE, I. M. (1988). The verb mutability effect: Studies of the combinatorial semantics of nouns and verbs. In S. L. Small, G. W. Cotrell, & M. Tenenhaus (Eds.), *Lexical ambiguity resolution in the comprehension of human language* (pp. 343-382). Los Altos, CA: Morgan Kaufman.

GENTNER, D., & LANDERS, R. (1985). Analogical reminding: A good match is hard to find. In *Proceedings of the International Conference on Systems, Man, and Cybernetics* (pp. 607-613). New York: IEEE.

GENTNER, D., & RATTERMANN, M. J. (1991). Language and the career of similarity. In S. A. Gelman & J. P. Byrnes (Eds.), *Perspectives on thought and language: Interrelations in development* (pp. 225-277). London: Cambridge University Press.

GENTNER, D., & TOUPIN, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, **10**, 277-300.

GICK, M. L., & HOLYOAK, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, **12**, 306-355.

GICK, M. L., & HOLYOAK, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, **15**, 1-38.

GREENO, J. G., MOORE, J. L., & SMITH, D. R. (1993). Transfer of situated learning. In D. K. Detterman & R. J. Sternberg (Eds.), *Transfer on trial: Intelligence, cognition, and instruction* (pp. 68-98). Norwood, NJ: Ablex.

GRICE, H. P. (1975). Logic in conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics: Vol. 3. Speech acts* (pp. 41-58). New York: Seminar Press.

HAYES, J. R., & SIMON, H. A. (1977). Psychological differences among problem isomorphs. In N. J. Castellan, Jr., D. B. Pisoni, & G. R. Potts (Eds.), *Cognitive theory* (Vol. 2, pp. 21-41). Hillsdale, NJ: Erlbaum.

HINSLEY, D. A., HAYES, J. R., & SIMON, H. A. (1977). From words to equations: Meaning and representation in algebra word problems. In M. A. Just & P. A. Carpenter (Eds.), *Cognitive processes in comprehension* (pp. 89-106). Hillsdale, NJ: Erlbaum.

HOLYOAK, K. J., & KOH, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, **15**, 332-340.

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

KOTOVSKY, K., HAYES, J. R., & SIMON, H. A. (1985). Why are some problems hard?: Evidence from Tower of Hanoi. *Cognitive Psychology*, **17**, 248-294.

MEDIN, D. L., GOLDSTONE, R. L., & GENTNER, D. (1993). Respects of similarity. *Psychological Review*, **100**, 254-278.

MEDIN, D. L., & ORTONY, A. (1989). Psychological essentialism. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 179-195). Cambridge: Cambridge University Press.

NOVICK, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 510-520.

ORTONY, A. (1979). Beyond literal similarity. *Psychological Review*, **86**, 161-180.

PAIGE, J. M., & SIMON, H. A. (1966). Cognitive processes in solving algebra word problems. In B. Kleinmuntz (Ed.), *Problem solving: Research, method, and theory* (pp. 51-119). New York: Wiley.

REED, S. K. (1987). A structure mapping model for word problems. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 124-139.

ROSS, B. H. (1984). Remindings and their effects in learning a cognitive skill. *Cognitive Psychology*, **16**, 371-416.

ROSS, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **13**, 629-639.

ROSS, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **15**, 456-468.

SMITH, L. B., & HEISE, D. (1992). Perceptual similarity and conceptual structure. In B. Burns (Ed.), *Percepts, concepts, and categories: The representation and processing of information* (pp. 233-272). New York: Elsevier.

STROHNER, H., & NELSON, K. E. (1974). The young child's development of sentence comprehension: Influence of event probability, nonverbal context, syntactic form, and strategies. *Child Development*, **45**, 567-576.

THAGARD, P. H., HOLYOAK, K. J., NELSON, G., & GOCHFELD, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, **46**, 259-310.

NOTES

1. The Hebrew version of the warning "Don't judge the book by its cover" is "Don't look at the jar but rather at what it contains." Interestingly, Anderson and Ortony (1975) demonstrated the validity of the

Hebrew version by showing its opposite yet complementary use: People rely on the contained stuff (e.g., coke vs. apples) to infer the type of its container (a bottle or a basket).

2. We want to thank Brian Ross for providing us with his experimental materials.

3. Because we report results that refer to an instructional situation, we use the labels "training" and "transfer" problems instead of the more general labels "base" and "target" problems that are typically used in the literature on analogical problem solving and which we have been using in the introduction.

4. The present experiment was conducted in conjunction with other experiments, and talk-aloud protocols were collected only for subjects who gave protocols in the other experiments. The performance of subjects who did and who did not talk aloud was similar, and we combined the results obtained from both groups. The lack of difference between these two groups is not surprising since the majority of subjects who gave protocols merely restated the problems. However, because there were very few meaningful protocols, we could not analyze the data from the protocols in a quantitative way. Hence, the protocol excerpts cited in the text should be viewed only as illustrative examples.

5. Frequent undoing of correlations between content and structure in a particular domain teaches people to treat the content in this domain as something superficial. For instance, many students learn to ignore the semantic information implied by the content of mathematical word problems and end up giving ridiculous answers as solutions (e.g., fractions of buses needed for a school trip).

APPENDIX A
EXAMPLES OF ISOMORPHIC PERMUTATION PROBLEMS IN EXPERIMENT 1

Training

Computers and Secretaries (Objects Assigned to People)

In a big publishing company, some secretaries will get to work on new personal computers. The company received a shipment of 21 computers, with serial numbers in a running order: from 10075 through 10095. There are 25 secretaries in this company that would like to work on a new computer. The names of the secretaries are listed in order of their work experience, from the most experienced secretary to the least experienced one. *The manager of the company randomly assigns computers to secretaries* according to the work experience of the secretaries. What is the probability that the three most experienced secretaries will get to work on the first three computers (10075, 10076, and 10077), respectively?

Doctors and Doctors (People Assigned to People)

In a medical meeting, doctors from a Minnesota Hospital will get to work in pairs with doctors from a Chicago Hospital. There is a list of 20 doctors from Chicago, arranged in alphabetical order. There are 16 doctors from Minnesota that would like to work with the doctors from Chicago. The names of the Minnesota doctors are listed in the order of their social security numbers, from highest to lowest. *The chairman of the meeting randomly assigns doctors from Minnesota to doctors from Chicago* according to the alphabetical order of the Chicago doctors. What is the probability that the first three doctors on the Minnesota Hospital's social security number list will get to work with the first three doctors on the Chicago Hospital's alphabetical list, respectively?

Transfer

Prizes and Students (People Assigned to Objects)

In a high school awards ceremony some students will receive prizes. There are 26 wrapped prizes, marked with the numbers one through 26. There are 30 honor students who would like to receive a prize in the ceremony. The names of the students are listed in the order of their GPA (grade point average), from highest to lowest. *The president randomly assigns students to prizes* according to the number order of the prizes. What is the probability that the first three students on the GPA list will receive the first three prizes (no. 1 through 3), respectively?

Children and Children (People Assigned to People)

In a certain after-school program, kids from Sweety nursery school will get to work in pairs with kids from Paradise nursery school. There is a list of 18 kids from Paradise, arranged in order of the kids' height: from tallest to shortest. There are 22 kids from Sweety that would like to work with the kids from Paradise. The names of the Sweety kids are listed in order of their age: from oldest to youngest. *The director of the after-school program randomly assigns kids from Paradise to kids from Sweety* according to the age of the Sweety kids. What is the probability that the three oldest kids from Sweety will get to work with the three tallest kids from Paradise, respectively?

APPENDIX B
TRAINING PROBLEM AND ONE EXAMPLE OF A TRANSFER PROBLEM IN EXPERIMENT 2

Training: Caddies Assigned to Golfers

At a country club, each golf player would like to receive a caddy to carry his clubs. There are 16 players, listed in the order of their membership, from the most recent to the oldest member. There is a list of 22 caddies, ordered according to their experience, from the most experienced to the least experienced. *The manager of the golf-course randomly assigns caddies to players* according to the membership order of the players. What is the probability that the three golf players with the most recent membership will get the three most experienced caddies, respectively?

Transfer: Carts Assigned to Caddies

At a country club, each caddy would like to receive a cart to carry clubs. There are 24 caddies, listed in the order of their experience, from the most experienced to the least experienced. There is a list of 18 carts, ordered from 001 to 018. *The manager of the golf-course randomly assigns carts to caddies* according to the experience of the caddies. What is the probability that the three most experienced caddies will get the first three carts (001, 002, 003), respectively?

Transfer: Caddies Assigned to Carts

At a country club, each caddy would like to receive a cart to carry clubs. There are 24 caddies, listed in the order of their experience, from the most experienced to the least experienced. There is a list of 18 carts, ordered from 001 to 018. *The manager of the golf-course randomly assigns caddies to carts* according to the number order of the carts. What is the probability that the three most experienced caddies will get the first three carts (001, 002, 003), respectively?

Note—Another problem involving students and prizes was used in Experiment 1 and appears in Appendix A.

(Manuscript recieved February 17, 1994;
revision accepted for publication June 23, 1994.)