Gentner, D., Loewenstein, J., Thompson, L., & Forbus, K. (2009) Reviving inert knowledge: Analogical abstraction supports relational retrieval of past events. *Cognitive Science*, 3, 1343-1382.

Cognitive Science 33 (2009) 1343–1382

ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/j.1551-6709.2009.01070.x

Reviving Inert Knowledge: Analogical Abstraction Supports Relational Retrieval of Past Events

Dedre Gentner,[a] Jeffrey Loewenstein,[b] Leigh Thompson,[c] Kenneth D. Forbus[d]

[a]*Department of Psychology, Northwestern University*
[b]*McCombs School of Business, The University of Texas at Austin*
[c]*Kellogg School of Management, Northwestern University*
[d]*Electrical Engineering and Computer Science Department, Northwestern University*

Received 26 October 2007; received in revised form 13 February 2009; accepted 4 May 2009

Abstract

We present five experiments and simulation studies to establish *late analogical abstraction* as a new psychological phenomenon: Schema abstraction from analogical examples can revive otherwise inert knowledge. We find that comparing two analogous examples of negotiations at recall time promotes retrieving analogical matches stored in memory—a notoriously elusive effect. Another innovation in this research is that we show parallel effects for real-life autobiographical memory (Experiments 1–3) and for a controlled memory set (Experiments 4 and 5). Simulation studies show that a unified model based on schema abstraction can capture backward (retrieval) effects as well as forward (transfer) effects.

Keywords: Analogical mapping; Memory retrieval; Relational structure; Schema abstraction; Inert knowledge

1. Introduction

The ability to apply relational knowledge across different contexts is of central importance in human cognition. One instance of this kind of ability is analogical transfer—mapping knowledge from a prior stored situation to a current situation (Gentner, 1983; Holyoak, Gentner, & Kokinov, 2001). There is abundant research demonstrating, first, that analogical transfer can lead to considerable insights *when it occurs*; and, second, that it very often fails to occur. When people succeed in accessing an appropriate prior example, they typically perform well in mapping the solution to the current problem (Pirolli & Anderson, 1985; Reed, 1987; Ross, 1989). However, people often fail to access prior cases in new

Correspondence should be sent to Dedre Gentner, Department of Psychology, Northwestern University, Evanston, IL 60208. E-mail: gentner@northwestern.edu

contexts (Gick & Holyoak, 1980, 1983; Reeves & Weisberg, 1994; Ross, 1987, 1989). Although people do show some degree of spontaneous relational transfer (Gentner, Rattermann, & Forbus, 1993; Johnson & Seifert, 1992), most remindings to prior situations appear to be driven largely by surface similarities, such as similar characters and settings, rather than by similarities in relational structure (Brooks, 1987; Brooks, Norman, & Allen, 1991; Catrambone, 2002; Gentner et al., 1993; Gick & Holyoak, 1980, 1983; Ross, 1984, 1987, 1989). This means that people often fail to apply past learning to new situations that share the same causal or mathematical principles.

This failure of relational transfer appears less severe for experts than for novices (Blanchette & Dunbar, 2001; Novick, 1988), suggesting that experts are more likely to encode situations in ways that can be relationally retrieved. However, this only highlights the question of how expertise is attained, given that novices often fail to access the very kind of prior knowledge that would favor achieving new insights. Relational retrieval is important not only as part of cognitive theory but also because the failure of appropriate retrieval—the *inert knowledge* phenomenon—is a major challenge in research on learning, education, and conceptual change (Barnett & Ceci, 2002; Bransford, 1979; Reeves & Weisberg, 1994; Weisberg, 1995). The ability to access relationally similar examples from one's own long-term memory is important for effective functioning in a broad range of contexts, both in education and in the workforce (Chipman, Segal, & Glaser, 1985; Pfeffer & Sutton, 2000).

Because of the importance of relational transfer, much research has gone into the question of what kinds of learning experiences promote relational retrieval. Conditions that promote intensive encoding of the initial examples, such as self-explanation (Chi & VanLehn, 1991) or assigning stories to a list of available themes (Seifert, McKoon, Abelson, & Ratcliff, 1986), can sometimes promote relational retrieval, although research on this point is inconsistent (see Weisberg, 1995). However, the best-established way of promoting relational transfer is for the learner to compare analogous examples during learning (Catrambone & Holyoak, 1989; Gentner, Loewenstein, & Thompson, 2003; Gick & Holyoak, 1983; Reeves & Weisberg, 1994; Ross & Kennedy, 1990; Seifert et al., 1986, Experiments 1 and 2). Comparison between two analogous examples acts to make their common relational structure more salient (Gentner & Medina, 1998; Gentner & Namy, 1999; Markman & Gentner, 1993). As Gick and Holyoak (1983) demonstrated, comparing two analogous examples can facilitate deriving a schema, which in turn can facilitate transfer to a structurally similar problem. By contrast, people who study cases one at a time (as is often the case during learning) tend to encode them in a more concrete, context-specific manner, with the result that later remindings are often based on surface similarities (Gentner et al., 1993; Holyoak & Koh, 1987; Loewenstein, Thompson, & Gentner, 1999).

There is a substantial body of evidence that comparing two analogous examples highlights their common relational structure (Catrambone & Holyoak, 1989; Gentner & Christie, 2007; Gentner & Clement, 1988; Gentner & Markman, 1997; Gentner & Namy, 1999; Gick & Holyoak, 1983; Markman & Gentner, 1993). For example, when people are asked to discuss *cigarettes*, they write something like ''cylindrical, paper-wrapped, filled with tobacco,'' whereas when people compare *cigarettes* and *time bombs*, they write something like ''They both do their damage after a period in which no damage is evident'' (Gentner &

Clement, 1988). Thus, comparison can lead learners to focus on deep relational commonalities rather than on specific, potentially idiosyncratic features of the particular examples (Gentner & Medina, 1998; Gentner & Namy, 1999; Ross & Kennedy, 1990). This could influence future transfer in at least two specific ways—*relational schema abstraction* and/or *learning to encode*.

On the *relational schema abstraction* account, the relational structure engendered by an analogical comparison is stored in memory (either instead of or along with the representations of the initial exemplars). This relational schema will be a fairly concentrated relational representation, with many of the initial item-specific features stripped away. Memory thus consists of representations at varying degrees of abstractness (Hintzmann, 1986; Markman, 1999). As memory retrieval hinges on the match between probe and stored item, the likelihood of retrieval increases when nonmatching features are removed (Forbus, Gentner, & Law, 1995; Hummel & Holyoak, 1997; Ross, 1989). On this account, relational transfer is more likely after comparison because a future relationally similar example will have a stronger match with the relational schema than it will with a specific example, which may differ in surface features from the probe.

On the *learning-to-encode* account, the relational highlighting that results from analogical comparison does not result in a stored abstraction but rather influences future transfer by altering the way in which future examples are encoded (Medin & Ross, 1989). That is, new examples will tend to be encoded using the relational insights derived from the comparison. This learning-to-encode account will lead over time to greater uniformity in relational encoding, which will tend to promote relational matches between new items and memory items (Forbus et al., 1995), and thus more relational transfer.

It seems highly likely that learning to encode is at least part of the explanation. There is independent evidence that learners change their encoding patterns as domain knowledge increases (Chi, Feltovich, & Glaser, 1981; Forbus & Gentner, 1986; Larkin, McDermott, Simon, & Simon, 1980; Medin & Ross, 1989). Another point in favor of the learning-to-encode account is that it naturally predicts the inert knowledge effect because it predicts low overlap between new encodings and old encodings made prior to the analogical insight. Thus, parsimony might favor adopting the learning-to-encode account as the sole explanation for the gains due to analogical learning.

However, the relational schema account makes a novel and testable prediction. On this account the comparison benefit stems from storing an abstract relational schema. This schema will differ from an example in two ways: It will have fewer potentially mismatching surface features, and its relations will gain more weight. (We expand on this explanation later, when discussing the simulation studies.) Hence, the schema, relative to an example, should form a stronger relational match with future analogs. As what counts is the degree of match between probe and memory item, the effect of relational schemas should apply equally well *whether they serve as memory items or as probes.* Thus, the schema resulting from comparing two analogs should not only increase ''forward'' transfer, but it should also increase ''backwards'' relational retrieval. This is the specific prediction we test here: that comparing two analogs at recall time (what we could think of as ''late analogical abstraction'') will result in more relational retrieval than that produced by either or both of the separate analogs.

The following experiments test the effects of late analogical abstraction on relational retrieval. We examine relational retrieval both by asking people to retrieve examples from their own autobiographical memories (Experiments 1–3) and by conducting a more standard memory test in which we provided the memory set (Experiments 4 and 5). We chose real-life memory as our main focus for two reasons. First, this allows for the possibility that analogical abstraction can ''reach backward'' over time spans of months or years—potentially far longer than those generally tested in laboratory experiments; second, the comparison with Experiments 4 and 5 allows us to ask whether our laboratory phenomena do in fact generalize to real-life learning.

We then present computational simulations to test our chief theoretical claim: that the schema resulting from comparing two analogs will increase relational retrieval whether it serves as a memory item (leading to transfer forward in time) or as a probe (leading to retrieval backwards in time). As a proxy for autobiographical memory, we used a large, varied set of examples—to our knowledge the largest set yet used in simulating human analogical retrieval. We ask whether the simulation parallels human performance in backwards retrieval and forward transfer, when given either single items or analogically derived schemas.

2. Overview of the experiments

In the current studies, we test whether comparing cases at recall time can facilitate relational retrieval of prior cases from long-term memory. Four studies examine people learning negotiation principles. Negotiation is an apt arena in which to investigate retrieval and transfer, for several reasons. First, negotiation principles apply across many different contexts, making it likely that participants have stored varied exemplars of these principles over time. Second, in this arena one can find highly motivated participants. In most of our studies, the participants are taking negotiation classes (often at significant personal cost), with a direct interest in raising their job effectiveness. Third, the learning must be applied in many different contexts, making transferability important to the learners.

The negotiation principle we focus on here is the idea of constructing contingent contracts: agreements whose terms depend on the outcome of a future event (Bazerman & Gillespie, 1999). These contracts allow people to reach agreement despite differences in opinion. Despite their usefulness, novice negotiators tend not to form contingent contracts, and typically continue to miss opportunities to form them after they have been introduced to the idea. People generally know subtypes of contingent contracts (e.g., bets and options contracts) but do not realize that it is a general category performing a general function in negotiations, and frequently fail to recognize their usefulness in nonroutine negotiation situations.

In the first three experiments, we test whether analogical abstraction of examples of contingent contracts at recall time facilitates relational retrieval of prior instances of contingent contracts from participants' own long-term memories. Such a ''late analogical abstraction'' effect would be impossible to explain on a pure learning-to-encode account. It would be clear evidence that relational schemas—even those formed after the fact—can facilitate relational retrieval. In Experiments 4 and 5, we extend these findings by using a controlled set of initial examples and by examining effects on knowledge transfer. Finding a late analogical abstraction effect in these studies would show that the account generalizes over laboratory-induced memory sets as well as naturalistic long-term knowledge.

3. Experiment 1

In Experiment 1, we tested whether analogical abstraction facilitates both forward relational transfer and backwards relational retrieval. If so, this would suggest that comparison can result in an abstract relational schema, and that this relational schema is a better match (relative to a single exemplar) not only to relationally similar potential future cases but also to relationally similar past stored exemplars.

The participants were management consultants, which means that their professional role is to draw on their prior experiences to solve current problems. They were highly motivated learners engaged in a professional education program to learn negotiation to perform their jobs more effectively. On average they had about 15 years of work experience. We examined whether they could learn a general negotiation contract structure, contingent contracts, from written examples. We gave all participants two analogous examples of negotiations in which parties formed contingent contracts (see Appendix A). We asked half the participants to compare them, and half to study the examples one at a time. Next, to test whether comparison yields an advantage for relational retrieval, we asked the consultants to recall an example from their own experience that illustrated the same principle as the initial examples. Finally, to test for transfer, as in our prior research (Loewenstein et al., 1999), we assigned participants to roles for a simulated test negotiation and examined whether they would use the principle to form better agreements. The test negotiation was analogous to the study cases (although all three cases were different on the surface) and was constructed so that using the studied principle would lead to a greater net gain across both participants.

3.1. Method

3.1.1. Participants

There were 124 participants, aged approximately 30–45 years. There were 64 participants in the comparison condition and 60 in the separate-case condition. All were full-time professional management consultants working at the same organization, who participated as part of a multiday negotiation training program. Because of the wide variety of experience involved in management consulting, this group had a large range of prior cases to draw on.

3.1.2. Materials and procedure

Participants read role materials to prepare to be either the buyer or seller in a negotiation role-playing scenario. Just prior to engaging in the role play, participants received a training packet. The first page concerned details about their upcoming negotiation. The next pages contained two cases exemplifying the contingent contract principle as shown in the Appendix. The comparison group read both cases and then was asked, ''What is going on in these negotiations? Think about the similarities between these two cases. What are the key parallels in the two negotiations? Please describe the solution and say how successful you think it is.'' The separate-case group received the following question after each case: ''What is going on in this negotiation? Please describe the solution and say how successful you think it is.'' The answers to these questions—participants' case descriptions—provide a measure of participants' understanding of the contingent contract principle.

For both groups, the next page of the training packet asked participants to recall an example like those they had just read: ''Please think of an example, preferably from your own experience, that embodies the same principle as that on the previous page.'' We then asked each participant to state the source of their example: ''own experience, colleagues/partners, popular press, prior class, other.'' Participants were assured of confidentiality and encouraged to use pseudonyms if desired. The answers to this question provide a measure of participants' retrievals of examples of contingent contracts (i.e., of their relational retrievals). Participants were not limited as to time and spent roughly 45 min on these three pages (an indication of their motivation to learn). We saw no time-on-task differences by condition. Then each participant was paired with another participant from the same condition in a face-to-face negotiation. The negotiation case was set in a context different from that in the training cases, and was designed such that a mutually beneficial contingent contract was an optimal solution for the test negotiation. Hence, forming (or not forming) a contingent contract in their negotiations provides a measure of knowledge transfer.

3.1.3. Scoring

The primary purpose of the scoring was to assess whether participants' responses captured the essential elements of the contingent contract: namely that there is a disagreement about a future event, and that the terms of the agreement vary depending on the outcome of that future event. Participants' negotiated agreements were scored as including a contingent contract if the terms of the written agreement were dependent on a future event (as opposed to fixed). Participants' case descriptions required more extensive coding. Two blind coders rated the degree to which the contingent contract principle was described, using a three-point scale: 0 = no elements of the principle were present, 1 = some elements, and 2 = all elements. Separate-case participants were assigned the maximum score they received for either example. The coders also rated whether participants linked the cases either by comments such as ''both cases involve…'' or by mentioning the first when writing about the second (as a manipulation check on the condition difference).

The coders also rated whether the examples participants recalled from their own memory were contingent contracts, using the same scale as above. In addition, they categorized the source domains in which participants' examples were set and whether the examples were literally similar to the initial cases (both surface and structurally alike), analogous (only structurally alike), or shallow matches (some surface similarities but only fragments of the common structure). To reduce coding biases, the ratings for the initial descriptions and the recalled examples were performed separately. Overall, there was high agreement on the coding (87%, min $\kappa = 0.72$); disagreements were resolved through discussion.

3.2. Results

To summarize the results, as predicted, the comparison group scored higher than the separate-case group on all three measures: quality of case descriptions (a measure of principle understanding), quality of contract (a measure of transfer), and quality of retrievals from long-term memory (a measure of relational retrieval). Specifically, comparison participants ($M = 1.45$) articulated the contingent contract principle in their case descriptions better than separate-case participants ($M = 0.98$), $t[122] = 2.97$, $p < .01$[1]; see Table 1). Not surprisingly, participants in the comparison condition (88%) were highly likely to link the two cases; however, only 17% of the participants in the separate-case condition did so, despite the fact that the cases were presented on consecutive pages. Comparison participants also exhibited greater knowledge transfer. In their face-to-face negotiations, pairs of comparison participants (59%) were nearly twice as likely to make contingent agreements as were separate-case participants (33%) ($\chi^2[1, N = 62] = 4.22$, $p < .05$). This bears out the finding of our previous research that comparison facilitates transfer (Fig. 1).

The key novel finding concerned memory retrieval: comparison participants ($M = 1.25$) were more likely to retrieve good examples of contingent contracts than were separate-case participants ($M = 0.82$) ($t[122] = 2.65$, $p < .01$). This is consistent with the relational schema abstraction account, whereby (a) analogical comparison yields a more abstract representation than does separate encoding of the cases; and (b) this more abstract representation is more likely to match a variety of prior stored instances of the contingency contract. A pure learning-to-encode account could not predict any such advantage of comparison on retrieval of prior items.

Table 1
The mean (*SD*) score (maximum 2) and the raw number of retrievals and principles scored as *full* (all elements of the principle present), *partial* (some elements present), or *none* (no elements present)

	Experiment 1		Experiment 2		Experiment 3	
	Compare Condition	Separate-case Condition	Compare Condition	Separate-case Condition	Compare Condition	Separate- and Single-case Conditions
Principles	1.45†† (0.85)	0.98 (0.91)	1.42†† (0.84)	0.53 (0.72)	1.33†† (0.79)	0.81 (0.81)
Full	44**	24	12**	2	19**	16
Partial	5	11	3	5	10	20
None	15	25	4	10	7	28
Retrievals	1.25†† (0.89)	0.82 (0.93)	1.47†† (0.90)	0.59 (0.94)	1.08†† (0.97)	0.50 (0.76)
Full	35*	21	14**	5	18*	10
Partial	10	7	0	0	3	12
None	19	32	5	12	15	42

Notes: The daggers (†) refer to the *F*- or *t*-tests for the mean principle or retrieval ratings, and the asterisks (*) to the chi-square analyses for the 2 (comparison vs. separate) × 3 (score: full, partial, none) tables for numbers of principles or retrievals at each scoring level.
*$p < .05$. †† or **$p < .01$.

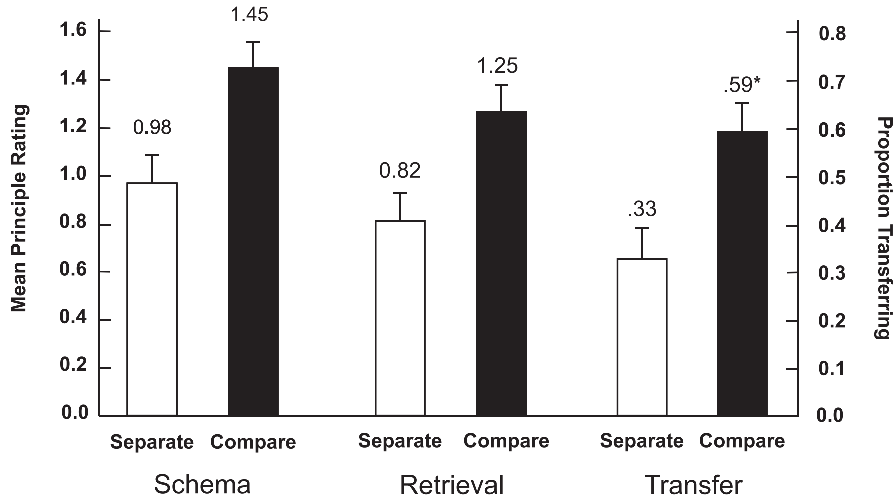

Fig. 1. Results of Experiment 1: Mean principle ratings, mean retrieval ratings, and the proportion of negotiated agreements using the principle.

Nearly all the 110 items that participants recalled were examples of negotiations, implying that the participants understood the requirements and could make at least a reasonable first-pass match (see Tables 2 and 3). Participants retrieved many analogies—52% of all responses. A few participants (6%) retrieved examples literally similar to our training examples. Participants also frequently (30%) retrieved shallow matches. These were not full contingent contracts but shared some part of their structure: They were stories about negotiations with unusual outcomes, negotiations in which trust or timing was critical, or negotiations in which there were disagreements about the future. Interestingly, 11% of the participants—highly motivated management consultants—wrote that they could not think of a good match, suggesting that they preferred to write nothing rather than something clearly inadequate.

A majority of the examples (66%) were from the participants' own experience or that of a colleague. A minority (17%) of the examples were from the popular press – all of which we were able to verify. The remainder of the participants (16%) did not specify their source.

Table 2
Types of Retrievals by condition in Experiment 1

	Similarity to the Study Cases		Proportions of Example Types Retrieved by Condition		
Type of Retrieval	Surface	Structural	Separate-case	Comparison	Total
Literal similarity	Yes	Yes	5% (3/60)	8% (5/64)	6% (8/124)
Analogy	No	Yes	42% (25/60)	63% (40/64)	52% (65/124)
Shallow matches	Yes	Partial	40% (24/60)	20% (13/64)	30% (37/124)
No reminding	–	–	13% (8/60)	9% (6/64)	11% (14/124)

Notes: Literal similarity matches include higher order relational structure, first-order relations, and some object attributes. Analogy matches include higher order and first-order relations. Shallow matches include first-order relations and some object attributes, but not higher order structure.

Table 3
Condensed examples of retrievals by type in Experiment 1 (all retrievals are to specific cases)

Type of Retrieval	Samples
Literal similarity	Travel by plane or boat? Which is more reliable? Bet on who pays for tickets contingent on weather
	Will cabinets be delivered and done on time? Penalty if not
Analogy	Value billing: low fee for services, but if services yield good results by objective measure, get paid fees proportional to gains
	Launching a website and variable payment to company, ''putting our money where our mouth is''
Shallow matches	A client will fail if no change is made, the change is made and they do better
	Helping a client find money in their budget to continue a project with an otherwise uncertain future
	Working out the timely shipping of an ordered product

The examples were mainly from the business domain (65%), with the rest drawn from daily life, such as betting on sports teams or uncertainty about the weather affecting a vacation activity. There were 19 participants (12 from the comparison condition) who recalled instances of *value billing*, an arrangement whereby a consulting firm bills clients a low base fee with a sizeable bonus structure based on the outcome of their work for their clients. It is a specialized and codified kind of contingent contract, but consultants do not refer to it as a contingent contract, and it appears that recognizing it as such requires a degree of insight.

3.2.1. Cross-measure associations

For the two tasks people did singly—writing case descriptions and writing retrievals—we found that the more highly rated the case description, the more highly rated the retrieval, $r_s = .18$[2] ($p < .05$). Investigating the association between the individual task outcomes (case descriptions and retrievals) and the negotiation task outcomes (which involved pairs of people) required some aggregation. Pairs in which at least one member had generated a completely matching retrieval (23 of 40, or 58%), were twice as likely to form a contingent contract as their negotiation task outcome than the remaining pairs (6 of 22, or 27%) ($\chi^2[1, N = 62] = 5.21$, $p < .05$). Despite showing a similar trend, negotiation task outcomes were not reliably linked to pairs in which at least one member had generated a complete contingent contract principle in their case descriptions (51%) and the remaining pairs (31%).

As a way of relating all three tasks, we considered the likelihood of finding high-performance pairs. A *high-performance pair* was defined as one in which at least one member wrote a complete principle, at least one retrieved a contingent contract, and the pair formed a contingent contract agreement in their negotiation. Multiplying the post hoc probabilities of performing each of these three tasks individually (79% wrote the contingent contract principle, 65% retrieved a contingent contract example, and 47% formed a contingent contract) yields the null hypothesis prediction that if there were no association between tasks, just 15 of the 62 pairs (or 24%) should fall into the high-performance group. We found that 21 pairs (34%) were in this group, a number greater than chance by a binomial test, $p < .05$, bearing out a link between the three tasks. Interestingly, reliably more comparison-condition pairs (47%) than separate-case pairs (20%) were high performing (χ^2[1, N = 62] = 4.99, $p < .05$). Thus, not only is performing well on the three tasks themselves linked, but jointly they are linked to having engaged in analogical comparison.

3.2.2. Distinguishing retrieval from invention

Before concluding that there is a comparison advantage in retrieval, it is important to assess whether participants were actually retrieving examples or were instead constructing them during the task. That is, it is possible that the retrieval advantage for the comparison group could stem simply from their using the derived principle to invent or reconstruct examples, rather than from recalling them. We offer two plausibility arguments for the time being; we return to these issues in Experiment 4. First, the proportion of structural remindings was the *lowest* (31%) for those failing to state a source, intermediate (45%) for those who stated nonverifiable sources, and the highest (68%) for those whose source we were able to verify. The opposite would have been expected on the fabrication account. Second, the difference between conditions among those whose source was verified (compare: 80%; separate-case: 56%) was comparable with the overall difference between conditions (59% vs. 33%).

3.3. Discussion

Comparing cases at recall time facilitated relational retrieval. Consultants who compared the two cases were more likely to retrieve an analogous example from their long-term memory than those who read the same two cases but did not compare them. Moreover, consistent with prior research, participants who compared the two cases were more likely to transfer the common relational structure to a later (analogous) negotiation situation. Thus, analogical abstraction produced two kinds of effects, one radiating backward and one radiating forward.

The fact that comparison led to an advantage in relational retrieval for items stored *prior to* the session and across years of life experience is consistent with the relational schema abstraction account. It is also evidence for the efficacy of late analogical abstraction. The results suggest that analogical comparison of two exemplars led participants to form a more abstract relational schema of the contingent contract principle. This schema was a better match for previously learned analogous exemplars (when serving as a probe) and was more reliably retrieved (when serving as a memory item) when participants received an analogous test case as probe. It appears that a relational schema makes a better relational match for relationally similar examples (relative to other analogous examples) whether it occurs before or after them.

Our results for the separate-case group are consistent with prior findings, suggesting that analogical retrieval is a relatively rare event (Catrambone & Holyoak, 1989; Gentner et al., 1993; Gick & Holyoak, 1980, 1983; Ross, 1987). Despite the consultants' considerable experience in the business world, and despite the fact that they had knowledge of value billing, almost half of them did not recall any examples of contingent contracts. Indeed, as noted above, 11% failed to write down any case at all. Against this background, the large gain in relational retrieval that resulted from simply comparing two probe analogs is all the more striking.

These results bring home both the power and fragility of comparison. As discussed earlier, comparison is a powerful learning mechanism *when it occurs*; however, it often fails to occur in situations where it could be effective—even when two analogous cases are analyzed in immediate succession, as here. Although our participants were consultants whose jobs depend on their ability to apply their knowledge in new situations, and who spent considerable time with the training materials, we saw little spontaneous comparison across examples in the group not explicitly told to compare.

In sum, these results show that analogical abstraction aids in relational retrieval from memory as well as in relational transfer. This is consistent with the schema abstraction view that analogical comparison promotes more abstract and relationally focused encodings. Further, the effects of alignment and abstraction in memory access are bidirectional. The representations that resulted from comparison were both more readily retrieved by future analogs and more effective as probes for prior analogs stored in memory than were the separate encodings.

4. Experiment 2

In Experiment 2 we sought to generalize the retrieval findings with a less expert population. The participants in this experiment were students in a Masters of Accounting program. Most of these students were recent college graduates (about 2 years out of college); they have considerably less work experience than the consultants in Experiment 1. Thus, this study tests whether the comparison advantage on retrieval will hold for a group that lacks extensive prior business experience. As comparison effects on transfer in negotiation across different levels of expertise have been amply demonstrated in our previous studies (e.g., Gentner et al., 2003; Loewenstein et al., 1999; Thompson, Gentner, & Loewenstein, 2000), in the remaining studies we focus on the effects of late analogical abstraction on retrieval from memory.

4.1. Method

4.1.1. Participants

Participants were 36 Masters of Accounting students participating through a course in negotiation. There were 19 in the comparison condition and 17 in the separate-case condition.

4.1.2. Procedure and scoring

The materials and procedure were as in Experiment 1. Participants were given the same training packets as in Experiment 1, except the retrieval instruction was as follows: ''Please think of an example, preferably from your own experience, that embodies the same principle as what you have just read.'' Participants' case descriptions (including whether they linked the two cases) and retrievals were evaluated as in the prior study. There was 86% agreement (min $\kappa = 0.70$) between two coders on these ratings and disagreements were resolved through discussion.

4.2. Results and discussion

As in the first study, the comparison group produced more complete formulations of the contingency contract principle in their case descriptions and were more likely to retrieve examples of contingent contracts from their own memory. The case descriptions of the comparison participants ($M = 1.42$) better articulated the contingent contract principle than did those of the separate-case participants ($M = 0.53$) ($t[34] = 3.41$, $p < .01$). Also as predicted, participants in the comparison group retrieved better examples of contingent contracts from their own life experience than did participants in the separate-case condition (mean ratings of 1.47 vs. 0.59) ($t[34] = 2.88$, $p < .01$). There was a clear association between extracting the principle and recalling a matching example, $r_s = .51$, $p < .005$. Finally, as in Experiment 1, nearly all (95%) the comparison participants linked the cases, whereas among separate-case participants only 18% did so (Fig. 2).

Most participants (86%) retrieved some kind of example. The retrieved examples were mostly from their own or a colleague's experience (74%), with a few from prior classes (3%), the popular press (3%), or a nonspecified source (19%). The majority of the retrievals were examples from daily life; only 35% were from business situations. (This contrasts with the consultants' pattern of 65% examples from business.)

We interpret these findings to mean (a) that comparing cases illustrating contingent contracts led participants to arrive at an abstract representation of the contingent contract principle and (b) that this abstract relational schema was a more effective probe for retrieving

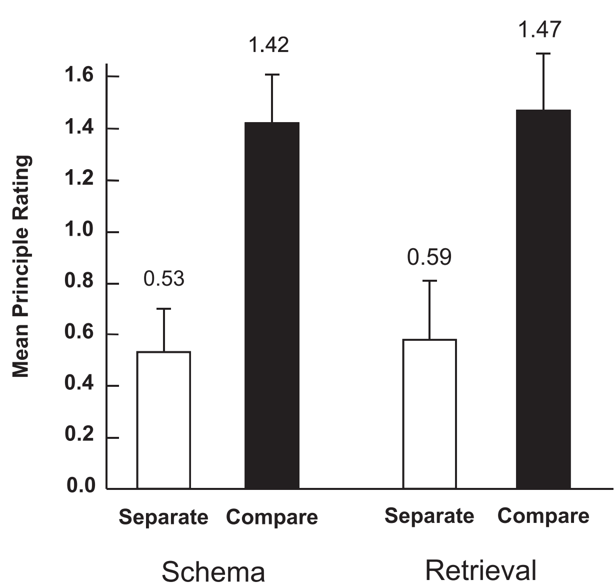

Fig. 2. Results of Experiment 2: Mean principle and retrieval ratings by condition.

further examples of the principle from long-term memory than were the separate-case representations. Thus, we have replicated the central finding from Experiment 1, that late analogical abstraction can facilitate relational retrieval, with a more novice population. Again, these findings suggest a symmetry: An abstract relational representation when used as a probe will lead to more accurate relational retrievals; and when stored in memory, they will be called forth more reliably by future relationally similar situations.

5. Experiment 3

Experiment 3 expanded the set of negotiation strategies considered. In addition to the contingent contract strategy studied thus far, we included examples illustrating a trade-off strategy. The trade-off strategy applies when the negotiating parties differ in the importance they place on the issues: for example, when one party's high-interest issue is the other party's low-interest issue. In this case, both parties can get what they most want by giving up something they consider less valuable (Froman & Cohen, 1970). A classic example (attributed to Follett, 1940) is the story of two sisters quarreling over an orange. They compromise and cut it in half, but they later realize that one wanted the juice and the other, the peel. Their compromise solution was suboptimal; both sisters could have had more of what they needed by adopting a trade-off in which one received all the juice and the other all the peel. Whenever two parties differ in their relative weightings of issues, a trade-off solution is likely to be superior to a compromise. Despite the apparent obviousness of this point, students taught this strategy in class often fail to apply it in test negotiations.

A further extension was that we added an additional control condition. It is conceivable that the comparison advantage stemmed from poor performance in the separate-case condition rather than from elevated performance in the comparison condition. Therefore, we included a condition in which people simply read one case and then were asked to retrieve an example.

5.1. Method

5.1.1. Participants

A total of 50 MBA students were each paid $10 for their participation. They ranged in age from about 25 to 35 years, and on average had about 5 years of work experience; hence, they are intermediate in expertise between the relatively novice Masters of Accounting students from Experiment 2 (with less than 2 years of work experience) and the more experienced consultants from Experiment 1 (who had about 15 years of work experience on average). There were 18 in the comparison condition, 18 in the separate-case condition, and 14 in the single-case condition.

5.1.2. Materials and procedure

Participants were given a packet of materials and allowed as much time as they needed to complete them. Most participants spent about 30 min on the task. The packets asked for case descriptions and retrievals first for contingent contract examples, and then for trade-off examples.

In the contingent contract section, participants first read the same brief cases exemplifying contingency contract agreements as in prior studies, with the same case description instructions for the comparison and separate-case groups. Participants in the single-case condition read one of the cases (half read one case, half read the other) and were given the same case description instructions as the separate-case group, but adapted to one case. Next, all participants were asked, ''Please think of an example, preferably from your own experience, that embodies the same principle as what you have just read.'' As before, there were further instructions as to the desired length and level of detail, as well as instructions to provide the source of the example. The trade-off section had the same format as the contingency contract section, but with new examples that exemplified trade-offs (see Appendix).

5.1.3. Scoring

Participants' case descriptions and retrievals were evaluated as to whether they adequately described contingency contract principles or trade-off principles (as appropriate), and they scored on the same three-point scale used in the prior studies. Also rated was whether participants explicitly linked the two cases in any way. There was 92% agreement (min $\kappa = 0.82$) between two coders on these ratings, and disagreements were resolved through discussion.

5.2. Results

Participants who compared two cases during learning were both better able to articulate the principles and better able to retrieve analogical matches from their own memories than those who did not. Overall, participants who analyzed two cases separately performed no better than those who studied just one case.

5.2.1. Case descriptions

Across both contingent contracts and trade-offs, comparison participants ($M = 1.33$) better captured the principles than did participants in the other two (control) conditions ($M = 0.81$; Fig. 3). A mixed measures ANOVA confirmed the comparison effect ($F[1,48] = 8.63$, MSE = 0.36, $p < .01$) and revealed no differences due to principle type ($F[1, 48] = 1.55$, MSE = 0.45, $p = .22$). Contrast tests showed a reliable advantage for the comparison group against both the separate-case group (Dunnet test, $p < .05$) and the single-case group (Dunnet test, $p < .05$). The comparison group relative to the combined control groups better articulated the contingent contract principle ($M_{\text{compare}} = 1.50$, $SD = 0.79$; $M_{\text{controls}} = 0.84$, $SD = 0.85$) ($t[46] = 2.70$, $p < .05$), but the difference for trade-off principle was not significant ($M_{\text{compare}} = 1.17$, $SD = 0.79$; $M_{\text{controls}} = 0.78$, $SD = 0.79$) ($t[46] = 1.66$, $p = .10$). Finally, we note that all (100%) of the comparison participants made a link between cases, whereas 36% of the separate-case participants did so.

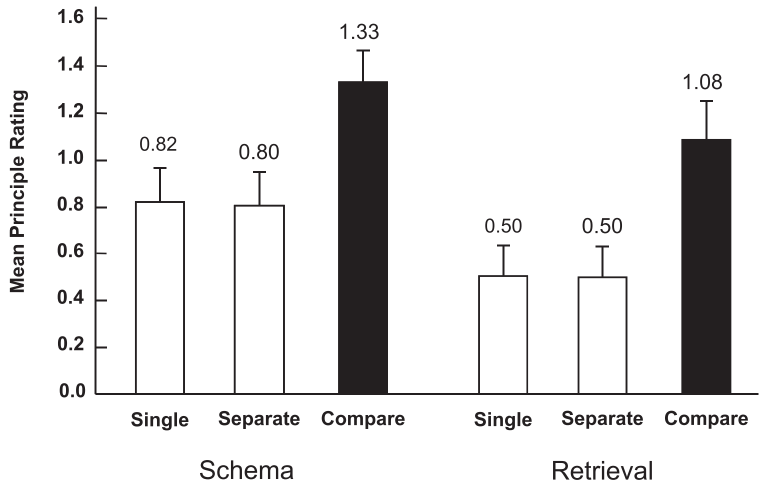

Fig. 3. Results of Experiment 3: Mean principle and retrieval ratings by condition.

5.2.2. Retrievals

Overall, comparison participants ($M = 1.08$) were more likely to retrieve a relationally matching example from memory than were participants in the two control conditions ($M = 0.50$; Table 4). A mixed-measures ANOVA confirmed the comparison effect ($F[1, 48] = 10.38$, MSE = 0.38, $p < .005$). There was no overall difference due to the type of principle ($F[1, 48] = 0.45$, MSE = 1.34, $p = .51$). Contrast tests showed that the comparison group performed reliably better than the separate-case group (Dunnet test, $p < .05$) and the single-case group (Dunnet test, $p < .05$). The comparison group showed more relational retrievals than the combined control groups for both the contingent contract examples ($M_{compare} = 1.00$, $SD = 0.97$; $M_{controls} = 0.47$, $SD = 0.72$) ($t[46] = 2.21$, $p < .05$), and the trade-off examples ($M_{compare} = 1.17$, $SD = 0.96$; $M_{controls} = 0.53$, $SD = 0.80$) ($t[46] = 2.47$, $p < .05$). We also found that the quality of participants' retrievals correlated with the quality of their case descriptions, both for contingent contracts ($r_s = .31$, $p < .05$) and for trade-offs ($r_s = .36$, $p < .05$). Overall, the data are consistent with the claim that drawing comparisons during learning facilitated participants' retrievals of relationally matching examples.

Table 4
Mean (*SD*) retrieval ratings (maximum 2) for the contingent contract and trade-off principles in Experiment 3

Condition	*n*	Contingent Contract	Trade-off	Total
Single-case	14	0.29* (0.61)	0.71 (0.91)	0.50* (0.52)
Separate-case	18	0.61 (0.78)	0.39* (0.70)	0.50* (0.73)
Comparison	18	1.00 (0.97)	1.17 (0.99)	1.08 (0.57)

Notes: Asterisks refer to paired contrasts with the comparison group in the same column.
*$p < .05$.

As before, the majority of the retrievals came from participants' own or a colleague's experience (73%), with others from the popular press (12%), or prior classes (4%). As in Experiment 1, only a small proportion of the retrievals (12%) were unspecified as to source, consistent with their being true retrievals. About half (54%) of the examples were from business situations, with the remainder being examples from daily life.

5.3. Discussion

As before, comparing analogous cases led to greater likelihood of analogical retrieval from participants' own long-term memories. The comparison advantage held for two different negotiation principles—contingent contracts and trade-offs—providing greater generalization of the effectiveness of late analogical abstraction. Strikingly, in the absence of comparison instructions, we found no significant advantage of studying two cases over studying one. Thus, it does not appear that the comparison advantage in the earlier studies stemmed from confusion in the separate-case condition. That reading two examples separately did not lead to better performance than reading one example again brings home the fact that opportunities for fruitful comparison are often missed—even between two analogous cases presented in succession.

6. Experiment 4

The studies so far have provided evidence that late analogical abstraction can facilitate relational retrieval from people's own long-term memory, as well as forward transfer to new situations. This is consistent with the possibility that relational schemas are operational in long-term domain learning because we have shown that they are effective at the scale of human domain experience. However, although the pattern of data suggests that participants were mostly reporting genuine retrievals (as discussed earlier), there still remains the possibility that some of these seeming retrievals were partly or wholly constructed at recall time to match the probes. To assess whether the retrieval benefits of comparison at recall time were real and that the retrieved items were encoded in an ordinary fashion, in Experiment 4 we manipulated participants' memory set.

As in prior studies of analogical retrieval, we provided participants with a set of examples and then cued them with analogous examples (e.g., Catrambone, 2002; Gentner et al., 1993; Wharton et al., 1994). Participants were given a set of negotiation examples, including one contingent contract example. After a filled delay, they were given two further contingent contract examples, with or without instructions to compare them, and asked to recall an earlier matching story. Thus Experiment 4 complements the prior three experiments by equating the initial memory set across groups and removing the possibility of confabulation. Our prediction is that participants who compare two probe examples will be more likely to recall the initial analogous story than will those who analyze the probe examples separately.

6.1. Method

6.1.1. Participants

Participants were 73 undergraduates at a large state university, participating for extra credit in an introductory management class. They were randomly assigned to the comparison ($n = 37$) or separate-case ($n = 36$) conditions.

6.1.2. Materials and procedure

Participants sat at individual computer terminals. They received a series of seven negotiation cases and were asked to indicate, for each case, whether the agreement was advantageous for one or both parties. Four of the seven examples were filler items, in the sense that they bore no clear resemblance (apart from being negotiations) to the later negotiation examples that served as memory cues. Of the remaining three, one described a contingent contract and was analogous to the later examples. The contingent contract was between a small software company and the manufacturer of a new video player; the software company would receive a share of profits only if their data compression algorithm was mentioned by trade reviews as key to the video player's success. The remaining two negotiation examples were not contingent contracts but had surface elements resembling the probe examples used as memory cues. One concerned two sisters running a gardening store (surface similar to the probe example of two brothers who inherited a farm), and the other concerned a Canadian buyer purchasing goods from a Southeast Asian merchant (surface similar to the probe example of an American buyer purchasing goods from an Asian merchant). After reading these seven negotiation examples, participants spent 30 minutes on various unrelated tasks.

Participants were then asked to read the Poor Brothers and Asian Merchant contingent contract negotiation examples used in prior experiments. The instructions were very similar to those in the prior experiments. Specifically, in the comparison condition, participants were told: ''What is going on in these negotiations? Think about their similarities. What are the key parallels in the two negotiations? Please describe the solution and say how successful you think it is.'' Then, on the next screen, participants were asked: ''Please recall an example presented earlier in the hour that best matches the cases you just read on the previous page. Please be sure to describe the key elements of the prior example in sufficient detail so that we can understand which one you mean.'' In the separate-case condition, participants were given one probe and asked: ''What is going on in this negotiation? Please describe the solution and say how successful you think it is.'' Then, on the subsequent screen, participants were asked: ''Please recall an example presented earlier in the hour that best matches the case you just read on the previous page. Please be sure to describe the key elements of the prior example in sufficient detail so that we can understand which one you mean.'' Then the separate-case participants were given the second probe and asked the same two questions. In both conditions, the order of presentation was counterbalanced across participants.

6.1.3. Scoring

Two raters, blind to the purpose of the study, scored participants' responses as to whether they had recalled the prior contingent contract example and the two prior surface similar examples. This was straightforward (min $\kappa = 0.94$), and the few disagreements were resolved through discussion.

6.2. Results and discussion

As predicted, participants in the comparison condition (27%) were more likely to recall the prior analogous example than were participants in the separate-case condition (6%) ($\chi^2[1, N = 73] = 6.12, p < .05$). This is consistent with the relational schema abstraction account: drawing a comparison and using the resulting schema as a memory probe facilitated analogical retrieval.

We also examined the retrieval of surface similar examples. Fewer comparison participants (30%) retrieved a surface similar example than separate-case participants (78%) ($\chi^2[1, N = 73] = 6.12, p < .05$). The remainder of the participants' responses was retrievals of other, less related examples or to say they could not think of a similar example.

The comparison and separate-case participants stored the same initial stories under the same conditions and read the same two stories at probe time. Yet the comparison group was roughly four times more likely to recall the analogous example. These results buttress the previous results from our more naturalistic studies of long-term retrieval from people's own memories (Experiments 1–3). Both in naturalistic conditions and when given a controlled memory set, comparing relationally similar examples at recall time increases the likelihood of a relational retrieval.

7. Experiment 5

A frequent outcome of analogical retrieval is analogical transfer—the application of a solution strategy from a prior example to a current problem (Gick & Holyoak, 1980, 1983; Novick, 1988). Therefore, a natural extension of the present research is to ask whether comparing items at recall time would facilitate problem-solving transfer based on a single prior example.

To examine the effects of comparison during the problem-solving phase, we extended work by Kurtz and Loewenstein (2007). Using materials from Gick and Holyoak's (1980, 1983) problem-solving task, they showed that the classic pattern—that even after hearing an analogous convergence story, participants mostly fail to solve Duncker's (1945) tumor-radiation problem—can be altered by providing participants with a comparison at recall time. The contrast with Gick and Holyoak's studies was that at test time, participants received *two* problems for which convergence solutions applied, instead of one. Participants who compared the two problems had a solution rate of about 50% for the tumor problem, reliably higher than the 15% solution rate found for participants who were not asked to compare the problems (and also higher than the 15% solution rate for those who received only one problem). As the encoding conditions for the initial story were the same across conditions, this pattern is evidence that analogical comparison at recall time facilitates analogical transfer. However, as participants were given the problems immediately after reading the initial story, it is possible that the results stemmed from continued working-memory activation of the initial story. Day and Gentner (2007) have found that over short delays, implicit transfer from prior analogs can occur without explicit remindings to the prior analog; furthermore, some studies have shown relational priming over short time periods (Bassok, Pedigo, & Oskarsson, 2008; Spellman, Holyoak, & Morrison, 2001).

To test whether late analogical abstraction can facilitate relational reminding and transfer over longer periods, in Experiment 5 we extended the Kurtz and Loewenstein (2007) method by instituting a filled delay after the initial story. In so doing, we extend the present studies by examining an effect of comparing current instances on analogical transfer based on a single stored example. A further extension was that, after posing solutions, we asked participants whether they were aware of recalling the initial analog—that is, we assessed whether analogical transfer was predicated on explicit recall.

7.1. Method

7.1.1. Participants

Participants were 227 undergraduates at a large state university, participating for extra credit in an introductory management class. They were randomly assigned to the comparison (n = 116) or separate-case (n = 111) condition.

7.1.2. Materials and procedure

Participants sat at individual computers and were presented with a series of five cases and questions. They were told that they were to read a series of cases and decision scenarios, and answer questions about them. The first case concerned Disney's marketing of Winnie the Pooh merchandise, and we asked: ''What implications for strategy does this story offer?'' The second was the fortress convergence story, as used by Gick and Holyoak (1980). We asked participants to ''please read the following simplified story about tactical moves as part of a larger military strategy. Make sure you are familiar enough with the story that could retell it in your own words.'' Following the story, participants were asked: ''What implications for strategy does this story offer?'' Participants then read three further cases and were asked multiple questions about each of them. On average, participants required about 25 minutes for the five cases.

Participants were then asked to read two analogous problems, counterbalanced for order of presentation: the tumor problem and the Red Adair problem (Gick & Holyoak, 1983). In the comparison condition, participants were told: ''After reading the following two problems carefully, please consider the parallels between them. What approach would you take to solve both of the following problems? Here's an important hint: The same strategy can be used to solve both problems.'' In the separate-case condition, participants were asked, for each problem separately: ''Please read the following problem carefully. What approach would you take to solve the following problem?''

To examine whether analogical transfer was due to a conscious analogical retrieval of the initial story, after solving (or failing to solve) the problems, participants in both conditions were asked whether they had been reminded of any of the prior cases when solving these problems. If they had, they were asked to list which one(s) and whether they were helpful.

7.1.3. Scoring

Two raters, blinded to the purpose of the study, scored participants' responses as to whether they had generated convergence solutions to the tumor problem ($\kappa = 0.91$). To be scored as a convergence solution, participants had to mention using multiple lower powered rays at the same time from different directions. Disagreements were resolved through discussion.

7.2. Results and discussion

As predicted, participants in the comparison condition (31%) were more likely to solve the tumor problem than were participants in the separate-case condition (5%) ($\chi^2[1, N = 227] = 24.71$, $p < .001$). A similar but nonsignificant trend held for the Red Adair problem, 41% versus 30% ($\chi^2[1, N = 227] = 2.89$, $p = .08$). We are not sure why the separate-case group performed so well in the Red Adair problem; possibly participants were more familiar with the behavior of water, hoses, and fire fighting than with that of rays, the strength of rays, and tumors. Nonetheless, the overall pattern supported a comparison advantage, consistent with the prediction that analogical abstraction between two current problems facilitated analogical transfer from a relationally similar prior example.

A further question was whether people reported explicitly being reminded of the earlier analogous story when solving the problems, and whether this was related to analogical transfer. To be clear, these recalls are not the same as the recall questions in the earlier experiments. Whereas in the earlier experiments the question was whether, after reading an example, people could recall a matching one, in this experiment the question was whether, *during problem solving*, people were spontaneously and explicitly reminded of an earlier analog. The overall rate of being reminded of the earlier analog was low and hardly differed between the comparison (18%) and separate-case (14%) conditions. However, the remindings seemed to have somewhat different effects for those in the two conditions, as shown by a binary logistic regression predicting whether people generated a convergence solution to at least one of the two problems. This analysis showed that there was a reliable effect of comparison[3] (Wald $\chi^2 = 5.12$, $p < .05$), a reliable effect of reported explicit recall of the earlier analog (Wald $\chi^2 = 13.22$, $p < .001$), and a reliable interaction between comparison and explicit recall (Wald $\chi^2 = 3.88$, $p < .05$). Of the 21 participants in the comparison condition who reported explicitly recalling the analog, 81% generated a convergence solution, whereas among those comparison participants who did *not* report recalling the analog, just 33% ($n = 95$) generated a convergence solution ($\chi^2[1, N = 116] = 16.56$, $p < .001$). By contrast, there was little difference in the proportion generating convergence solutions within the separate-case condition based on whether people reported explicitly recalling the analog (44%, $n = 16$) or not (31%, $n = 95$) ($\chi^2[1, N = 111] = 1.09$, $p = .30$). Thus, being spontaneously reminded of the earlier analog coincided with analogical transfer for those in the comparison condition but did not clearly do so for those in the separate-case condition.

Experiment 5 showed that late analogical abstraction facilitates retrieval analogical transfer even when the initial memory set is controlled, and after a filled delay. The comparison and separate-case participants stored the same initial stories under the same conditions, and later read the same two problems at recall time. Yet the comparison group was more likely to solve the test problem.

8. Simulating backward retrieval and forward transfer

The five experiments support two important phenomena: (a) that analogical abstraction at recall time promotes relational *retrieval* of prior exemplars from long-term memory and (b) (in concert with prior findings) that analogical abstraction at learning time promotes relational *transfer* to future exemplars. We have hypothesized that analogical comparison and abstraction of relational schemas underlies both backward relational retrieval and forward relational transfer. In the following simulation studies, we provide a unified account showing that relational schemas produced by analogical abstraction can account for both the backward retrieval and forward transfer effects. As the memory set, we used a large, varied set of exemplars as a proxy for long-term memory, as further described below.

To account for the observed advantage of analogical abstraction on both backward retrieval and forward transfer, we draw on two properties of relational schemas: *diminished surface competition* and *stronger relational matches*. The first property, *diminished surface competition*, stems from the fact that a relational schema will have few surface features. If a relational schema is a probe, its few surface features should mean that few surface feature matches are retrieved. By contrast, if an ordinary example is a probe, surface feature matches typically out-compete and crowd out relational matches (Catrambone & Holyoak, 1989; Gentner et al., 1993; Novick, 1988). Thus, diminished surface competition should be capable of accounting for the effect of analogical abstraction on backward relational retrieval. However, diminished surface competition cannot explain the forward effect: that a relational schema stored in memory leads to greater relational transfer to a specific probe. In this situation, the probe is a specific example, and therefore liable to retrieve surface matches. A relational schema stored in memory will not change this; indeed, its lack of surface feature matches to the probe example only lessens its chances of being retrieved.

To account for the forward effects of analogical abstraction, a second property of relational schemas must be invoked: namely *stronger relational matches*. Some models (including MAC/FAC, Forbus et al., 1995; and LISA, Hummel & Holyoak, 1997) require that the total weight allotted to a memory representation be constant. If this is the case, the relations in a relational schema gain greater weight than those same relations would if embedded in an example because the example also must allot weight to its surface features. Thus, stronger relational matches, unlike diminished surface competition, can potentially account for the well-attested finding that relational schemas facilitate forward relational transfer.

We propose that these two properties of relational schemas—diminished surface competition and stronger relational matches—permit a unified model of analogical abstraction and retrieval that naturally explains the advantage of relational schemas in both backward retrieval and forward transfer found in the behavioral experiments. In the four simulation studies that follow, we use MAC/FAC (Forbus et al., 1995) to test this account. If these assumptions can account for both relational retrieval (Simulation studies 1 and 2) and relational transfer (Simulation studies 3 and 4), this will support the claim that the effects of analogical comparison arise through relational schema abstraction. These simulations are the first attempt at modeling the effects of relational schemas as probes to long-term memory. They are also the largest simulations to date of the effects of relational schemas on forward transfer.

8.1. MAC/FAC: Many are called but few are chosen

Forbus et al. (1995) designed MAC/FAC to account for six constraints on a computational model of similarity-based retrieval and processing. To account for findings on mapping and reasoning, MAC/FAC (a) stores s*tructured representations* and (b) (in FAC) incorporates *structure-mapping processes.* To account for findings on memory retrieval, the MAC part of the model captures (c) the *primacy of the mundane*: literal similarity matches (matches high in both structural and surface commonalities) are the most frequent kind of reminding; (d) the *surface superiority* criterion: retrievals based on surface similarity are more frequent than analogical matches; (e) the *rare insights* criterion: although infrequent, relational remindings occasionally occur; and, crucial to a model of human memory; and (f) the *scalability* criterion: the model must be capable of operating over a vast memory.

The challenge to memory models is to capture both the large-scale fast, parallel, and rather stupid process of memory retrieval (the last four criteria) and the discerning, structure-sensitive processes of mapping and inference that occur once the memory item has been retrieved. MAC/FAC captures both sides of this divide by having two stages: a computationally cheap initial filtering process (MAC) that operates in parallel to find a subset of likely candidates for more expensive processing (King & Bareiss, 1989; Waltz, 1989), and a structure-sensitive, computationally more expensive process (FAC) that operates over the candidates provided by the MAC stage.

The MAC stage relies on *content vectors*: flat feature vectors composed only of the content of the memory representation and not its structure. These content vectors are unit vectors; thus, each is given the same total amount of weight. The content vectors contain the constituent predicates of their representations, each weighted according to how many instances appear in the representation. This means that the vector representing an abstract schema gives more weight to its relations than does a specific example that has to distribute its weight among not only its relations but also its surface attributes.

The MAC stage of memory access consists of computing, in parallel, dot products between the content vector of the probe and the content vectors of items in memory. In this way MAC captures the standard assumption that memory access is largely determined by the match between the probe representation and the stored representation(s) (Gillund & Shiffrin, 1984; Hintzmann, 1986). The dot-product computation is simple enough to make it feasible to compute many matches in parallel, allowing MAC to satisfy the scalability criterion.

MAC passes its top three candidates to the FAC stage, which filters them further. The FAC stage carries out a structure-mapping between the probe representation and each candidate using the structure-mapping engine (SME).[4] The FAC stage then returns its top one or a few candidates[5] as the final output of the retrieval process. MAC/FAC is not guaranteed to retrieve the best possible match. The ensuing simulations ask whether its probability of doing so when given relational schemas matches the human pattern.

8.2. Overview of the simulation studies

These simulations test the role of relational schemas in both backward retrieval and forward transfer using a consistent modeling framework. Simulations 1 and 2 examine the backwards retrieval effect demonstrated in our behavioral experiments by using relational schemas as probes for retrieving examples from LTM. Simulations 3 and 4 examine the forward transfer effect we demonstrated in Experiment 1 (as well as by much prior research) by testing whether relational schemas stored in LTM are well retrieved by future probes.

To simulate retrieval and transfer from LTM requires a large, varied memory set containing all the distinct kinds of matches that we know to be important in analogical retrieval and mapping. For this purpose, we used the Karla the Hawk story set (Gentner et al., 1993), which serves as a standard set of materials in simulations of analogical processing and relational retrieval (e.g., Catrambone, 2002; Gentner et al., 1993; Ramscar & Pain, 1996; Ramscar & Yarlett, 2003; Salvucci & Anderson, 2001; Thagard, Holyoak, Nelson, & Gochfeld, 1990). In addition to providing comparability with other modeling work, the Karla the Hawk set is the largest set of controlled materials available for simulating similarity-based retrieval. It includes seven story sets, each containing literally similar matches, shallow matches, and analogous matches. For each story set, we created relational schemas (as described below); we then tested their behavior as probes and as items-to-be-retrieved. As the issue of scalability is crucial in modeling LTM, we also included a large number of filler items in memory.

8.3. Simulation 1

In Simulation 1, we simulate the kind of situations we examined in Experiments 1–4: the use of either examples or relational schemas derived from comparisons as probes to memory. Available in memory were analogous examples, shallow matches, and filler items. Experiments 1–4 suggest that a relational schema as a probe, relative to an example as a probe, retrieves more analogous cases and fewer shallow matches. The question for Simulation 1 is whether MAC/FAC, like our human participants, will do so, thereby establishing the plausibility of our processing account for the newly demonstrated backward retrieval phenomenon.

8.3.1. Materials

We used seven Karla the Hawk story sets,[6] each containing six stories: two base stories used as probes, and four further stories used as memory items. The two base stories were constructed to be analogous to each other: that is, they had the same relational structure (both first-order and higher order relations) but different entities and contextual features. The memory items included an *analogous* story (a relational match to both bases) that had the same relational structure (both first-order and higher order relations) as the bases, but differed from both bases in its entities and contextual features. There were also two *shallow matches*,[7] one for each of the bases, which each had the same surface features and first-order relations (specific actions and events) as one of the bases but different higher order relational structure from the bases. Lastly, there was a first-order relational match—a story that matched in events but not in higher order relational structure nor in object attributes. As these share common actions, they are plausible foils and may be retrieved at low levels. In addition to these memory items that were related to the base stories, we added 48 additional examples with no systematic similarities to the base stories. This served as a test of MAC/FAC's ability to retrieve appropriate examples from a large and varied memory set. Across all the stories, we used the same small set of higher order relational predicates (e.g., CAUSE, IMPLIES, and FOLLOWS), although, critically, the structures they were used to form were different. Thus, just as in the Gentner et al. (1993) experiments described earlier, preferentially retrieving an analogous story over a shallow match story hinges on distinguishing a story that most completely matches the higher order relational structure in the probe.

For each story set we generated the relational schema resulting from comparing the two base stories. To create the schemas, we used SME to carry out a structural alignment between a pair of examples and stored the resulting relational schema. The schema consists of the common structure, including the consistent predicates that form the structure, the generic entities that play roles in the structure, and any consistent attributes shared by entities playing the same role. Nonmatching attributes and predicates are stripped away. The resulting schema can then be treated as an item in memory.

8.3.2. Procedure

We gave MAC/FAC a memory set of 76 items: the analogical match, the two shallow matches, and the first-order relational match from each of the seven story sets, plus the unrelated examples. Then we assessed MAC/FAC's retrieval pattern when given as a probe either a base story (14 instances) or the SME relational schema (seven instances). We used the standard setting determined through prior sensitivity analyses (Forbus et al., 1995): MAC and FAC each return up to three matches.

8.3.3. Results

First, inspection of SME's output confirmed that its relational schemas contain a higher proportion of relational predicates than the individual base stories (roughly 80% vs. 40%, respectively). This of course should follow from SME's analogical abstraction process, but it is important for the logic of the overall results to establish that SME's relational schemas are stronger relational matches.

Tables 5a–d show the results for both stages of MAC/FAC. We begin with the FAC results, which represent the final retrievals produced by the simulation. When individual stories (the base stories) were used as probes, the dominant response was shallow matches

(Story 1: 86%; Story 2: 100%). Analogical matches were rarely retrieved (Story 1: 14%; Story 2: 0%). By contrast, when SME's relational schemas were used as probes, analogical retrievals dominated (86%); relational schemas also occasionally retrieved the shallow match (29%) and, less often, first-order relational matches (14%). Thus, relational schemas are more likely to produce relational retrievals than are individual probes, even when those probes are relationally similar to items in the database. This matches the pattern found for humans.

Table 5a
Simulation 1: Mean proportion across seven story sets of each item type retrieved from memory (columns) for each probe type (rows), for the MAC and FAC retrieval phases

	MAC Retrieval				FAC Retrieval			
Probe	Analogy	Shallow	FOR	Other	Analogy	Shallow	FOR	Other
Story 1	0.14	1.00	0	0	0.14	0.86	0	0
Story 2	0	1.00	0	0	0	1.00	0	0
Schema	1.00	0.86	0.29	0	0.86	0.29	0.14	0

Notes: Memory contained seven each of analogies, shallow matches, first-order relational (FOR) matches, and 48 other, unrelated examples. There were seven instances of each kind of probe (base story and schema).

Table 5b
Simulation 2: Mean proportion across seven story sets of each item type retrieved from memory (columns) for each probe type (rows), for the MAC and FAC retrieval phases

	MAC Retrieval					FAC Retrieval				
Probe	Lit Sim	Analogy	Shallow	FOR	Other	Lit Sim	Analogy	Shallow	FOR	Other
Story	1.00	0	1.00	0	0	1.00	0	0.29	0	0
Schema	1.00	0.86	0.71	0.14	0	0.71	0.86	0.14	0	0

Notes: Memory contained seven each of literal similarity matches, analogies, shallow matches, first-order relational (FOR) matches, and 48 other, unrelated examples. There were seven instances of each kind of probe (base story and schema).

Table 5c
Simulation 3 compared with Simulation 1: Mean proportion of each item type retrieved from memory (columns) given a specific story as probe, with a relational schema present (Simulation 3) or absent (Simulation 1) in long-term memory

	MAC Retrieval					FAC Retrieval				
Schema in LTM	Analogy	Shallow	FOR	Other	Schema	Analogy	Shallow	FOR	Other	Schema
Present: Sim 3	0.14	1.00	0	0	0.57	0.14	0.57	0	0	0.57
Absent: Sim 1	0.14	1.00	0	0	–	0.14	0.86	0	0	–

Notes: Memory contained seven each of shallow matches, analogies, schemas, literal similarity matches, first-order relational (FOR) matches, and 48 other, unrelated examples. There were seven instances of each kind of probe (base story and schema). Note the dominance of shallow matches in FAC when no schema is present in long-term memory (Simulation 1).

Table 5d
Simulation 4 compared with Simulation 2: Mean proportion of each item type retrieved from memory (columns) given a specific story as probe

	MAC Retrieval					FAC Retrieval				
Schema in LTM	Lit Sim	Analogy	Shallow	FOR	Schema	Lit Sim	Analogy	Shallow	FOR	Schema
Present: Sim 4	1.00	0	1.00	0	0.43	0.86	0	0.14	0	0.29
Absent: Sim 2	1.00	0	1.00	0	–	1.00	0	0.29	0	–

Notes: Memory contained seven each of shallow matches, analogies, schemas, first-order relational (FOR) matches, and 48 other, unrelated examples. There were seven instances of each kind of probe (base story and schema). When literal similarity matches are available, they dominate during long-term memory retrieval regardless of whether schemas are present.

Turning to MAC, we find that when an individual story was used as probe, the primary retrieval at the MAC stage was its shallow match; the analogous story was rarely included in the MAC retrieval set. Thus, FAC, having no better match to choose from, generally returned the shallow match. By contrast, when the schema was used as a probe, MAC typically retrieved both the analogous story and the shallow match for that story set—setting the stage for FAC to select the analogous match. Poor matches (first-order relational matches and unrelated items) were rarely retrieved.

MAC/FAC's pattern matches the psychological finding of an advantage for relational schemas over specific examples as probes for relational retrieval. Internally, MAC/FAC's operation is consistent with both of the potential explanations discussed earlier for this advantage: diminished surface competition and stronger relational matches. As relational schemas contain many fewer specific object attributes and other low-level features, they attract fewer surface competitors (shallow matches) at the MAC stage. Because of the unit-vector normalization, the relations in a schema have greater weight than the relations in a specific example (even one with precisely the same relational structure); thus, MAC's dot-product score will be higher for a match between a schema and an example than for one between two analogous examples. Other matches are sometimes retrieved, but once the candidate retrievals are passed to FAC, relational similarity dominates. The bottleneck for MAC/FAC, as for humans, is whether potential relational retrievals are ever brought into consideration.

8.4. Simulation 2

As Forbus et al. (1995) point out, literally similar exemplars are highly retrievable from memory; one can even speculate that memory retrieval is evolutionarily geared to literally similar examples: for example, a tiger should reliably remind you of a tiger (Gentner, 1989). Therefore, it is important to know whether the pattern of data changes if people have a literally similar example stored in memory. Consequently, Simulation 2 added literally similar items to the memory set. As there were no meaningful differences between the two base stories in Simulation 1, we used the original base story from Gentner et al. (1993) for each set and compared its results with those for the SME schema. The predictions, based on human results, were that: (a) base stories should primarily retrieve their literal similarity matches (which are the maximally close matches), with few if any analogies or other partial matches and (b) schemas should retrieve many more analogy matches than should specific base stories, but fewer literally similar matches.

8.4.1. Results

For the concrete base probe, MAC retrieved examples with matching object attributes—both literal similarity matches (100%) and shallow matches (100%). FAC then rejected most of the shallow matches, retaining the literal similarity matches (which match the probe in relational structure as well as in surface features). Given that literal matches were available, the final FAC retrievals for the specific base story contained only literal similarity matches; there were no first-order relational matches, spurious matches, nor analogy matches. When SME's relational schema was used as probe, MAC retrieved high levels of literally similar matches (100%), analogy matches (86%), and shallow matches (71%). The FAC stage then largely rejected the shallow matches (14%), returning the literal similarity matches (71%) and analogy matches (86%).

Simulations 1 and 2 provided evidence for the predicted pattern, over a large number of competing matches and a large memory set (Fig. 4). First, specific story probes are highly likely to successfully retrieve literally similar matches, when they are available (100% in Simulation 2). This accords with Gentner et al.'s (1993) data and with the earlier MAC/FAC simulations (Forbus et al., 1995). However, if memory does not contain literally similar examples, as in Simulation 1 (and as often occurs in real life), then concrete probes are likely to retrieve mainly shallow matches. Second, schemas as probes retrieve a wider range of relational matches than do concrete examples. This is because schemas retrieve not only items that are literally similar to one of the schema's inputs but also other items that are analogs. The detailed tracing of the operation of the simulation provides support for both reduced surface competition and stronger relational matches as two factors that operate to create the schema advantage. Taken together, the two simulation studies support our processing account for the backward retrieval phenomenon we documented in the behavioral experiments.

8.5. Simulations 3 and 4

One goal in generating the simulation studies was to test a unified processing account of both the newly demonstrated backward retrieval effect and the well-known forward transfer effect (replicated in Experiment 1). Simulations 3 and 4 test the forward transfer effect by examining whether relational schemas are more likely than base stories to be *retrieved from* memory when a specific analog is used as probe. Such an advantage would support the claim that relations are weighted more heavily in relational schemas than in specific analogs. This is because the other contributor to the superiority of schemas as probes—reduced surface competition—does not apply if the probe is a specific case.

The earlier simulations showed that an example retrieves from memory, in order of decreasing likelihood: (a) literally similar cases, (b) shallow matches, and (c) analogous cases. The key question is where relational schemas will fit in this ordering. As they lack the

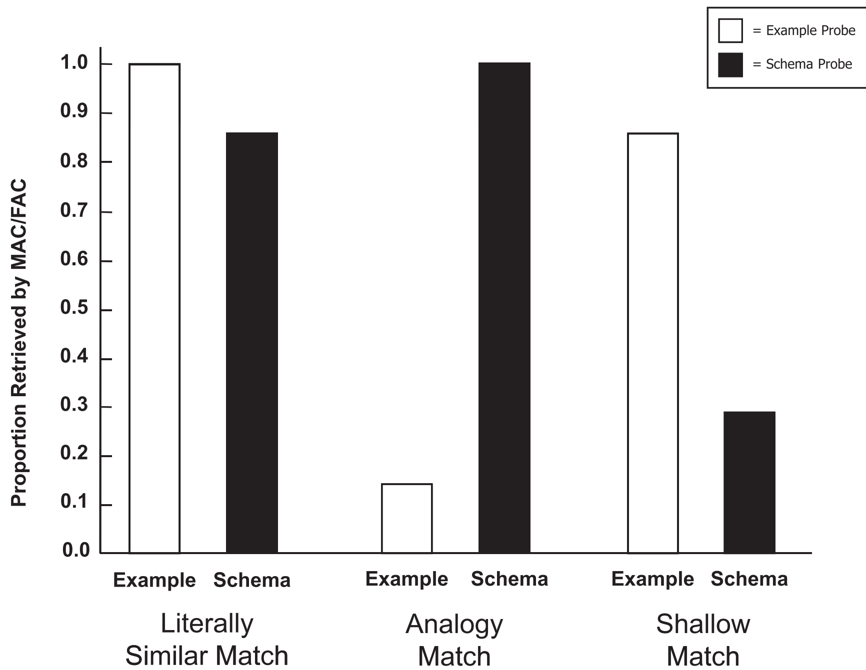

Fig. 4. Results of Simulations 1 and 2: Proportions of literally similar, analogical, and shallow matches retrieved by example and by schema probes.

added similarity of matching surface features, schemas should be retrieved less often than literally similar cases. However, as relations receive greater weight in schemas than in cases, schemas should be retrieved more often than analogous examples from memory. Thus, like shallow matches, relational schemas should rank somewhere between literal similarity matches and analogous matches in likelihood of retrieval given a probe example. We test this in two simulations that are parallel to Simulations 1 and 2, except with SME's relational schemas added to the memory set.

Some computational evidence that relational schemas are highly retrievable (given a relationally matching probe example) comes from a simulation by Hummel and Holyoak (1997). They ran the LISA model using a specific example (the Tumor problem) as a probe, varying which other item was stored in memory. The LISA model's retrieval index was the highest when the Surgeon problem (a highly similar example) was the memory item. The match was lower when the schema for the convergence situation was the memory item, and the lowest for the Fortress problem (an analogous example). However, LISA's match scores were generated one at a time, with only one item in long-term memory. As the essence of long-term memory retrieval is the challenge of finding one item among a large number of potential matches, in simulating memory retrieval it is desirable to have multiple items in memory. In the following simulations, we examine retrieval from a memory pool that includes the relational schema as well as over 75 items, so that the model must deal with multiple potential matches to the probe.

8.5.1. Materials and procedure

In Simulation 3, a base example was used as a probe; memory consisted of SME-generated relational schemas, analogies, shallow matches, FOR matches, and unrelated examples. In Simulation 4, the literally similar examples were added to memory. The base examples served as probes throughout. We then compared the output with that of Simulations 1 and 2, when we used the same base probes and the same memory sets except for the relational schemas.

8.5.2. Results

In Simulation 3, MAC frequently advanced the schemas (57% retrieved) along with the shallow matches (100%). FAC returned both relational schemas (57%) and shallow matches (57%). Analogy matches were rarely returned (14% through MAC and FAC). Thus, a base probe retrieved 57% of the relational schemas and only 14% of the analogical matches. We can compare these results with those of Simulation 1, which had no relational schemas in the memory set, and which also showed a low number of analogical matches (14%). This mirrors the pattern of ''rare relational retrieval'' of specific analogs found in human data, as well as the finding (extrapolating from transfer studies) that relational schemas are more readily retrieved.

In Simulation 4, literally similar examples were added to the memory set. As predicted (and as in human retrieval), they dominated memory retrieval (100% retrieved by MAC and by FAC). The shallow matches (100%) and some of the relational schemas (43%) were selected by MAC, but not by FAC (relational schemas: 29%; shallow matches: 14%). This is because both the SME schemas and the shallow matches were outscored by the literally similar examples, which match both in relational structure and in surface attributes. Analogical matches were not retrieved at all, being outscored by these other kinds of matches.

To summarize, when a specific example is used as probe, literally similar examples are by far the most likely to be retrieved, in accord with Forbus et al.'s (1995) ''primacy of the mundane'' principle. If literally similar matches are not present in the memory set, relational schemas and shallow matches are both well retrieved. Of course, they are retrieved for different reasons: Relational schemas are retrieved because their relations are weighted heavily and shallow matches are retrieved because of their many surface feature matches. Analogical matches are dominated by all three of these match types: by literally similar examples because they do not have matching features, by relational schemas because their relations are weighted less strongly (they are diluted by weight given to their unrelated features), and by shallow matches because of surface feature overlap. Hence, analogical matches are unlikely to be retrieved by a base example probe. This fits with human patterns. As Gentner (1989) noted, a tiger reliably reminds us of a tiger, but only occasionally of a shark.

Taken together, the four simulation studies support a consistent processing account of both backward retrieval and forward transfer. The simulations capture the pattern found in the studies (e.g., Experiment 1)—that analogical abstraction leads both to elevated relational transfer and to elevated relational retrieval from long-term memory. We showed that a relational schema derived from comparing analogous examples was more likely than a single example to retrieve a relationally similar example from memory, thereby simulating the backward retrieval effect shown in the behavioral experiments. We also showed that a relational schema stored in memory was more likely to be retrieved by a relationally similar probe than was a single example stored in memory, thereby simulating the forward transfer effect using the same processing account. Analogical abstraction resulting in relational schemas thus provides a parsimonious account of two important phenomena in learning and memory.

9. General discussion

These studies help establish *late analogical abstraction* as a new psychological phenomenon: Relational retrieval from long-term memory can be facilitated by abstractions formed after initial learning. Across the studies, analogical comparison at retrieval time—well after the recalled instances were stored—led to increases in relational retrieval. These findings have important implications for learning and teaching, in that they provide a pathway whereby people can retrieve and use formerly inert knowledge. They also bear on basic theories of learning and memory. Analogical comparison was positively associated with abstracting a schema, and schema quality was associated with the likelihood of relational retrieval. Simulation studies support two specific explanations for the relational retrieval advantage of schemas relative to examples as probes to memory: diminished surface competition and stronger relational matches. The simulations further support the idea that schemas formed via comparison during learning will be highly retrievable by relationally similar future probes, by virtue of having stronger relational matches than other analogs residing in memory.

In Experiment 1, participants who compared two examples of a contingent contract were better able to retrieve relationally similar exemplars from long-term memory. Further, these participants were also better at transferring to a future relationally similar example. In Experiment 2 we found a retrieval advantage due to late analogical abstraction with relatively inexperienced participants. Experiment 3 generalized the finding to a second kind of negotiation strategy, the trade-off strategy, and added a further control group presented with one probe example instead of two. Strikingly, receiving two analogous cases produced no advantage over receiving just one, unless the examples were compared. Experiment 4 showed that the comparison advantage for relational retrieval held for a controlled memory set. Finally, Experiment 5 extended the effects of late analogical abstraction to problem-solving transfer: comparing two problems facilitated generating solutions similar to a single example presented at an earlier time.

Simulation studies, using Forbus et al.'s (1995) MAC/FAC model of similarity-based retrieval, bear out the plausibility of this account. We found that relational schemas generated by SME resulted in superior relational retrieval over individual cases. This was true regardless of whether the relational schemas served as probes to memory (consistent with the current findings) or as items in memory (consistent with much prior research, as well as to the transfer findings of Experiment 1). The simulation studies also lend specificity to the explanation: Comparing two analogous examples yields a representation that is more abstract than those of the initial cases, in that many of the specific features of the examples will not be included. This means that during the initial parallel retrieval stage (the MAC stage): (a) there will be fewer surface-similar candidate retrievals, and thus less chance that true relational retrievals are crowded out (relative to single examples as probes) and (b) the match between the abstract probe and the LTM examples will show a stronger weighting for the common relational structure.

Taken together, the studies provide robust evidence that the process of comparing two analogous cases leads to superior relational retrieval over analyzing one case or even two analogous cases separately. We conclude that analogical abstraction facilitates relational access to memory in both directions: in retrieval of past experience as well as in transfer to future experience.

9.1. Explaining the analogical abstraction effect

The starting point for this work was the empirical generalization that comparing two analogous instances leads to better relational transfer (relative to studying the cases separately) to future cases with different surface features (Catrambone & Holyoak, 1989; Gentner et al., 2003; Gick & Holyoak, 1983; Loewenstein et al., 1999). We considered two potential explanations as to why comparison leads to better relational transfer: *learning to encode* and storing *relational schemas*. On the *learning-to-encode* account, analogical comparison does not lead to a stored abstraction; rather, the insight that results from a comparison influences the way future exemplars are encoded (Medin & Ross, 1989). To the extent that subsequent exemplars are encoded using the same relational structure, this uniform relational encoding will result in a larger number of relational matches with memory (Forbus et al., 1995). This account has the advantage of parsimony, in that schemas are not required. However, relational retrieval will only occur for exemplars learned after the insight. Inert knowledge is truly inert.

Although learning to encode is clearly part of the phenomenon of learning and conceptual change, the current results suggest that it cannot be the whole story. Relational schema formation is an important part of how comparison influences learning. Our results suggest that analogical abstraction favors the formation of a relational schema, which then acts as an effective retrieval probe for prior relationally matching examples. Specifically, we found that people who compared analogous exemplars (a) showed better transfer to future relationally similar exemplars; (b) showed better retrieval of past relationally similar exemplars; and (c) were more likely to map an inference from a prior relationally similar exemplar. Two further pieces of evidence that relational schemas are directly involved in these phenomena are that (d) people who compared analogous exemplars generated better relational descriptions of the principle than did those who received separate exemplars; and (e) people who wrote relationally accurate schemas were more likely to retrieve a relationally similar prior exemplar.

These phenomena argue strongly for the psychological reality of abstract schemas and suggest a plausible process—analogical abstraction—by which they can arise. It is not clear how a pure exemplar account would deal with these findings. We conclude that relational schemas formed via analogical comparison lead to elevated relational retrieval from long-term memory and to elevated relational transfer to future cases.

9.2. Open questions

In considering why relational schemas are more effective probes than single examples, we have focused on two specific factors: diminished competition from surface feature matches, and greater weight on relational matches. It is also possible that analogical abstraction invokes two further factors that increase relational retrieval: candidate inferences and re-representation. *Candidate inferences* are often produced when comparing examples. If these inferences seem correct, they may be incorporated into the common schema. In this case the resulting relational schema will be a more complete representation of the relational structure than (at least one of) the contributing examples. The second potential contributor, *re-representation*, refers to the idea that during comparison, people may generalize the relations in the examples to increase the match—*divorce* and *divest* may become *get rid of*, and *confiscate* and *buy* may become *obtain* (see Gentner & Kurtz, 2006). These more abstract relations could match a broader array of examples (cf. Clement, Mawby, & Giles, 1994).

A natural question at this point is whether the effect of comparison could be achieved more simply and directly through other means. For example, if abstracting relational schemas is important, perhaps people can simply be given abstract statements rather than comparing examples. However, we have found that participants who are provided directly with the relevant negotiation principle plus an accompanying example fail to show the forward transfer effect unless they are told to compare the two (Colhoun, Gentner & Loewenstein, 2008; Gentner, Loewenstein & Thompson, 2004). This is consistent with other studies showing that a concrete example and a statement of an abstract principle can be effective for learning, *if* the learner aligns and integrates them (Goldstone & Wilensky, 2008; Ross & Kilbane, 1997). However, although the process of comparison is important, it is not all that matters; the examples compared should align in such a way as to yield a relevant schema. As Gick and Holyoak (1983) demonstrated, comparing nonanalogous examples does not facilitate transfer.

Finally, we might ask whether comparison is simply an instance of elaboration or depth of processing, in which case, other ways of increasing the depth of processing of a probe could also facilitate relational retrieval. There are some grounds for pessimism here. Instructions that should have increased depth of processing—including describing the solution in an example and evaluating its effectiveness, as in the present studies, and also asking participants to provide advice to the protagonists (Thompson et al., 2000)—have not led to successful relational transfer. However, it remains an open question whether other methods of encouraging deep processing, such as self-explanation (Ahn, Brewer, & Mooney, 1992; Chi & VanLehn, 1991), focusing on predictive features (Johnson & Seifert, 1992), attempting to derive the principle, or providing perceptual simulation experience (Goldstone & Wilensky, 2008) could be used at probe time to improve relational retrieval.

Simulation work as well as experimental research will be important in further advancing these issues. We found that MAC/FAC and SME, without modification, were capable of successfully modeling the empirical results. It would be interesting to see whether and how other analogy systems, such as AMBR (Kokinov & Petrov, 2001), CAB (Larkey & Love, 2003), DRAMA (Eliasmith & Thagard, 2001), EMMA (Ramscar & Pain, 1996), or LISA (Hummel & Holyoak, 1997) would capture these findings.

9.3. Analogical retrieval: A good match is still hard to find

One striking aspect of the current results is how poorly participants perform at analogical retrieval and transfer in the absence of analogical comparison. Although this finding of poor relational retrieval has been found repeatedly in the laboratory, the current research is unusual in using people's real-life memories as the memory set (see also Blanchette & Dunbar, 2001). Thus, the pattern cannot be dismissed as an artifact of laboratory studies. In Experiment 1, the participants were management consultants paying thousands of dollars to take part in multiday negotiation training seminars. They were extrinsically and intrinsically motivated to learn to negotiate better. Further, all of them knew multiple potentially good analogs—most strikingly, one from their own profession, namely, value billing. Yet, in the separate-case condition, the majority of participants failed to write a good structural match, and 13% failed to write any example at all (Table 2). Separate-case participants very often retrieved a partial structural match rather than a full structural match, even though (as the debriefing after their face-to-face negotiation made clear) these participants knew better matches and failed to retrieve them. Thus, comparison at retrieval had a strong impact on the likelihood of retrieving prior relevant knowledge from the full spectrum of people's own life memories.

9.4. From laboratory to life

Although the failure of relational retrieval just discussed may be disheartening in one sense, from a scientific standpoint it is useful to know that the basic phenomenon discovered in the laboratory (e.g., Gentner et al., 1993; Gick & Holyoak, 1983; Keane, 1988) holds up quite well in access to autobiographical memory. People were relatively unlikely to retrieve pure relationally similar items when given single cases. Further, as documented in both laboratory and in more natural contexts, people are better able to do so when they have derived a relational abstraction via comparison.

9.5. Broader implications

The current findings suggest an expansion of the role for analogical comparison processes in learning and conceptual change. The established finding that analogical comparison can facilitate transfer to future relationally matching examples suggests ways of learning and teaching that promote broad applicability to future situations. The current findings go beyond this in showing that analogical comparison can also have a role in reviving past experience.

From an instructional point of view, the key point is that experience stored prior to understanding the principles of a domain is not necessarily lost or wasted. Of course, there will be limitations on this process: If the initial encoding is completely wrong (rather than simply overly concrete), then the case is unlikely to be retrieved by relational schemas formed in later learning. Still, in many cases, instructors may be capable of capitalizing on learners' prior knowledge by encouraging analogical comparisons that will lead students to form appropriate relational schemas. Students may then be better able to retrieve their own examples of key phenomena, thus permitting the prior examples to be reinterpreted in light of current knowledge. This also has implications for adults in work contexts because it means that people need not have learned something perfectly when they first encountered it—they may be capable of retrieving useful prior knowledge by developing their current understanding.

More broadly, an important implication of this finding is that even though initial learning is often contextually bound, learners may still be capable of connecting their concrete experience with abstract domain principles at a later time. There is a way out of the inert knowledge dilemma—a process by which reflection can retrieve and reorganize knowledge. The current results suggest that gains in knowledge can propagate not only forward in transfer but backward to illuminate prior learning.

Notes

1. Throughout this paper, we use the mean ratings of the participants' schemas and retrievals on a scale from 0 to 2 as dependent measures. If instead we assess the comparison effect by using nonparametric tests of association on the number of schemas or retrievals at each level of rated quality, we find the same pattern of results. Specifically, the comparison group's responses had reliably higher schema ratings than the separate-case groups in the first two studies, and from both of the other groups in Experiment 3 (Experiment 1: $\chi^2[2, N = 124] = 10.51$, $p < .01$; Experiment 2: $\chi^2[1, N = 36] = 10.13$, $p < .01$; Experiment 3: $\chi^2[2, N = 100] = 9.06$, $p < .05$), and reliably higher retrieval ratings as well (Experiment 1: $\chi^2[2, N = 124] = 7.22$, $p < .05$; Experiment 2: $\chi^2[1, N = 36] = 7.06$, $p < .01$; $\chi^2[2, N = 100] = 13.71$, $p < .01$). Table 1 summarizes the p-values for both the F- and t-tests reported in the main results sections for Experiments 1–3, as well as these nonparametric tests on the raw counts.
2. Across Experiments 1–3, we have four estimates of the association between schema and retrieval quality. Listed in order of magnitude they are $r_s = .18, .31, .36$, and .51. We are uncertain as to why Experiment 1 yielded a lower correlation than the subsequent experiments. Cross-experiment analyses showed that the confidence intervals for the middle two include the first and last; even without correcting for post hoc exploration, the first and last correlations are only marginally different ($p = .054$).
3. The following tests use Wald chi-squares, which are Wald statistics tested on a chi-squared distribution. They are used to test the effect of individual variables in logistic regressions, the same way t-tests are used to test the effect of individual variables in ordinary multiple regressions (e.g., Agresti, 2007).
4. SME (Falkenhainer, Forbus, & Gentner, 1989) takes two structured representations as input and outputs a structural alignment, an evaluation of the match (its Structural Evaluation Score), candidate inferences arising from the match. (For a more complete description, see Forbus et al., 1995; Gentner & Markman, 1997.)

5. The FAC stage can return up to three candidates, if their structural evaluations are very close to that of the top candidate; however, it rarely does so.
6. We used Forbus et al.'s (1995) original base, analogy and shallow match stories in our simulation experiment, and constructed for each set an additional base and its shallow match.
7. All stories in a set shared first-order relations—that is, events and states applying to entities, such as SHOOT(hunter, hawk). As there is some degree of relational overlap in all the stories, we use the term ''shallow matches'' rather than ''surface matches''; the latter term is often used for pairs that match solely in object features. The shallow matches used here correspond better to the partial relational matches that our participants produced when they failed to retrieve strong relational matches.

Acknowledgments

The research was supported by the Office of Naval Research, award number N00014-02-10078, to the first author, a fellowship from the Dispute Resolution Research Center to the second author, and a grant from the National Science Foundation, SES-9870892, to the third author.

References

Agresti, A. (2007). *An introduction to categorical data analysis* (2nd ed.). Hoboken, NJ: Wiley.

Ahn, W. K., Brewer, W., & Mooney, R. (1992). Schema acquisition from a single example. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*, 391–412.

Barnett, S. M., & Ceci, S. J. (2002). When and where do we apply what we learn? A taxonomy for far transfer. *Psychological Bulletin*, *128*, 612–637.

Bassok, M., Pedigo, S. F., & Oskarsson, A. T. (2008). Priming addition facts with semantic relations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*, 343–352.

Bazerman, M. H., & Gillespie, J. J. (1999). Betting on the future: The virtues of contingent contracts. *Harvard Business Review*, *77*, 155–160.

Blanchette, I., & Dunbar, K. (2001). Analogy use in naturalistic settings: The influence of audience, emotion, and goals. *Memory & Cognition*, *29*, 730–735.

Bransford, J. D. (1979). The role of prior knowledge. In J. D. Bransford (Ed.), *Human cognition: Learning, understanding, and remembering* (pp. 129–165). Belmont, CA: Wadsworth.

Brooks, L. R. (1987). Decentralized control of categorization: The role of prior processing episodes. In U. Neisser (Ed.), *Concepts and conceptual development: The ecological and intellectual factors in categorization*. (pp. 141–174). Cambridge, England: Cambridge University Press.

Brooks, L. R., Norman, G. R., & Allen, S. W. (1991). Role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, *120*, 278–287.

Catrambone, R. (2002). The effects of surface and structural features matches on the access of story analogies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 318–334.

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

Chi, M. T. H., Feltovitch, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.

Chi, M. T. H., & VanLehn, K. A. (1991). The content of physics self-explanations. *The Journal of the Learning Sciences*, *1*, 69–105.

Chipman S. F., J. W. Segal, & R. Glaser (Eds.) (1985). *Thinking and learning skills: Current research and open questions*, Vol. 2. (pp. 485–517). Hillsdale, NJ: Erlbaum.

Clement, C. A., Mawby, R., & Giles, D. E. (1994). The effects of manifest relational similarity on analog retrieval. *Journal of Memory and Language*, *33*, 396–420.

Colhoun, J., Gentner, D., & Loewenstein, J. (2008). Leaving abstract principles through principal-case comparison. In B. C. Love, K. McRae & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the cognitive Science Society* (pp. 1659–1664). Austin, TX: Cognitive Science Society.

Day, S. B., & Gentner, D. (2007). Nonintentional analogical inference in text comprehension. *Memory & Cognition*, *35*, 39–49.

Duncker, K. (1945). On problem-solving (L. S. Lees, Trans.). *Psychological Monographs*, *58*, Whole No. 270. (Original work published 1935).

Eliasmith, C., & Thagard, P. (2001). Integrating structure and meaning: A distributed model of analogical mapping. *Cognitive Science*, *25*, 245–286.

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

Follett, M. P. (1940). Constructive conflict. In H. C. Metcalf & L. Urwick (Eds.), *Dynamic administration: The collected papers of Mary Parker Follett* (pp. 30–49). New York: Harper.

Forbus, K. D., & Gentner, D. (1986). Learning physical domains: Toward a theoretical framework. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*, Vol. 2 (pp. 311–348). Los Altos, CA: Kaufmann.

Forbus, K. D., Gentner, D., & Law, K. (1995). MAC/FAC: A model of similarity-based retrieval. *Cognitive Science*, *19*, 141–205.

Froman, L. A., & Cohen, M. D. (1970). Compromise and logroll: comparing the efficiency of two bargaining processes. *Behavioral Science*, *30*, 180–183.

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 and analogical reasoning* (pp. 199–241). London: Cambridge University Press.

Gentner, D., & Christie, S. (2007). Learning new relations via structural alignment. In S. Vosniadou & D. Kayser (Eds.). *Proceedings of the Second European Cognitive Science Conference*. New York: Taylor & Francis.

Gentner, D., & Clement, C. (1988). Evidence for relational selectivity in the interpretation of analogy and metaphor. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Vol. 22 (pp. 307–358). New York: Academic Press.

Gentner, D., & Kurtz, K. (2006). Relations, objects, and the composition of analogies. *Cognitive Science*, *30*, 609–642.

Gentner, D., Loewenstein, J., & Thompson, L. (2003). Learning and transfer: A general role for analogical encoding. *Journal of Educational Psychology*, *95*, 393–408.

Gentner, D., Loewenstein, J., & Thompson, L. (2004). Analogical encoding: Facilitating Knowledge transfer and integration. In K. D. Forbus, D. Gentner & T. Reiger (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 452–457. Austin, TX: Cognitive Science Society.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, *52*, 45–56.

Gentner, D., & Medina, J. (1998). Similarity and the development of rules. *Cognition*, *65*, 263–297.

Gentner, D., & Namy, L. (1999). Comparison in the development of categories. *Cognitive Development*, *14*, 487–513.

Gentner, D., Rattermann, M. J., & Forbus, K. D. (1993). The roles of similarity in transfer: Separating retrievability and inferential soundness. *Cognitive Psychology*, *25*, 524–575.

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.

Gillund, G., & Shiffrin, R. M. (1984). A retrieval model for both recognition and recall. *Psychological Review*, *91*, 1–67.

Goldstone, R. L., & Wilensky, U. (2008). Promoting transfer through complex systems principles. *Journal of the Learning Sciences*, *26*(1), 465–516.

Hintzmann, D. L. (1986). Schema abstraction in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.

Holyoak, K. J., Gentner, D., & Kokinov, B. N. (2001). Introduction: The place of analogy in cognition. In D. Gentner, K. J. Holyoak, & B. N. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science* (pp. 1–19). Cambridge, MA: MIT Press.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, *15*, 332–340.

Hummel, J. E., & Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, *104*(3), 427–466.

Johnson, H. M., & Seifert, C. M. (1992). The role of predictive features in retrieving analogical cases. *Journal of Memory and Language*, *31*, 648–667.

Keane, M. T. (1988). *Analogical Problem solving*. Chichester, England: E. Horword.

King, J., & Bareiss, R. (1989). Similarity assessment and case-based reasoning. In *Proceedings: Case-Based Reasoning Workshop* (pp. 67–71). San Mateo, CA: Kautmann.

Kokinov, B., & Petrov, A. (2001). Integrating memory and reasoning in analogy-making: The AMBR model. In D. Gentner, K. Holyoak, & B. Kokinov (Eds.), *The analogical mind: Perspectives from cognitive science*. (pp. 59–124). Cambridge, MA: MIT Press.

Kurtz, K. J., & Loewenstein, J. (2007). Converging on a new role for analogy in problem solving and retrieval: When two problems are better than one. *Memory & Cognition*, *35*, 334–341.

Larkey, L. B., & Love, B. C. (2003). CAB: Connectionist analogy builder. *Cognitive Science*, *27*, 781–794.

Larkin, J., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, *208*, 1335–1342.

Loewenstein, J., Thompson, L., & Gentner, D. (1999). Analogical encoding facilitates knowledge transfer in negotiation. *Psychonomic Bulletin & Review*, *6*, 586–597.

Markman, A. B. (1999). *Knowledge representation*. Mahwah, NJ: Erlbaum.

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

Medin, D. L., & Ross, B. H. (1989). The specific character of abstract thought: Categorization, problem-solving, and induction. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence*, Vol. 5 (pp. 189–223). Hillsdale, NJ: Erlbaum.

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

Pfeffer, J., & Sutton, R. I. (2000). *The knowing–doing gap*. Boston: Harvard Business School Press.

Pirolli, P. L., & Anderson, J. R. (1985). The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology*, *39*, 240–272.

Ramscar, M., & Pain, H. (1996). Can a real distinction be made between cognitive theories of analogy and categorisation? In G. W. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 346–351). Hillsdale, NJ: Erlbaum.

Ramscar, M., & Yarlett, D. (2003). Semantic grounding in models of analogy: an environmental approach. *Cognitive Science*, *27*, 41–71.

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

Reeves, L. M., & Weisberg, R. W. (1994). The role of content and abstract information in analogical transfer. *Psychological Bulletin*, *115*, 381–400.

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, and 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, and Cognition*, *15*, 456–468.

Ross, B. H., & Kennedy, P. T. (1990). Generalizing from the use of earlier examples in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *16*, 42–55.

Ross, B. H., & Kilbane, M. C. (1997). Analogical mapping views and the effects of principle explanations. *Journal of Experimental Psychology Learning, Memory and Cognition*, *23* (2), 427–440.

Salvucci, D. D., & Anderson, J. R. (2001). Integrating analogical mapping and general problem solving: The path-mapping theory. *Cognitive Science*, *25*, 67–110.

Seifert, C. M., McKoon, G., Abelson, R. P., & Ratcliff, R. (1986). Memory connections between thematically similar episodes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 220–231.

Spellman, B. A., Holyoak, K. J., & Morrison, R. G. (2001). Analogical priming via semantic relations. *Memory & Cognition*, *29*, 383–393.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, *46*, 259–310.

Thompson, L., Gentner, D., & Loewenstein, J. (2000). Avoiding missed opportunities in managerial life: Analogical learning improves case-based transfer. *Organizational Behavior and Human Decision Processes*, *82*, 60–75.

Waltz, D. (1989). Is indexing used for retrieval? In R. Bareiss (Ed.), *Proceedings: Case-Based Reasoning Workshop* (pp. 41–45). San Mateo, CA: Morgan Kaufmann.

Weisberg, R. W. (1995). Prolegomena to theories of insight in problem solving: A taxonomy of problems. In R. J. Sternberg & J. E. Davidson (Eds.), *The nature of insight* (pp. 157–196). Cambridge, MA: MIT Press.

Wharton, C. M., Holyoak, K. J., Downing, P. E., Lange, T. E., Wickens, T. D., & Melz, E. R. (1994). Below the surface: Analogical similarity and retrieval competition in reminding. *Cognitive Psychology*, *26*, 64–101.

Appendix: Cases

Contingent contract cases

The Asian Merchant Case

Syd, a recently promoted head buyer of a small retail clothing chain in the United States, had bought some wholesale leather boots from a large Asian manufacturer. All aspects of the deal had been successfully negotiated except the shipment of the boots. Although Syd thought there might be potential difficulties in dealing with a large multinational conglomerate, she had decided to purchase boots from this company—rather than from several smaller available manufacturers—because of their aggressive pricing. The sales representative from the Asian company told Syd that they would pay to ship the boots by boat. Syd was concerned because the United States has announced that a trade embargo would likely be placed on all goods from that country in the near future. The Asian sales representative told Syd not to worry because the boat would arrive at the U.S.

dock before the embargo occurred. Syd, however, thought that the boat would be late and that the multinational conglomerate's policy to ship goods was not friendly to the needs of small businesses. Syd wanted the merchant to pay to ship the boots by air freight (which was substantially more expensive). The Asian sales representative refused because of the higher cost. They argued about when the boat would arrive. Syd considered breaking off the deal and buying the boots from her next best option, a small manufacturer. However, she did not like their product line as well and also wanted to try to establish a relationship with the large multinational conglomerate.

Finally, after a lengthy discussion in which it seemed that the negotiations might break off unsuccessfully, Syd suggested a new proposal to the sales representative from the Asian manufacturer. They would send the boots by air freight but both sides would watch when the boat actually docks in the United States. If the boat happened to arrive on time (as the sales representative believed it would), Syd would pay for the added cost of air freight. However, if the boat were to arrive late (as Syd believed it would), the Asian manufacturer would pay the air freight bill. Syd and the Asian manufacturer were pleased with this proposal, and each party got the arrangement they wanted.

The Two Poor Brothers Case

Two fairly poor brothers, Ben and Jerry, had just inherited a working farm whose main crop has a volatile price. Ben wanted to sell rights to the farm's output under a long-term contract for a fixed amount rather than depend upon shares of an uncertain revenue stream. In short, Ben was risk averse. Jerry, on the other hand, was confident that the next season would be spectacular and revenues would be high. In short, Jerry was risk seeking. The two argued for days and nights. Ben wanted to sell immediately because he believed the price of the crop would fall; Jerry wanted to keep the farm because he believed the price of the crop would increase. Jerry could not afford to buy Ben out at the time, but the strain on their family relationship over their disagreements was becoming too large. Ben had always trusted Jerry's instincts in the past, but this time felt Jerry was being overly optimistic.

In an effort to settle the matter and close up the growing rift between them, the two brothers agreed to meet with a family business advisor. Following the consultation, Jerry proposed a possible agreement to his brother: They would keep the farm for another season. If the price of the crop fell below a certain price (as Ben thought it would), then they would sell the farm and Ben would get 50% of the farm's current value, adjusted for inflation; Jerry would get the rest. However, if the price of the crop were to rise (as Jerry thought it would), Jerry would buy Ben out for 50% of the farm's current value, adjusted for inflation, and would get to keep all of the additional profits for himself. Jerry was delighted when his brother told him he could agree to this arrangement, thereby avoiding further conflict.

Trade-off cases

The Annual Meeting Case

The Sales and Marketing divisions of a large health maintenance corporation were trying to decide where to hold their annual meeting. There were substantial disagreements between the two divisions that were beginning to create conflict between them. The Sales division wanted to go to a lodge in the mountains. They had researched this possibility already and were anxious to reserve a location as soon as possible. The Marketing division, on the other hand, wanted to go a major city. They had already generated materials on the potential exposure of their company in several urban markets in preparation. The two divisions considered the compromise of holding two annual meetings, but both the added cost and the hectic travel schedules of the executives involved made this option unfeasible. There was added pressure from company leaders to hold a single annual meeting given the potential benefits of building better working relationships across divisions which had a history of tense and competitive relations. Indeed, one objective of the annual meeting was to foster better relations between the two divisions.

One member of the Sales division was appointed to negotiate with one member of the Marketing division concerning the location of the annual meeting. As the division planners gained trust in each other, important understandings emerged. The Sales division desired to run the annual meeting as a retreat to talk about important company issues, which required having a location suitable to focusing on the work at hand. The Marketing division saw the annual meeting as an opportunity to promote the company image. The two divisions resolved to create a well-publicized annual meeting located in the mountains, and they agreed to maintain open lines of communication between divisions.

The Videogame Sales Case

Vortex, Inc., a small video-arcade software firm, had a promising new line of special forces videogames. Keppel and Co., a major manufacturer of video-arcade equipment in Europe, was working with Vortex to produce the hardware needed for the special forces games. They were negotiating over how to share revenues from their joint product. The deal was mostly going smoothly—Vortex wanted to broaden the market for its products and Keppel needed a boost in sales to meet their shareholders expectations for the year. However, the two companies were struggling with how to split sales revenues. Keppel was demanding a high percentage from sales to finance the added expense of a custom-made action control for Vortex's games. Further, Keppel knew that it had the greatest resources to get Vortex's special forces games on the market. On the other hand, Vortex was also demanding a high percentage from sales on the grounds that what was being sold was their games, they had the patent on the new action control, and Keppel was simply one of several available manufacturers. Having negotiations at a standstill was bad for both companies because Keppel needed to increase their sales by the end of the year and Vortex needed to get their products out while they were still state of the art.

The breakthrough came when negotiators from Keppel and Vortex began discussing the differing needs of their companies. The negotiation teams reached the following agreement: Vortex would give up some of its share of revenue for the remainder of the year to cover Keppel's production costs and to aid their current financial situation. In return, Keppel would give up a comparable share of revenue in future fiscal years for these products, and Vortex still maintained their patent on the new control device.