

Class Inclusion, the Conjunction Fallacy, and Other Cognitive Illusions

VALERIE F. REYNA

University of Arizona

Class inclusion, the conjunction fallacy, and other phenomena are presented to illustrate cognitive illusions, that reasoning can be misdirected so that errors of judgment violate known principles. Errors in class inclusion are traced to accessing and implementing logical knowledge, rather than logical deficits, memory overload, or misleading questions. It is argued that superordinate-set cuing does not afford subset comparisons, but, rather, facilitates implementation of logical knowledge. According to fuzzy-trace theory, successful performance involves applying the cardinality principle, but also suppressing solutions based on judgments of relative numerosity, in which horizontal relationships between classes are apt to usurp the role of vertical ones. © 1991 Academic Press, Inc.

In the class-inclusion paradigm, children judge the relative numerosity of a superordinate set compared to its proper subset. Until the surprisingly advanced age of 9 or 10, most children make the transparently illogical inference that subsets are more numerous than superordinate sets. Brainerd and Reyna (1990a) proposed a new theory of this phenomenon whose central premise was that numerical information, which children readily use to make most types of quantitative comparisons, creates a powerful illusion of judgment when such comparisons involve inclusion hierarchies. Consequently, they argued, although numerical comparisons are requested on class-inclusion problems, mature reasoning suppresses such comparisons. Thus, the illusion is avoided by engaging in qualitative, rather than quantitative, reasoning. This explanation was shown to also imply a new account of perceptual salience effects involving the configuration of information in memory (rather than attention or encoding), both in the class-inclusion paradigm and in cognitive development generally.

Howe and Rabinowitz's (1991) commentary on Brainerd and Reyna's theory is not so much a point-for-point critique of that theory as it is a presentation of an alternative view. In this rejoinder, clarifications, criticisms, and their alternative view are discussed under the following headings: varieties of inclusion illusions; a fuzzy-trace theory of inclusion illusions; the cardinality principle: logical competence; retrieval necessity

Please send reprint requests to: Valerie Reyna, Department of Educational Psychology, University of Arizona, College of Education, Tucson, AZ 85721.

and processing interference; explanations of superordinate-set cuing as subclass comparison; working-memory explanations of class-inclusion performance; question wording and linguistic explanations of class inclusion; and conclusions.

VARIETIES OF INCLUSION ILLUSIONS

Class-inclusion errors should not occur. The questions are simple, in principle; children should be competent to answer them. Reviewing the literature, Winer (1980) concluded that this mystery had not been explained: "Moreover, the fact remains that children still have great difficulty with items that they certainly have knowledge of and should . . . be able to classify" (p. 324). Piaget made a special point of the fact that children who failed class-inclusion tasks, nevertheless, understood the inclusion relation, but that this competence was not sufficient to solve such problems. It is perhaps unfortunate that researchers continue to use the phrase "class-inclusion reasoning" to describe this paradigm when the understanding of class inclusion is not at issue.

At the level of data, it is apparent that some ways of asking about inclusion relations are transparent, and easy to solve, whereas others produce difficulties even for adults. To illustrate, consider the following problem that Tversky and Kahneman (1983, p. 297) posed to adults. Subjects read the description, and then ranked the alternatives in terms of probability.

Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Linda is a teacher in an elementary school.

Linda works in a bookstore and takes Yoga classes.

Linda is active in the feminist movement. [*F*]

Linda is a psychiatric social worker.

Linda is a member of the League of Women Voters.

Linda is a bank teller. [*B*]

Linda is an insurance salesperson.

Linda is a bank teller and is active in the feminist movement. [*B and F*]

The critical alternatives are *F* ("feminist"), *B* ("bank teller"), and *B and F*; the other items are fillers. Of course, subjects ranked *F* as probable and *B* as improbable. The important finding, however, is that they ranked *B and F* as more probable than *B* alone. This is called the conjunction fallacy because the class of bank tellers includes the class of bank tellers who are feminists; therefore, being a bank teller who is a feminist cannot be more probable than being a bank teller.

Note that the Linda problem is analogous to the class-inclusion problem, "Are there more animals or more cows?" (where there are 10 ani-

mals of which 7 are cows). In both problems, an inclusion relation stipulates a quantitative relationship, numerosity in class inclusion and probability in conjunction problems. Thus, ignoring some philosophical niceties, if we know that class B properly includes class A, we can conclude that B is more numerous than A (what we have called the cardinality principle, not to be confused with the cardinality concept), and we can conclude that being a B is more probable than being a B which is also an A.

The analogy between the Linda problem and the class-inclusion problem is instructive because it shows that adults are also subject to inclusion illusions. It is uncontroversial that adults understand inclusiveness, and can apply this understanding in making quantitative judgments. Neither memory overload as proposed by Howe and Rabinowitz nor logical deficits as proposed by Piaget would apply here to adults. Indeed, the superiority of older subjects, especially adults, on these dimensions is presumed in these theories. Consonant with this idea, when the conjunction fallacy is pointed out to adults, they immediately appreciate that their judgments violate the implications of class inclusion (Tversky & Kahneman, 1983).

There are other examples of inclusion errors to which adults fall prey (e.g., Falmagne, 1975). In syllogistic reasoning, for instance, the phrasing *All A are included in B* leads to more errors than *B includes all A*, and *All As are Bs* is sometimes misinterpreted as *All Bs are As*. So, reasoners lose track of possibilities, inferring (wrongly) from *Some A are B* and *All C are B* that *Some A are C*. In the latter syllogism, elements of C can be mapped onto B, but there may be other Bs that are not Cs. Such "parentheticals," the Bs that may not be Cs, are often neglected in reasoning. Moreover, whenever a transitive chain links elements of A, B, and C, the chain is salient, overshadowing parentheticals outside of the chain. Concrete content can help or hinder processing by making neglected alternatives more or less salient, respectively. Thus, adults can lose track of Cs that are not Bs (above) in syllogistic reasoning just as children lose track of the Bs that are As in class inclusion.

Procedural errors in accessing and implementing competence should be distinguished from lack of logical competence (see Overton, 1990). Procedural errors can be diminished by, for example, providing a notational scheme such as Venn diagrams for syllogistic reasoning, or superordinate-set tags in class inclusion. This reasoning is no less logical than unaided reasoning, however, though the mechanics may be simpler. As Falmagne (1975) points out, "awareness of the isomorphism between . . . the logical properties of "some" or "all" and the corresponding properties of the diagram representation. . . is not a trivial matter. . . Venn diagrams are conventional representations of a formal theory. . . there

may be no essential difference in terms of 'logicality' between a process such as this one and the apparently more formal, linguistically based processes" (p. 256). Therefore, class inclusion, conjunction fallacy, and (at least some) syllogistic reasoning errors do not demonstrate that reasoners lack the ability to reason logically about classes, since the requisite operations can be performed in other tasks. Difficulties in all three paradigms boil down to a failure to invoke logical knowledge in context (retrieval), or to apply it coherently (processing).

Cognitive illusions exist when compelling interpretations of problems misdirect reasoning away from relevant knowledge. One clue that subjects are in the grip of cognitive illusions is that performance is not at chance, but systematically below chance, which is a standard finding with both the conjunction fallacy and class inclusion. The fact that inclusion illusions are not random or idiosyncratic allows short-change artists to dupe victims into thinking that they are exchanging equivalent parts for wholes. Although victims soon realize that they have less money than they started with, they are not sure how it happened. Interestingly, the illusion of equivalence is often achieved by having the victim make the calculations himself. As I shall discuss, the very automaticity of certain judgments of relative quantity is partly responsible for leading reasoners down the garden path.

Therefore, one could quibble about whether children possess the competence necessary for class-inclusion reasoning, but adults undeniably have such competence, and, yet, they too are subject to inclusion illusions. The datum to be explained, then, is why reasoners are led astray in class-inclusion problems when they have all of the ingredients for successful performance. To say, as Howe and Rabinowitz do, that subjects make subset comparisons because they are "simpler," or because it is "their favorite reasoning strategy," or because they "prefer" subset comparisons to superordinate-subset comparisons is to beg the question.

A FUZZY-TRACE THEORY OF INCLUSION ILLUSIONS

According to fuzzy-trace theory, inclusion illusions have to do with the way the information in standard class-inclusion problems is presented. Because human beings reason intuitively, changes in problems that seem superficial from, say, a Bayesian or computational view, can produce wide variations in reasoning (Brainerd & Reyna, 1990b; Tversky & Kahneman, 1983). Intuition in fuzzy-trace theory is not "equate[d] . . . with the use of the cardinal-ordering principle" (Howe & Rabinowitz), but with a type of thought that is based on gist, rather than literal details, that is parallel, rather than linear as in logic, and is fuzzy or qualitative, rather than precise as in computational models (Reyna & Brainerd, 1991). Neither is gist imprecisely defined. Table 2 of Reyna and Brainerd (1991), for

example, spells out gist representations for seven standard tasks, including class inclusion. Memory for gist has long been contrasted with verbatim memory in linguistics and psycholinguistics, and is discussed in introductory textbooks in those areas (e.g., see "memory for substance" in Clark & Clark, 1977).

In brief, fuzzy-trace theory explains class-inclusion reasoning as follows. At some point prior to solution, reasoners typically encode both relational gist, or subset relationships (e.g., that there are more cows than horses), and inclusional gist (e.g., that the cows are animals and the horses are animals), the latter perhaps when they are asked the inclusion question (e.g., Are there more animals or more cows?). The evidence for encoding of inclusion relations is overwhelming, beginning with Piaget, and most recently, Experiment 4 of Brainerd and Reyna showing that forced encoding of the inclusion relations prior to questioning was indistinguishable from optional encoding. Studies that increase the salience of the inclusional hierarchy, without making processing any easier, generally show little improvement in performance (e.g., Winer, 1980).

Therefore, contrary to the position ascribed by Howe and Rabinowitz to Brainerd and Reyna, encoding is *not* taken to be a source of difficulty. Instead, subjects are assumed to note, store, and have accessible the inclusion relations among classes, although the relational gist is more salient. Studies of other paradigms have shown that children typically encode relational gist (e.g., that there are more cows than horses) automatically whenever sets of objects vary in quantity (Brainerd, 1981). Thus, the presence of quantitative information is pivotal in creating the illusion because it makes relational gist salient, contributing to a tendency to judge relative numerosity, as opposed to ignoring numerosity information and directly applying the cardinality principle.

THE CARDINALITY PRINCIPLE: LOGICAL COMPETENCE

In class-inclusion problems, quantitative information is misleading in the sense that it is not necessary to solve the problem, and, worse, it induces a systematic illusion of judgment. However, the cardinality principle guarantees logically that judgments of relative numerosity should be *consistent* with inclusional relations. One procedure for solving the problem, direct inference from cardinality, is not inherently superior to the other, quantitative judgment. Both are algorithmic in principle, and the former entails the latter, although, in practice, logical inference may be unreliable (Reyna and Brainerd, 1990). In fact, children may use quantitative information to verify that the relationships among classes satisfy the cardinality principle. Such children will experience the same difficulties as those who simply judge relative numerosity, however. The levels in the inclusion hierarchy, the *vertical* relationships among classes, are

not distinctly represented, and this lack of clarity makes it harder to map the principle onto the facts.

Therefore, someone who knows the cardinality principle should have no trepidation, on the contrary, in answering the class-inclusion question as stated (i.e., by judging relative numerosity). The child who understands cardinality has no reason to suspect that approaching the question straightforwardly (as a relative numerosity judgment), or attempting to verify that the problem is an instance of the cardinality principle, creates systematic processing derailments. (Note that we do not claim that children count when making judgments of relative numerosity.) Consistent with this account, merely asking children to wait 15 s before responding enhances performance (see Winer, 1980). Clearly, logical competence could not be developing within such a period. On the other hand, reflection encourages reasoners to reconsider their approach to problems; competence enables them to notice that something is amiss in their solution. Correction, however, depends on abandoning what ought to be a valid procedure.

Thus, successful performance involves applying the cardinality principle, but also suppressing solutions based on numerical information. As Brainerd and Reyna showed, children can demonstrate their knowledge of the cardinality principle outside of the misleading context of the class-inclusion paradigm. Howe and Rabinowitz counter that children can answer such questions as whether there are more people in Tucson or in Arizona by "simple retrieval of domain-specific knowledge." There are three arguments against such a conclusion. First, and most obvious, generally, children of the age in question simply do not have such knowledge. As survey after survey in the United States shows, high school students, let alone kindergartners, do not know such facts as populations, or facts that might allow one to infer relative populations (without using the cardinality principle). For cardinality questions that were not like the Tucson-Arizona example, it is not clear how specific knowledge would be helpful in answering correctly. Note that arguing that young children might know the number of children in their class (but probably not the number in their school), but they know that there must be more children in the school than in their class is the same as arguing that they apply the cardinality principle. Nothing more is meant by the cardinality principle other than this qualitative idea that when one class includes another, it has more (or the same).

Perhaps the kind of knowledge that children retrieve that allows them to solve cardinality questions is their experiences, for example involving "Tucson" and "Arizona." If children retrieve pairings of people and Tucson (or people and the word "Tucson," or people and locations in Tucson, etc.) and compare their frequency to corresponding pairings of

people and Arizona, however, the former would outnumber the latter. Experiences and verbal references are more likely to be linked to less inclusive locales (one's town as opposed to country; one's class as opposed to school district), so instances paired with the less inclusive class should be more numerous in one's experience. Similarly, for nonlocation questions, basic-level words are used more frequently than superordinates, so instances labeled with the less inclusive word (e.g., "cat") should have been experienced more often than corresponding superordinate words (e.g., "animal"), especially for children. Therefore, availability in memory of experiences or verbal references to hierarchically related classes should be inversely proportional to their numerosity. Retrieval of experiences, then, would be misleading regarding the relative numerosity of nested classes.

Second, there is no evidence that factual knowledge about the sizes of classes can be used to inform relative numerosity judgments involving those classes; indeed, there is evidence to the contrary (Brainerd & Kaszor, 1974; Brainerd & Reyna, 1990a). If children use knowledge about the counts in classes to inform relative numerosity judgments, they should be able to solve class-inclusion problems. Numerous studies have shown that children accurately remember the counts of both subsets and the superordinate set, and yet they fail class inclusion.

Finally, failing to appreciate the quantitative implications of inclusion relations does not explain poor performance in class-inclusion tasks. Adults understand such implications, but their judgments still violate those implications, as in the conjunction fallacy. Howe and Rabinowitz's data (their Table 1) confirm that "it was not until adolescence (grade 9 and college)" that subjects could pass class-inclusion tests. Seventh graders were still "treating the question as a subclass-subclass problem." These results are not unique. As discussed in Brainerd and Reyna, the class-inclusion paradigm is distinguished from other so-called operational tasks in showing late development. In some studies, even error rates for college students are quite high (Brainerd, 1978). In order to use failure to appreciate the cardinality principle to explain the difficulty of class inclusion, one must assume that older subjects do not understand the cardinality principle, an implausible hypothesis. The difficulties experienced as late as seventh grade in the class-inclusion paradigm and by adults in the conjunction fallacy paradigm must be explained by something other than an inability to understand the implications of inclusion for quantitative judgments.

RETRIEVAL NECESSITY AND PROCESSING INTERFERENCE

As concluded in Brainerd and Reyna (1990a), therefore, it seems likely that encoding and knowledge of the cardinality principle are not the major

sources of difficulty in class inclusion. Knowing and being able to apply the cardinality principle, however, is not the same as retrieving it in context (Overton, 1990; Reyna & Brainerd, 1991). As Brainerd and Reyna demonstrated, younger children benefited from retrieval cues prompting the cardinality principle. For 5-year-olds, retrieval prompts alone increased the number of correct responses (out of 20) to class-inclusion questions from 2.4 to 8.6 in Experiment 5 and to 9.2 in Experiment 3. Note that the low level of performance for 5-year-olds under standard conditions (e.g., 2.4 in Experiment 5) indicates that retrieval of the correct problem-solving principle was essentially nil. For these young children, who apparently do not spontaneously retrieve the relevant principle, gist manipulations were ineffective in the absence of retrieval cues (Experiments 1, 2, and optional encoding conditions of Experiment 3). Such results indicate that retrieval is a necessary step, since without it problem solving does not go forward. On the other hand, when spontaneously higher levels of retrieval in older children (as shown by their data) are then augmented by retrieval prompts, and retrieval of the cardinality principle is combined with gist salience manipulations, performance becomes nearly perfect (e.g., for 7-year-olds, 19.6, and 19.2 out of 20). Thus, what separates seventh-graders who do not "understand" class inclusion at all, according to Rabinowitz et al., from 7-year-olds who score almost perfectly are just two types of manipulations: A retrieval prompt, and a cuing manipulation that makes inclusion relations easier to process.

Our characterization of successful reasoning in the standard task, then, consists of encoding the inclusion relations that are implicit in the problem information, knowing and retrieving the cardinality principle, and combining the inclusion relations with the cardinality principle to infer relative numerosity (processing). In the standard paradigm, reasoners encode both relational gist (subset comparisons) and inclusional gist (superordinate-subset comparisons). The relational gist, however, is more salient than the inclusional gist. The inclusion relations, though well understood, are implicit, whereas the subset relations are represented explicitly; the latter can be "seen." The relational gist, then, competes with, and dominates, the inclusional gist as a way to see the problem information (Reyna & Brainerd, 1991). Contrary to Howe and Rabinowitz's description of our theory, relational gist is not assumed to be selectively encoded or selectively stored: It is not that reasoners are unaware of the inclusion relations, but rather that *one set of relations is figure and the other is ground*.

Moreover, seeing one set of relations as figure tends to drive the other toward ground, and vice versa; these are incompatible ways of viewing the problems. This is what we intended to convey when we stated that "it

is hard to preserve a numerical whole in memory when one of its parts has been dissevered from it," that there is an inherent difficulty produced by part-whole relationships. One tends to view the cows as cows (as opposed to horses) or as animals, but it is difficult to see them as both *simultaneously*. Once the cows have been assigned as cows, they are no longer available to be assigned as animals. So, as processing fixes on animals, cows and horses recede, but as processing shifts to cows (or to horses), animals recede. This difficulty should produce longer response times to either minor or major subclass-superordinate class comparisons, compared to subclass-subclass comparisons (see Rabinowitz et al. Table 5, pp. 389-390, and pp. 401 for confirmatory data).

Thus, if one attempts to use inclusional gist to solve the problem, since all of the items are members of the superordinate set, when the superordinate set is considered, subclasses recede, and there is no other set accessible for comparison. Conversely, once processing focuses on the subset mentioned in the question, the superordinate set recedes, and the question *appears* to involve nothing more than the relational gist, a subset-subset comparison. Therefore, even if the reasoner retrieves and attempts to apply the cardinality principle, the inclusional gist is elusive. In standard class-inclusion tasks, reasoning seizes on the most salient representation of the problem, for which relative numerosity is only too "obvious."

EXPLANATIONS OF SUPERORDINATE-SET CUING AS SUBCLASS COMPARISON

When superordinate-set cuing is used, mental bookkeeping is simplified. The levels of the inclusion hierarchy are explicitly tagged. Howe and Rabinowitz claim, as others have (e.g., Dean et al., 1981), that superordinate-set cuing allows the child to make subclass comparisons, where the tags are a "third subclass," rather than making subclass-superordinate comparisons. There are four arguments against such a conclusion, however. First, children know that cows are animals (i.e., that animals are not another class on the same level as cows, but includes cows). Therefore, when children answer a question about cows and animals by using tags to invoke the cardinality principle, the tags are merely placeholders for "animals." Unless children suddenly forget that animals are a more inclusive set than cows, the tags cannot refer to a class on the same level as cows. In other words, children are aware of three classes (e.g., animals, cows, and horses), *but they realize that animals are a more inclusive set than cows*. Because the child is not thinking of the third class as another subclass, on the same order as cows, the comparison is not subclass to subclass.

Second, the argument for subclass comparisons is based on the suppo-

sition that, in the cued condition, children reinterpret the standard question "Are there more animals or more cows?" as, for example, "Are there more things with hats or more cows?". (The question form used in Brainerd and Reyna is the standard one, not the form used in Wilkinson, 1976, in which cues were explicitly mentioned.) To make such a conversion, however, the child must substitute "things with hats" for "animals," which presupposes that the child identifies the things with hats with animals. If things with hats were yet another subclass, along with cows and horses, it would remain to be explained how they had become identified with animals in the child's mind. Since the question *mentions nothing about hats*, in order to compare hats as a subclass to cows, some connection must be made between hats and the word "animals" in the question. The obvious explanation, and the one we favor, is that children see that all the animals have hats, so they use hats as placeholders for animals. If we ignore the obvious explanation, and assume instead that "animals" is ambiguous and receives a subclass interpretation here (as in "other animals"), such an interpretation would prohibit an identification with hats. All of the objects have hats, including, notably, the cows to which animals are to be compared. There is no third class, no "other animals," *to which hats could be linked* in order to interpret the question as referring to subclasses. Note that if the "other animals" were horses (ignoring for the moment that the cows are wearing hats too), children would give the wrong response (cows), and they give the right one (animals). Such arguments amount to claiming that children do not know that the things with hats are animals, despite the fact that they connect "animals" in the question to hats, or, alternatively, that "animals" is somehow reinterpreted so that even though hats are connected to "animals," "animals" really refers to a subclass. The former arguments suffer from a failure to explain how the question is understood as referring to hats, and the latter arguments founder on the absence of any reinterpretation as subclass that corresponds to the data.

As in Wilkinson (1976), Dean et al.'s (1981) class-inclusion questions mentioned perceptual features. However, the tasks were constructed such that focusing on features led to incorrect responses. Prior to each task, children were required to count the misleading features. In addition, the mapping of features to classes switched across tasks. For example, for some questions, windows were mapped onto the superordinate and doors onto the subclass, but, for other questions, the more numerous windows characterized all but one of the houses, so the windows-doors mapping was reversed, and not in a way that was immediately noticeable. (Some facilitating effect of cuing may be present in Dean et al.'s data, but baseline performance cannot be determined due to the way the data are reported.)

Due to the nontrivial changes made to the procedure, therefore, Dean et al.'s conditions cannot be directly compared to Wilkinson's superordinate-set cuing condition. Brainerd and Reyna's salience conditions, however, are comparable to Wilkinson's, except for the change in question, and a facilitating effect of superordinate-set cuing was retained (see Experiments 1, 2, and 6). Thus, Brainerd and Reyna's results are not subject to the criticism that the question allows the child to compare features. Standard questions (e.g., "Are there more animals or more cows?"), used in Experiments 1, 2, and 6, mention only classes, rather than perceptual cues for either subclass or superordinate. Therefore, Dean et al.'s results do not support a subclass-comparison explanation for superordinate-set cuing, and Howe and Rabinowitz's subclass explanation cannot account for children's ability to interpret questions in which cues are not mentioned. (For arguments regarding what knowledge is required for connecting classes to cues, I refer the reader to the discussion above.)

The last, and most telling, argument against subclass comparisons is the Age \times Cuing interaction. If superordinate set cuing permitted subclass comparisons, a less advanced strategy, it should produce greater improvement in younger children. The opposite trend is observed; cuing is more effective in older than in younger children. This is inconsistent with any account suggesting that cuing allows the problems to be solved with a less advanced strategy.

Manipulations that simplify processing, but do not affect retrieval, then, are bound to have a bigger effect on older children. Enhancing retrieval, on the other hand, should improve performance as long as some retrieval failure is occurring. Therefore, retrieval of the cardinality principle is necessary, but not sufficient, for correct performance. Retrieval enhancements can produce correct performance when processing difficulties can be circumvented, however, either by disentangling the processing of relative numerosity (e.g., by tagging the superordinate set), or by directly invoking the cardinality principle while resisting the dissonant conclusion suggested by quantitative relationships. The ability to resist interference increases developmentally (e.g., Dempster, 1990; Reyna & Brainerd, 1989). The effects of manipulations in Brainerd and Reyna can thus be neatly summarized as retrieval enhancements (effective throughout the age range tested) or processing enhancements (more effective in older subjects).

WORKING-MEMORY EXPLANATIONS OF CLASS-INCLUSION PERFORMANCE

Howe and Rabinowitz claim that the source of reasoning failures in children and adults is their inability to process the class-inclusion ques-

tion because of "an excessive cognitive workload" that demands willed attention. They agree with Brainerd and Reyna that "young children do encode, store, and retrieve subclass-superordinate class relations in class-inclusion tasks." They also say that it is implausible that young children fail to possess the *cardinal ordering principle*, citing Gelman and Gallistel (1978), and go on to say that it is not clear what this principle has to do with problems of class inclusion which have to do, instead, with inclusiveness. The ability to use inclusiveness, in turn, depends on subject and task characteristics, with younger subjects unable to understand and interpret inclusiveness questions, and tasks more difficult when they place heavy demands on willed attention.

Howe and Rabinowitz also suggest that "salience" (presumably cuing) and retrieval manipulations might also affect and interact with encoding or "what is stored," and cite evidence from paradigms other than class inclusion in which "stimulus context" and instructions affected encoding, storage, and/or retrieval. Howe and Rabinowitz then state that willed attention is unnecessary when processing becomes automatic, and that processing becomes more automatic with age. They introduce a five-parameter model with two "memory" parameters (e and d) and three processing parameters (u , s , and i). Performance is explained by the processing parameters u and s , which give the probability that subjects correctly understood the question or misinterpreted it as a subclass comparison, respectively. They then argue that memory load manipulations do not affect the memory parameters, but instead affect the processing parameters because of a tradeoff between reasoning and remembering; i.e., subjects conserve encoding and storage of information at the expense of their reasoning strategies. The erroneous strategy children settle on, according to Howe and Rabinowitz, is subclass comparison because it is "simpler."

First, although Howe and Rabinowitz claim no difficulty due to encoding, storage, or retrieval of inclusion relations, they ascribe beneficial effects of cuing to encoding, and effects of questions to encoding, retrieval, and storage of problem information. (Again, contrary to Howe and Rabinowitz's description, Brainerd and Reyna do not attribute difficulties to "encoding of the initial background facts," nor to erroneous storage or representation of background facts.) If children already encode, store, and retrieve background facts, in particular, inclusion relations, as Howe and Rabinowitz claim, improvements due to cuing or questions cannot be due to encoding, storage, or retrieval of such information. Numerous studies have shown that encoding manipulations produce little or no effect, in contrast to the cuing and question manipulations of Brainerd and Reyna. Also, according to Rabinowitz et al., for conditions comparable to Brainerd and Reyna's, "the estimated value of e and

d (the memory parameters) in these cases was .99 which does not differ reliably from 1" (p. 402). If encoding and memory are not deficient, it is unclear how manipulations that improve performance are improving encoding and memory. Further, the fact that such effects occur in paradigms other than class inclusion is immaterial. Therefore, it seems imprudent to give up the simplicity of the fuzzy-trace theory account of such manipulations for a more complicated version for which there is no evidence.

Second, the cardinal ordering principle of Gelman and Gallistel is not the cardinality principle of which we spoke. Neither do we mean simply inclusiveness, a relation among classes, the understanding of which is not at issue, nor numerosity per se. Instead, the cardinality principle refers to an implication of inclusiveness for relative numerosity, an implication that children understand (see above).

According to Howe and Rabinowitz, class-inclusion reasoning taxes working memory, placing demands on willed attention, the central construct of their theory. However, memory load manipulations had no effect on Rabinowitz et al.'s memory parameters: "Even though memory load was varied, only 4 of 14 possible memory-parameter comparisons between the No Load and either the Load or the Detection group were significant. Furthermore, one of these significant effects . . . reflected more accurate same-different encoding in the Load than in the No Load group" (pp. 403-404). Thus, the measures of memory for information pertinent to reasoning did not vary across Load and No Load conditions. In addition, models containing memory parameters failed necessity tests (although the usual class-inclusion errors occurred); a simpler model, without memory parameters, was able to account for class-inclusion performance under conditions comparable to those of Brainerd and Reyna's. Thus, the behavior of the memory parameters suggests that reasoning was independent of memory in Rabinowitz et al.'s data, although they never directly measured memory. Brainerd and Kingma (1985), however, measured both memory for problem information and reasoning in class inclusion, and found that they were independent.

Moreover, although error rates for reasoning in Rabinowitz et al. were high, memory was quite good, as the memory parameter values (.99) and results from their pilot study attest. In the pilot study, children were asked about information presented in class-inclusion problems, and "inasmuch as they almost always correctly stated the number of items in the superordinate class immediately following these incorrect class-inclusion answers, it appears that they appropriately encoded the superordinate class" (p. 381). Although memory parameters did not vary in Rabinowitz et al.'s experiments, the processing parameters did, however: "In contrast to the small and inconsistent effects of memory load on the memory parameters, the effects of memory load on the interpretation parameters

were large and consistent'' (p. 404). As in Brainerd and Reyna's Experiment 6, Load (or array absence) does not appear to cause memory deficits that interfere with performance, although processing can be affected under some conditions. As Howe and Rabinowitz acknowledge, variations in performance, and the primary sources of children's failures in class-inclusion reasoning, were linked to the processing parameters (u and s).

In brief, the data that Howe and Rabinowitz refer to show that, for conditions comparable to Brainerd and Reyna's, (a) memory parameters did not vary, and models without them were sufficient, (b) the probability of memory failure was nil and could not explain reasoning errors, and (c) the effects that were obtained were associated with processing parameters. Howe and Rabinowitz's data make a strong case that memory, or willed attention, has little to do with performance in class-inclusion problems. I wholeheartedly agree with Howe and Rabinowitz that formal models representing precisely defined constructs are desirable. Clarity is not achieved simply by having a formal model, however. The model also consists of the interpretation of parameters, and in this respect, willed attention is wanting.

Indeed, Howe and Rabinowitz's data, and the conclusions about processing that follow from them, are highly consistent with those of Brainerd and Reyna. By fourth grade, the youngest age level tested by Rabinowitz et al., Brainerd and Reyna's results indicated that retrieval (of the cardinality principle) was no longer a source of difficulty, and that performance deficits under standard conditions could be explained by processing problems (Experiments 1 and 2). Thus, providing fourth-graders with superordinate-set cues brings their performance to near-ceiling levels. It is only in subjects younger than fourth-graders (not tested by Rabinowitz et al.) that retrieval failure appears to be a factor in performance.

QUESTION WORDING AND LINGUISTIC EXPLANATIONS OF CLASS INCLUSION

Howe and Rabinowitz also suggest that "linguistic cuing manipulations" used in Brainerd and Reyna improve performance in class-inclusion tasks by allowing subclass comparisons. Although they do not go into detail, "labeling the common perceptual feature" appears to be the key to facilitation, in their view. Brainerd and Reyna used *two* retrieval manipulations (combining them in Experiment 6), only one of which involved mentioning perceptual cues in questions. I will call the latter the "garden-path" question (used in Experiments 3 and 6), and return to it shortly. The other retrieval manipulation involved presenting numerical information illustrating the cardinality principle (e.g., "There

are eight animals, five cows, and three horses"; Experiment 5 and all conditions of Experiment 6). This information highlights the quantitative relationships among nested classes. As Brainerd and Reyna noted, virtually any theory of retrieval suggests that confronting an instance of a principle (e.g., that animals outnumber cows) increases the probability of retrieving that principle (e.g., that the more inclusive class outnumbers, or equals, its subclasses).

Winer (1974) employed a similar manipulation with like effects, calling it "verbal facilitation." Winer also obtained a classic retrieval pattern of results: Children who did not receive numerical cues in the first half of the problems improved when given cues in the second half, but those who received numerical cues first continued to do well in the second half of the problems, in which cues were absent. Once retrieval of a principle is successfully cued, its relevance to the context noted, that realization would be expected to influence subsequent performance.

Winer (1974), Howe and Rabinowitz, and others, however, ascribe such facilitative effects to linguistic factors, i.e., the interpretation of class-inclusion questions, rather than retrieval. Winer, for example, discussed the possibility of an "exclusive or" interpretation of the class-inclusion questions (e.g., cows as opposed to "other animals"), which 30% of adults favored. Shipley (1979) also argued that superordinate classes are given a narrower interpretation than intended, so the standard question is understood as involving coordinate classes (subclasses). I do not dispute that children give erroneous answers that reflect subclass relationships. But, it is not that linguistic principles cause the question to be interpreted so that children make errors. Rather, the question is ambiguous, and reasoning, in the grip of a cognitive illusion, results in an interpretation of the problem that is not prohibited by the question. This conclusion about the standard question is supported by three arguments: (a) that linguistic analyses suggest that the intended reading is more basic than the narrower interpretation, (b) that data from the class-inclusion paradigm rule out linguistic factors as causal, and (c) that, if the effect were due to such linguistic principles, developmental trends would be the opposite of those observed.

On its face, the linguistic explanation is compelling. Unfortunately, a fallacious leap is made from the fact that the question permits an interpretation in which the superordinate excludes the subclass to the conclusion that this causes the effect. Knowing that children seem to compare subclasses may make it appear that such an interpretation is linguistically favored for class-inclusion questions. Shipley (1979) made such an argument on the grounds that children misinterpret class-inclusion questions as "ungrammatical requests for distributive comparisons" (p. 321) when they are actually collective comparisons. Interpreting the question as a

distributive comparison invokes the constraint that such structures "usually refer to elements of coordinate classes" (p. 320). Therefore, class-inclusion questions would be interpreted as involving comparisons of coordinate classes. The implausible aspect of this argument is that children would give up a grammatical interpretation of a question (that they are capable of achieving) for an ungrammatical one. Shipley pointed out that, if children made this ungrammatical overgeneralization, such a result would parallel their performance in class-inclusion tasks, but she offered no evidence that children engage in ungrammatical overgeneralizations.

Shipley also attempted to show that, given ungrammatical distributive comparisons (rather than grammatical collective comparisons as in class inclusion), children interpret classes as coordinate. The data were inconclusive, however, because, according to Shipley, it is impossible to tell from children's responses what they were comparing, and there were few "revealing justifications" (p. 322). Linguistic manipulations involving the connotations of collective nouns also proved inconclusive in Dean et al. (1981). Finally, data from memory questions given after class-inclusion questions argue against any linguistic explanation assuming misinterpretation of superordinates as exclusive. For example, when asked what the class-inclusion question was, memory errors did not involve substituting subclasses for superordinates, as the linguistic theory would expect (Brainerd & Kingma, 1985). That is, even though children responded in class inclusion as if they were asked whether there were more cows than horses, they did not recollect the question as being about cows and horses, rather than cows and animals (Brainerd & Kingma, 1985). In addition, if superordinate terms were interpreted exclusively, memory questions about the number in that class should elicit the number of a subclass. Answers to such questions revealed, instead, that children interpreted the referent of the superordinate term as the inclusive set (Brainerd & Kingma, 1985).

Thus, neither linguistic analysis nor data support Shipley's thesis. In commenting on her own efforts, Shipley provided an excellent summary of research on linguistic manipulations in this paradigm. As she noted, linguistic factors are not the primary source of difficulties in the class-inclusion task because "a majority of children do persist in making errors when conventional wording is introduced. This suggests there is some factor in addition to wording which interferes with correct performance in the class-inclusion task" (p. 317).

Linguistic analysis, in fact, does not suggest that class-inclusion questions are stacked in favor of coordinate-class or exclusive "or" interpretations. Winer's data, for example, show that the majority of adults (70%) interpreted class-inclusion questions as involving an inclusive "or" (in

which the superordinate included the subclass). As Sadock (1978) explains, "consider the well-known property of English *or* that it can, but need not, be used in an exclusive understanding. . . . *or* need not be treated as an ambiguous lexical item but may be assigned a conventional meaning exactly like logical (inclusive) disjunction. The fact that it often suggests the exclusion of conjunction would then be a conversational implicature" (p. 291). Thus, the literal meaning of "or" is inclusive, but pragmatic principles might allow one to infer an exclusive interpretation in context.

On this analysis, to account for developmental trends, one would have to argue that young children are more apt to make pragmatic inferences that go beyond what is literally stated than adults are, and data on the development of pragmatic linguistic inferences suggest the opposite. That is, adults would be more likely to suppose that the questioner could not be asking a question to which he already knows the answer (because it is implicit in the meaning of the words), so "animals" must refer to "other animals," excluding cows. Given such an interpretation, adults would commit the class-inclusion "error" of responding with the larger subclass. Moreover, since the ability to make such pragmatic inferences (that depend, in part, on assessing what the speaker already knows) increases with age, larger-subclass interpretations should increase with age, and they decrease. Even if young children were as willing as adults to make such inferences, linguistic inferences do not explain why they are more likely than adults to make subclass errors.

For the standard class-inclusion question, therefore, linguistic analyses suggest an inclusive interpretation of the superordinate class as the basic reading. Adjectives, however, can be used to induce an exclusive interpretation of a category, as in the garden-path questions used in Experiments 3 and 6 of Brainerd and Reyna. As Clark and Clark (1977) explain, "adjectives normally narrow down the category the noun picks out. *Car* names a broad category of vehicles, while *red car* narrows this category down to one with a specific color. . . . [The use of such adjectives] depends on what the target object is to be distinguished from. If a car is the target and it is to be distinguished from a truck, *car* alone will do. But if it is to be distinguished from another car, speakers must find a dimension along which the two differ—for example, color" (pp. 472–473).

The garden-path question (e.g., Are there more red animals or more cows?), then, indicates that there is a category narrower than the superordinate, for example, that there are red animals as opposed to nonred animals. However, all of the members of the superordinate have the feature described in the question (e.g., all of the animals are red). The array contradicts the implications of the adjective (or adjectival phrase) in the question. Therefore, when the reasoner attempts to identify the target

objects (e.g., red animals as opposed to nonred ones), she is led to discover that the “distinguishing” feature characterizes all of the objects. All of the objects are red animals, for example, and there are more red animals than anything else. Knowing that cows are animals, the reasoner can see that there are more red animals than cows; the “moreness” of the red animals is apparent in the display—as long as the observer includes cows among the red animals.

Thus, the question suggests an exclusive interpretation of the superordinate that is cancelled by the display. However, the process of comprehending the question, picking out the target objects, divulges the fact that the distinguishing feature tags the entire superordinate set. This realization prompts the reasoner to see the cardinality principle exemplified (e.g., that there are more red animals than cows), and, as before, an instance encourages the retrieval of a principle. And processing is facilitated as well in these circumstances because the superordinate set is cued. Ironically, with the array present, an exclusive interpretation of the superordinate that turns out to be mistaken (hence the term garden-path) facilitates performance on class inclusion.

CONCLUSIONS

Class-inclusion errors are a curious phenomenon because they should not happen at all, and, yet, they persist long after other operational tasks have been mastered. Prior to fuzzy-trace theory, a variety of intuitively compelling explanations were offered for this phenomenon. I have discussed some of these explanations here, including logical incompetence, working-memory overload, and misleading questions. Experimental manipulations from these perspectives have only nibbled at the edges of the effect, however, accounting for small amounts of variance. Worse, data have been generated that disallow these explanations, however compelling they seem. Thus, observers have agreed that children make subclass comparisons in the standard task, but the reasons for such erroneous comparisons remained an enigma.

The confusion in class inclusion, we have suggested, does not come from confusing wording because children initially understand the question as intended; it is not a result of an overtaxed memory because, on the contrary, memory is well preserved; and, it is not produced by a failure to appreciate the implications of inclusion relations for relative quantity since that appreciation can be demonstrated outside of the class-inclusion task. Instead, a particular confluence of factors creates a cognitive illusion that leads reasoners astray. Any effort to judge relative numerosity in such problems will founder on the difficulty of processing the vertical class relations. There is no placeholder unambiguously corresponding to the superordinate level of the inclusion hierarchy. Therefore, although

children understand the vertical relations well enough, it is difficult to apply that understanding in processing. In contrast, the horizontal relationships are salient; they are derived automatically, and their levels are conspicuous. The elusive vertical relationships, then, are supplanted by the salient horizontal ones. Meanwhile, the ambiguity of the class-inclusion question allows the substitution to take place, although it does not create the illusion.

Such explanations allow us to understand how people can fail to realize the implications of their own knowledge. Successful problem solving involves selecting from among many relationships given as background facts, retrieving some among many principles that could be applied to such relationships, and, finally, applying the principles coherently. In class inclusion, responding to a question about which has more, younger children sensibly retrieve the relative numerosity principle. Older children get a bit further. Realizing that inclusion bears on the judgment of relative quantity, they retrieve the cardinality principle, but err in its application. For both younger and older, however, "correct" principles are applied to "true" facts from the problem.

By this account of class-inclusion errors, however, I do not mean to imply that children's answers are really correct, or that cognitive illusions absolve the reasoner from error. On the contrary, no matter how compelling the illusions, animals are more numerous than cows and being a bank teller is more probable than being a feminist bank teller. Paradoxically, the victims of such illusions adhere to the principles expressed in these examples. Class-inclusion errors, the conjunction fallacy, and other cognitive illusions, then, teach us to look beyond incompetence for explanations of reasoning errors, to the ordinary modes of thought of competent reasoners.

REFERENCES

- Brainerd, C. J. (1978). *Piaget's theory of intelligence*. Englewood Cliffs, NJ: Prentice Hall.
- Brainerd, C. J. (1981). Working memory and the developmental analysis of probability judgment. *Psychological Review*, 88, 463-502.
- Brainerd, C. J., & Kaszor, P. (1974). An analysis of two proposed sources of class-inclusion errors. *Developmental Psychology*, 4, 633-643.
- Brainerd, C. J., & Kingma, J. (1985). On the independence of short-term memory and working memory in cognitive development. *Cognitive Psychology*, 17, 210-247.
- Brainerd, C. J., & Reyna, V. F. (1990a). Inclusion illusions: Fuzzy-trace theory and perceptual salience effects in cognitive development. *Developmental Review*, 10, 365-403.
- Brainerd, C. J., & Reyna (1990b). Gist is the grist: Fuzzy-trace theory and the new intuitionism. *Developmental Review*, 10, 3-47.
- Clark, H. H., & Clark, E. V. (1977). *Psychology and language: An introduction to psycholinguistics*. New York: Harcourt Brace Jovanovich.
- Dean, A. L., Chabaud, S., & Bridges, E. (1981). Classes, collections, and distinctive features: Alternative strategies for solving inclusion problems. *Cognitive Psychology*, 13, 84-112.

- Dempster, F. (1990). *The rise and fall of the inhibitory mechanism: Toward a unified theory of cognitive development and aging*. Manuscript submitted for publication.
- Falmagne, R. J. (Ed.) (1975). *Reasoning: Representation and process in children and adults*. Hillsdale, NJ: Erlbaum.
- Gelman, R., & Gallistel, R. (1978). *The child's understanding of number*. Cambridge, MA: Harvard University Press.
- Howe, M. L., & Rabinowitz, F. M. (1991). Gist another panacea? Or just the illusion of inclusion. *Developmental Review*, 11, 305-316.
- Overton, W. F. (1990). Competence and procedures: Constraints on the development of logical reasoning. In W. F. Overton (Ed.), *Reasoning, necessity, and logic: Developmental perspectives*. Hillsdale, NJ: Erlbaum.
- Rabinowitz, F. M., Howe, M. L., & Lawrence, J. A. (1989). Class inclusion and working memory. *Journal of Experimental Child Psychology*, 48, 379-410.
- Reyna, V. F., & Brainerd, C. J. (1989). Output interference, generic resources, and cognitive development. *Journal of Experimental Child Psychology*, 47, 42-46.
- Reyna, V. F., & Brainerd, C. J. (1990). Fuzzy processing in transitivity development. *Annals of Operations Research*, 23, 37-63.
- Reyna, V. F., & Brainerd, C. J. (1991). Fuzzy-trace theory and children's acquisition of mathematical and scientific concepts. *Learning and Individual Differences*, 3, 27-59.
- Sadock, J. M. (1978). On testing for conversational implicature. In P. Cole (Ed.), *Syntax and semantics* (Vol. 9). New York: Academic Press.
- Shipley, E. F. (1979). The class-inclusion task: Question form and distributive comparisons. *Journal of Psycholinguistic Research*, 8, 301-331.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293-315.
- Wilkinson, A. (1976). Counting strategies and semantic analyses as applied to class inclusion. *Cognitive Psychology*, 8, 64-85.
- Winer, G. A. (1974). An analysis of verbal facilitation of class-inclusion reasoning. *Child Development*, 45, 224-227.
- Winer, G. A. (1980). Class-inclusion reasoning in children: A review of the empirical literature. *Child Development*, 51, 309-328.

RECEIVED: August 7, 1990; REVISED: August 20, 1990.