Category-Based Induction

Daniel N. Osherson
Massachusetts Institute of Technology

Edward E. Smith
University of Michigan

Ormond Wilkie
Massachusetts Institute of Technology

Alejandro López
University of Michigan

Eldar Shafir
Princeton University

An argument is categorical if its premises and conclusion are of the form *All members of C have property P,* where *C* is a natural category like FALCON or BIRD, and *P* remains the same across premises and conclusion. An example is *Grizzly bears love onions. Therefore, all bears love onions.* Such an argument is psychologically strong to the extent that belief in its premises engenders belief in its conclusion. A subclass of categorical arguments is examined, and the following hypothesis is advanced: The strength of a categorical argument increases with (a) the degree to which the premise categories are similar to the conclusion category and (b) the degree to which the premise categories are similar to members of the lowest level category that includes both the premise and the conclusion categories. A model based on this hypothesis accounts for 13 qualitative phenomena and the quantitative results of several experiments.

The Problem of Argument Strength

Fundamental to human thought is the confirmation relation, joining sentences $P_1 \ldots P_n$ to another sentence C just in case belief in the former leads to belief in the latter. Theories of confirmation may be cast in the terminology of argument strength, because $P_1 \ldots P_n$ confirm C only to the extent that $P_1 \ldots P_n / C$ is a strong argument. We here advance a partial theory of argument strength, hence of confirmation.

To begin, it will be useful to review the terminology of argument strength. By an *argument* is meant a finite list of sentences, the last of which is called the *conclusion* and the others its *premises.* Schematic arguments are written in the form $P_1 \ldots P_n / C$, whereas real arguments are written vertically, as in the following examples:

Grizzly bears love onions.
Polar bears love onions.
All bears love onions.   (1)

Owls prey on small rodents.
Rattlesnakes prey on small rodents.   (2)

An argument *A* is said to be *strong* for a person *S* just in case *S*s believing *A*'s premises causes *S* to believe *A*'s conclusion. Mere

belief in the conclusion of an argument (independently of its premises) is not sufficient for argument strength. For this reason, Argument 1 is stronger than Argument 2 for most people, even though the conclusion of Argument 2 is usually considered more probable than that of Argument 1. An extended discussion of the concept of argument strength is provided in Osherson, Smith, and Shafir (1986). It will be convenient to qualify an argument as strong, without reference to a particular person *S*, whenever the argument is strong for most people in a target population (e.g., American college students). We also say that $P_1 \ldots P_n$ confirm *C* if $P_1 \ldots P_n / C$ is strong.

An illuminating characterization of argument strength would represent a long step toward a theory of belief fixation and revision. Unfortunately, no general theory is yet in sight, and even partial theories are often open to elementary counterexamples (see Osherson et al., 1986). This article offers a hypothesis about the strength of a restricted set of arguments, exemplified by Arguments 1 and 2. The premises and conclusions of such arguments attribute a fixed property (e.g., *preys on small rodents*) to one or more categories (e.g., OWL and RATTLESNAKE).[1] The present study focuses on the role of categories in confirmation; the role of properties is not systematically investigated. In this sense, the model we advance concerns induction that is category based.

Category-based induction was first examined by Rips (1975). He studied the strength of single-premise arguments involving categories such as RABBIT and MOUSE, or EAGLE and BLUEJAY. The present investigation builds on one of the models that Rips discusses and applies it to a larger class of arguments.

Our discussion proceeds as follows. After defining the class of arguments to be considered in this article, and introducing some relevant terminology, we document a set of 13 qualitative

[1] We use capitals to denote categories. Properties are italicized.

phenomena that must be deduced by any adequate theory of argument strength. Our own theory is then presented and shown to account for all of the phenomena. Next, we describe several experiments designed to test the theory quantitatively. Refinements and alternatives to the theory are discussed in the final section.

Arguments To Be Considered

An argument is called *categorical* just in case its premises and conclusion have the logical form *all members of X have property Y*, where X is a (psychologically) simple category like FALCON, VEHICLE, or MAMMAL, and Y remains fixed across premises and conclusion. Arguments 1 and 2 are categorical in this sense. The arguments discussed in this article are all categorical.

The property ascribed to the categories figuring in Argument 1 is *loves onions*. Subjects are likely to have prior beliefs about the kinds of animals that love onions, as well as prior beliefs about properties that are correlated with this one, such as *eats a wide variety of fruits and vegetables*. Beliefs such as these can be expected to weigh heavily on argument strength, defeating our goal of focusing on the role of categories in the transmission of belief from premises to conclusions. For this reason, the arguments to be examined all involve predicates about which subjects in our experiments have few beliefs, such as *requires biotin for hemoglobin synthesis*. Such predicates are called *blank*. Although blank predicates are recognizably scientific in character (in the latter case, biological), they are unlikely to evoke beliefs that cause one argument to have more strength than another.

In summary, the theories discussed below bear on categorical arguments involving natural kinds and blank predicates. An example of such an argument is

> Mosquitoes use the neurotransmitter Dihedron.
> Ants use the neurotransmitter Dihedron.
> ───
> Bees use the neurotransmitter Dihedron.   (3)

Henceforth, the term *argument* is to be understood in the foregoing sense.

Terminology and Notation

We assume that subjects (and experimenters) largely agree with each other about facts related to the hierarchical level of natural-kind concepts. To illustrate, wide agreement is presupposed about the following judgments:

1. FALCON and PELICAN are at the same hierarchical level;
2. BIRD is one level above both FALCON and PELICAN; and
3. ANIMAL is one level above BIRD.

We cannot presuppose universal agreement about such levels. For example, some subjects might take BIRD-OF-PREY to be one level above EAGLE, whereas others (who make fewer distinctions) might take BIRD to be one level above EAGLE. This kind of individual difference about fine-grained categories will be harmless in what follows. It is sufficient that agreement exists above salient categories such as EAGLE, BIRD, and ANIMAL.

Recall that all premises and conclusions to be discussed have the form *all members of X have property Y*. Given such a prem-

ise P or conclusion C, we denote the category that figures in P or C by CAT(P) or CAT(C). Thus, if P is the first premise of Argument 1 above, then CAT(P) = GRIZZLY BEAR. If C is the conclusion of Argument 3, then CAT(C) = BEE.

Let Argument $A = P_1 \ldots P_n/C$ be given. A is called *general* if CAT(P_1) ... CAT(P_n) are all properly included in CAT(C). For example, Argument 1 is general. A is called *specific* if any category that properly includes one of CAT(P_1) ... CAT(P_n), CAT(C) also properly includes the others. For example, Argument 3 is specific. By this definition, no argument is both general and specific. A is called *mixed* if A is neither general nor specific. The following argument is mixed:

> Flamingoes require titanium for normal muscle development.
> Mice require titanium for normal muscle development.
> ──
> All mammals require titanium for normal muscle   (4)
> development.

Argument 4 is not general because FLAMINGO is not included in MAMMAL. It is not specific because BIRD properly includes FLAMINGO but not MOUSE or MAMMAL. Argument 2 is also mixed.

Phenomena

General Remarks

Even within the restricted class of arguments at issue in this article, a variety of phenomena can be discerned that must be accounted for by any adequate theory of category-based induction. Each phenomenon signals the importance of a given variable in argument strength when other variables are held more or less constant. The phenomena should thus be conceived as tendencies rather than strict laws determining confirmation. We now present 13 such phenomena and illustrate each with a contrasting pair of arguments. The first argument in each pair is claimed to be stronger than the second, in conformity with the phenomenon that the pair illustrates. At the end of this section, we describe a study that empirically documents all of these claims about relative argument strength.

Phenomena Concerning General Arguments

Let general argument $P_1 \ldots P_n/C$ be given.

Phenomenon 1 (premise typicality). The more representative or typical CAT(P_1) ... CAT(P_n) are of CAT(C), the more $P_1 \ldots P_n$ confirm C. Because robins are more typical than penguins of BIRD, this phenomenon is illustrated by the following pair of arguments:

> Robins have a higher potassium concentration in their blood than humans.
> ───
> All birds have a higher potassium concentration in   (5a)
> their blood than humans.

> Penguins have a higher potassium concentration in their blood than humans.
> ───
> All birds have a higher potassium concentration in   (5b)
> their blood than humans.

The foregoing arguments have single premises. Multiple-prem-

ise illustrations of the same point are easy to construct. A premise typicality effect for social categories has been reported by Rothbart and Lewis (1988, see also Collins & Michalski, 1989).

Phenomenon 2 (premise diversity). The less similar CAT(P_1) ... CAT(P_n) are among themselves, the more P_1 ... P_n confirm C. Thus, since hippos and hamsters differ from each other more than do hippos and rhinos, the following arguments illustrate the premise diversity phenomenon:

> Hippopotamuses have a higher sodium concentration in their blood than humans.
> Hamsters have a higher sodium concentration in their blood than humans.
> ———————————————————————————
> All mammals have a higher sodium concentration in their blood than humans. (6a)

> Hippopotamuses have a higher sodium concentration in their blood than humans.
> Rhinoceroses have a higher sodium concentration in their blood than humans.
> ———————————————————————————
> All mammals have a higher sodium concentration in their blood than humans. (6b)

Observe that Argument 6a is stronger than Argument 6b even though hamsters are less typical than rhinoceroses of MAMMAL. Thus, the greater diversity of the premise categories in Argument 6a outweighs the greater typicality of the premise categories in Argument 6b.

Phenomenon 3 (conclusion specificity). The more specific is CAT(C), the more C is confirmed by P_1 ... P_n. Thus, because BIRDS is a more specific category than ANIMAL, this phenomenon is illustrated by the following pair of arguments:

> Bluejays require Vitamin K for the liver to function.
> Falcons require Vitamin K for the liver to function.
> ———————————————————————————
> All birds require Vitamin K for the liver to function. (7a)

> Bluejays require Vitamin K for the liver to function.
> Falcons require Vitamin K for the liver to function.
> ———————————————————————————
> All animals require Vitamin K for the liver to function. (7b)

A phenomenon related to conclusion specificity is reported by Gelman (1988, p. 78) in a developmental study.

Phenomenon 4 (premise monotonicity). For general arguments, more-inclusive sets of premises yield more strength than less inclusive sets. The following pair of arguments illustrates this kind of monotonicity:

> Hawks have sesamoid bones.
> Sparrows have sesamoid bones.
> Eagles have sesamoid bones.
> ———————————————————————————
> All birds have sesamoid bones. (8a)

> Sparrows have sesamoid bones.
> Eagles have sesamoid bones.
> ———————————————————————————
> All birds have sesamoid bones. (8b)

Premise monotonicity has been investigated by Carey (1985).

Phenomena Concerning Specific Arguments

Let specific argument P_1 ... P_n/C be given.

Phenomenon 5 (premise-conclusion similarity). The more similar CAT(P_1) ... CAT(P_n) are to CAT(C), the more P_1 ... P_n

confirm C. Because robins and bluejays resemble sparrows more than they resemble geese, this phenomenon is illustrated by the following pair of arguments:

> Robins use serotonin as a neurotransmitter.
> Bluejays use serotonin as a neurotransmitter.
> ———————————————————————————
> Sparrows use serotonin as a neurotransmitter. (9a)

> Robins use serotonin as a neurotransmitter.
> Bluejays use serotonin as a neurotransmitter.
> ———————————————————————————
> Geese use serotonin as a neurotransmitter. (9b)

The present phenomenon was originally reported by Rips (1975) for single-premise argument. (see also Collins & Michalski, 1989).

Phenomenon 6 (premise diversity). The less similar CAT(P_1) ... CAT(P_n) are among themselves, the more P_1 ... P_n confirm C. Illustration is provided by the following pair inasmuch as lions are less similar to giraffes than than they are to tigers.

> Lions use norepinephrine as a neurotransmitter.
> Giraffes use norepinephrine as a neurotransmitter.
> ———————————————————————————
> Rabbits use norepinephrine as a neurotransmitter. (10a)

> Lions use norepinephrine as a neurotransmitter.
> Tigers use norepinephrine as a neurotransmitter.
> ———————————————————————————
> Rabbits use norepinephrine as a neurotransmitter. (10b)

Phenomenon 6 corresponds to Phenomenon 2 for general arguments. Observe that Argument 10a is stronger than Argument 10b even though giraffes resemble rabbits no more than do tigers.

Phenomenon 7 (premise monotonicity). More inclusive sets of premises yield more strength than less inclusive sets, provided that the new premise is drawn from the lowest level category that includes the old premises and conclusion. The following pair of arguments illustrates this kind of monotonicity.

> Foxes use Vitamin K to produce clotting agents in their blood.
> Pigs use Vitamin K to produce clotting agents in their blood.
> Wolves use Vitamin K to produce clotting agents in their blood.
> ———————————————————————————
> Gorillas use Vitamin K to produce clotting agents in their blood. (11a)

> Pigs use Vitamin K to produce clotting agents in their blood.
> Wolves use Vitamin K to produce clotting agents in their blood.
> ———————————————————————————
> Gorillas use Vitamin K to produce clotting agents in their blood. (11b)

Phenomenon 8 (premise-conclusion asymmetry). Single-premise arguments are not symmetric, in the sense that P/C may not have the same strength as C/P. This kind of asymmetry is illustrated by the following pair of arguments:

> Mice have a lower body temperature at infancy than at maturity.
> ———————————————————————————
> Bats have a lower body temperature at infancy than at maturity. (12a)

Bats have a lower body temperature at infancy than at maturity.
———————————————————————————————————
Mice have a lower body temperature at infancy than at maturity.   (12b)

Premise–conclusion asymmetry was first discussed by Rips (1975).

Phenomena Concerning Mixed Arguments

Phenomenon 9 (nonmontonicity-general). Some general arguments can be made weaker by adding a premise that converts them into mixed arguments. This kind of nonmonotonicity is illustrated by the following contrast:

Crows secrete uric acid crystals.
Peacocks secrete uric acid crystals.
———————————————————————————————————
All birds secrete uric acid crystals.   (13a)

Crows secrete uric acid crystals.
Peacocks secrete uric acid crystals.
Rabbits secrete uric acid crystals.
———————————————————————————————————
All birds secrete uric acid crystals.   (13b)

Phenomenon 10 (nonmonotonicity-specific). Some specific arguments can be made weaker by adding a premise that converts them into mixed arguments. This kind of nonmonotonicity is illustrated as follows.

Flies require trace amounts of magnesium for reproduction.
———————————————————————————————————
Bees require trace amounts of magnesium for reproduction.   (14a)

Flies require trace amounts of magnesium for reproduction.
Orangutans require trace amounts of magnesium for reproduction.
———————————————————————————————————
Bees require trace amounts of magnesium for reproduction.   (14b)

A Phenomenon Involving Both General and Specific Arguments

Phenomenon 11 (inclusion fallacy). A specific argument can sometimes be made stronger by increasing the generality of its conclusion. Because BIRD includes OSTRICH, this phenomenon is illustrated as follows:

Robins have an ulnar artery.
———————————————————————————————————
Birds have an ulnar artery.   (15a)

Robins have an ulnar artery.
———————————————————————————————————
Ostriches have an ulnar artery.   (15b)

The choice of Argument 15a as stronger than Argument 15b is counternormative and may be termed an *inclusion fallacy.* For discussion and analysis of inclusion fallacies in another context, see Shafir, Smith, and Osherson (in press).

Two Limiting-Case Phenomena

The last two phenomena to be discussed have an evident character, and no data are needed for their documentation.

Phenomenon 12 (premise–conclusion identity). Any argu-

ment of the form Q/Q is perfectly strong. One such argument is as follows:

Pelicans have property Y.
———————————————————————————————————
Pelicans have property Y.   (16)

Phenomenon 13 (premise–conclusion inclusion). Suppose that statements P and C are such that the conclusion category is included in the premise category. Then the argument P/C is perfectly strong. For example:

All animals have property Y.
———————————————————————————————————
All birds have property Y.   (17)

Table 1 summarizes all 13 phenomena.

Empirical Documentation of Phenomena 1–11

Two studies were performed to empirically document Phenomena 1–11. In Study 1, subjects were presented a 12-page booklet. The first page contained instructions, and each of the following pages contained one of the contrasting pairs of arguments used above to illustrate Phenomena 1–11. The instructions were as follows:

> We are interested in how people evaluate arguments. On each page of your booklet there will be two arguments labeled "a" and "b." Each will contain one, two, or three statements separated from a conclusion by a line. Assume that the statements above the line are facts, and choose the argument whose facts provide a better reason for believing its conclusion. These are subjective judgments; there are no right or wrong answers.

On each subsequent page, the contrasting pair of arguments was arranged vertically; across all subjects each argument appeared equally often in the upper and lower positions. The order of the argument pairs was randomized anew for each subject. The subjects were 80 University of Michigan undergraduates who were paid for their participation and tested in groups of 20.

The results of Study 1 are presented in Table 1. For each contrasting pair, the number of subjects choosing a given argument is shown in brackets next to the argument. In all cases but one, the majority choice is overwhelmingly for the argument we claimed in the preceding section to be stronger (these differences are significant at the .01 level by a two-tailed sign test). The sole exception is Phenomenon 8, premise–conclusion asymmetry, in which there is roughly an equal preference for the two contrasting arguments.

Postexperimental comments by some subjects suggested that the arguments constituting Phenomenon 8 were treated differently than other arguments. Because the arguments MOUSE/BAT and BAT/MOUSE contain identical statements, subjects apparently reasoned that there could be no difference in strength between them. This metacognitive strategy may have obscured the underlying difference in strength in which we are interested.

To respond to the foregoing difficulty, Study 2 was performed. Subjects were presented with a four-page booklet in which the first page contained instructions and each of the three test pages contained a contrasting pair of arguments. One contrasting pair was that used above to illustrate premise–conclusion asymmetry (Phenomenon 8) (viz., MOUSE/BAT versus BAT/MOUSE). The other two pairs were fillers. The instructions were designed to

Table 1
Summary of the 13 Phenomena

Phenomenon	Stronger argument (Version a)	Weaker argument (Version b)
General arguments		
1. Premise Typicality	ROBIN/BIRD [73]	PENGUIN/BIRD [7]
2. Premise Diversity	HIPPO, HAMSTER/ MAMMAL [59]	HIPPO, RHINO/MAMMAL [21]
3. Conclusion Specificity	BLUEJAY, FALCON/ BIRD [75]	BLUEJAY, FALCON/ANIMAL [5]
4. Premise Monotonicity	HAWK, SPARROW EAGLE/BIRD [75]	SPARROW, EAGLE/BIRD [5]
Specific arguments		
5. Premise–Conclusion Similarity	ROBIN, BLUEJAY/ SPARROW [76]	ROBIN, BLUEJAY/GOOSE [4]
6. Premise Diversity	LION, GIRAFFE/ RABBIT [52]	LION, TIGER/RABBIT [28]
7. Premise Monotonicity	FOX, PIG WOLF/GORILLA [66]	PIG, WOLF/GORILLA [14]
8. Premise–Conclusion Asymmetry	MICE/BAT [41] (40)	BAT/MICE [39] (20)
Mixed arguments		
9. Nonmonotonicity–General	CROW, PEACOCK/ BIRD [68]	CROW, PEACOCK RABBIT/BIRD [12]
10. Nonmonotonicity–Specific	FLY/BEE [51]	FLY, ORANGUTAN/BEE [29]
General and specific arguments		
11. Inclusion Fallacy	ROBIN/BIRD [52]	ROBIN/OSTRICH [28]
Limiting-case arguments		
12. Premise–Conclusion Identity		PELICAN/PELICAN
13. Premise–Conclusion Inclusion		ANIMAL/BIRD

Note. Number of subjects in Study 1 preferring each argument is given in brackets.
Entries in parentheses are results of Study 2.

suppress the metacognitive strategy. They were the same as the instructions of Study 1, except that the last sentence was replaced by

> Although the two arguments in a pair may sometimes seem very similar, there is always a difference in how much reason the facts of an argument give to believe its conclusion. However small this difference may be, we would like you to indicate for which argument the facts provide a better reason to believe the conclusion.

As before, each argument in a pair appeared equally often in the upper and lower positions, and the order of arguments was varied across subjects. The subjects were 60 University of Michigan undergraduates, paid for their participation and tested in groups of 20. None had participated in Study 1.

Most subjects preferred the MOUSE/BAT argument to the BAT/MOUSE argument. The difference, 40 versus 20, is significant at the .01 level by a sign test (two-tailed).

Replications

We have replicated the foregoing results—including the preference for MOUSE/BAT over BAT/MOUSE—in several studies using sets of argument pairs that overlap those described earlier. In addition, we have documented the phenomena with an alternative methodology, as follows. Forty University of Michigan undergraduates were given 24 arguments in individually ran-

domized order. They were asked to estimate the probability of each conclusion on the assumption that the respective premises were true. Twenty-two of the arguments corresponded to the 11 contrasts of Phenomena 1–11. Illustrating with nonmonotonicity–general (Phenomenon 9), 2 of the arguments were as follows:

Terriers secrete uric acid crystals.
_____
All canines secrete uric acid crystals. (18a)

Terriers secrete uric acid crystals.
Mustangs secrete uric acid crystals.
_____
All canines secrete uric acid crystals. (18b)

Twenty-four of the 40 subjects assigned a higher conditional probability to Argument 18a than to Argument 18b; 10 showed the reverse judgment, and 6 assigned the same conditional probability to each conclusion. This bias in favor of nonmonotonicity is significant by a sign test ($p < .01$, one-tailed). The other phenomena tested in the original experiments have been similarly replicated except for premise–conclusion asymmetry, which was tested without special instructions and yielded only a nonsignificant difference in the predicted direction.

Theory

Two Variables in Confirmation

The theory developed below claims that confirmation varies directly with the following two variables: a = the degree to

which the premise categories resemble the conclusion category; and b = the degree to which the premise categories resemble members of the lowest-level category that includes both the premise and conclusion categories. These variables can be illustrated with Argument 9b, in which the premise categories are ROBIN and BLUEJAY and in which the conclusion category is GOOSE. Variable a corresponds to the similarity between robins and bluejays on the one hand, and geese on the other. Regarding variable b, observe that BIRD is the lowest-level category that includes ROBIN, BLUEJAY, and GOOSE. Hence, b corresponds to the similarity between robins and bluejays on the one hand, and all birds on the other. This variable is intended to capture the following kind of reasoning:

> "Since robins and bluejays have the property, it may be the case that all birds have the property. Geese are birds. So maybe geese have the property too."

Although we do not claim that such reasoning is consciously produced by the typical subject, we do claim that it represents a thought process that is central to inductive judgment (see also Carey, 1985).

Rips (1975) has already proposed that variables a and b are fundamental to category-based induction. Our goal is to formulate this idea in a way that applies to a broader set of arguments than the single-premise, specific arguments considered by Rips. We show that the resulting model is consistent with all 13 phenomena discussed above and provides a reasonable fit of quantitative data described later.

Extended Similarity Functions

Both variables a and b invoke similarity as the underlying mechanism of confirmation judgment. Accordingly, our model rests on an extended notion of similarity. For a given subject S, we suppose the existence of a function SIM_s defined on any pair of elements that are at the same hierarchical level within some natural category. Pairs of this kind include (BEE, MOSQUITO), (APPLE, WATERMELON), and (FALCON, CHIMPANZEE)—the last pair consists of elements at the same hierarchical level within the category ANIMAL. Given such a pair (k, g), $SIM_s(k; g)$ is assumed to return a real number between 0 and 1 that reflects the similarity that S perceives between k and g, where values near 0 and 1 represent low and high similarity, respectively. Although $SIM_s(k; g)$ need not always equal $SIM_s(g; k)$ (see Tversky, 1977), symmetry does seem to be approximately true for the stimuli that figure in our discussion.

The SIM function is directly relevant only to single-premise, specific arguments P/C. For such arguments $SIM_s(CAT(P); CAT(C))$ represents the similarity between the categories in P and C. To be relevant to multiple-premise arguments, both general and specific, we extend the domain of the SIM function as follows. Let $k_1 \ldots k_n$, g be elements that are at the same hierarchical level within some natural category. We define $SIM_s(k_1 \ldots k_n; g)$ to be the maximum of $\{SIM_s(k_1; g) \ldots SIM_s(k_n; g)\}$. In words, $SIM_s(k_1 \ldots k_n; g)$ is the greatest similarity that S perceives between g and some one of $k_1 \ldots k_n$. To illustrate, if S's intuitions conform to ours, SIM_s(robin, crow; sparrow) = SIM_s(robin; sparrow). When $n = 1$, the extended SIM function reduces to the original one.[2]

The foregoing use of the MAX function can be motivated by the following considerations about similarity and confirmation. Consider the argument

$$\frac{\text{Rhinos have BCC in their blood.}}{\text{Antelopes have BCC in their blood.}} \qquad (19)$$

Although not exceptionally strong Argument 19 has nonnegligible strength, partly because of the similarity of RHINO to ANTELOPE. The same remarks apply to

$$\frac{\text{Elephants have BCC in their blood.}}{\text{Antelopes have BCC in their blood.}} \qquad (20)$$

However, combining the two premises, as in

$$\frac{\begin{array}{l}\text{Rhinos have BCC in their blood.}\\ \text{Elephants have BCC in their blood.}\end{array}}{\text{Antelopes have BCC in their blood.}} \qquad (21)$$

yields an argument that seems to be only slightly stronger than Argument 19 or 20, not twice as strong. The similarity of premise categories to conclusion categories thus appears not to summate when overall confirmation is mentally computed. This lack of additivity cannot be due to a mechanism that averages the similarity of premises to conclusion because averaging is inconsistent with premise monotonicity for specific arguments (see Phenomenon 7). In particular, since elephants resemble antelopes more than do monkeys, an averaging mechanism would render the single-premise Argument 20 stronger than

$$\frac{\begin{array}{l}\text{Elephants have BCC in their blood.}\\ \text{Monkeys have BCC in their blood.}\end{array}}{\text{Antelopes have BCC in their blood.}} \qquad (22)$$

But Argument 22 is clearly stronger than is Argument 20.

These considerations point to a maximizing principle in computing the similarity of multiple premises to a specific conclusion. With regard to Arguments 19–21, such a principle allows the similarity of RHINO to ANTELOPE to be eclipsed by the greater similarity of ELEPHANT to ANTELOPE (or vice versa, if RHINO is more similar to ANTELOPE than is ELEPHANT), which explains why Argument 21 is not twice as strong as Argument 19.

A theory based on maximization must also account for the greater strength of Argument 22 compared with 20. Our theory achieves this, not by considering the additional similarity of MONKEY to ANTELOPE, but by considering the greater "coverage" of the category MAMMAL by the set {ELEPHANT, MONKEY} than by {ELEPHANT}. In the same way, the somewhat greater coverage by {RHINO, ELEPHANT} compared to {ELEPHANT} of the category MAMMAL accounts for the somewhat greater strength of Argument 21 compared with 20. The need to formalize this notion of coverage leads us to a second extension of the SIM function.

We extend the SIM function so that it applies to tuples of the form $(k_1 \ldots k_n; G)$, where $k_1 \ldots k_n$ are at the same hierarchical

---

[2] Outside the context of confirmation judgment, subjects may rate the similarity of a set to an object by averaging rather than by taking MAX. The present definition of $SIM_s(k_1 \ldots k_n; g)$ is not intended to apply outside the domain of confirmation.

level of some natural category, and G is at a higher level. In this case, we define $SIM_s(k_1 \ldots k_n; G)$ to be the average of

$\{SIM_s(k_1 \ldots k_n; g) \mid S$ believes that g is at the same level as

$k_1 \ldots k_n$ and that g belongs to $G\}$.

In words, $SIM_s(k_1 \ldots k_n; G)$ is the average similarity that S perceives between $k_1 \ldots k_n$ and members of G at the level of $k_1 \ldots k_n$. To illustrate, suppose that all the songbirds that S can think of appear in the list: ROBIN, SPARROW, FINCH, CARDINAL, BLUEJAY, ORIOLE. Then, SIM_s(CARDINAL, SPARROW; SONGBIRD) is the average of

MAX{SIM_s(CARDINAL; ROBIN), SIM_s(SPARROW; ROBIN)}

MAX{SIM_s(CARDINAL; SPARROW), SIM_s(SPARROW; SPARROW)}

MAX{SIM_s(CARDINAL; FINCH), SIM_s(SPARROW; FINCH)}

MAX{SIM_s(CARDINAL; CARDINAL), SIM_s(SPARROW; CARDINAL)}

MAX{SIM_s(CARDINAL; BLUEJAY), SIM_s(SPARROW; BLUEJAY)}

MAX{SIM_s(CARDINAL; ORIOLE), SIM_s(SPARROW; ORIOLE)}

If we make the further supposition that S's intuitions are like our own, SIM_s(CARDINAL, SPARROW; SONGBIRD) equals the average of

SIM$_s$(SPARROW; ROBIN)

SIM$_s$(SPARROW; SPARROW)

SIM$_s$(SPARROW; FINCH)

SIM$_s$(CARDINAL; CARDINAL)

SIM$_s$(CARDINAL; BLUEJAY)

SIM$_s$(CARDINAL; ORIOLE).

As a second illustration, SIM_s(RABBIT; MAMMAL) equals the average of SIM_s(RABBIT; ELEPHANT), SIM_s(RABBIT; MOUSE), and so forth. SIM_s(RABBIT; PACHYDERM) does not figure in this average because PACHYDERM is not at the same level as RABBIT.[3]

It may be helpful to provide an intuitive interpretation of such expressions as SIM_s(CARDINAL, SPARROW; SONGBIRD). SIM_s returns a high value on (CARDINAL, SPARROW; SONGBIRD) to the extent that every songbird (retrieved by S) is similar to either cardinals or sparrows or both. Conversely, the value is low if there are many songbirds that are similar to neither cardinals nor sparrows. Thinking of similarity (solely as an aid to intuition) as a decreasing function of metric distance in a space of instances, SIM_s(CARDINAL, SPARROW; SONGBIRD) is large to the extent that the set {CARDINAL, SPARROW} "covers" the space of songbirds, in the sense that every songbird is near some member of {CARDINAL, SPARROW}.

Finally, it is worth pointing out that we could not properly reconstruct the coverage conception if we had earlier defined $SIM_s(k_1 \ldots k_n; g)$ to be the sum rather than the maximum of $\{SIM_s(k_1; g) \ldots SIM_s(k_n; g)\}$. To see this, consider a seven-member category $\Gamma = \{A, B, C, D, E, F, G\}$, the similarities of which are represented (inversely) by linear distance in the following diagram:

A B   C D E   F G

Intuitively, {B, D, F} covers Γ better than {C, D, E} does. This

intuition conforms to the MAX-version of $SIM_s(k_1 \ldots k_n; g)$ because every member of Γ is near some member of {B, D, F}, whereas some members of Γ—namely, A, B, F, and G—are far from every member of {C, D, E}. That is, SIM_s(B, D, F; Γ) > SIM_s(C, D, E; Γ). However, the intuition that {B, D, F} covers Γ better than does {C, D, E} is violated if $SIM_s(k_1 \ldots k_n; g)$ is computed as $SUM\{SIM_s(k_1; g) \ldots SIM_s(k_n; g)\}$. For, by measuring distances in the diagram the reader can verify that if the SUM version is used, the average of $\{SIM_s(C, D, E; \gamma) \mid \gamma \in \Gamma\}$ exceeds the average of $\{SIM_s(B, D, F; \gamma) \mid \gamma \in \Gamma\}$. Thus, the sum version of $SIM_s(k_1 \ldots k_n; g)$ counterintuitively declares {C, D, E} to provide better coverage of Γ than does {B, D, F}. The same counterintuitive result obtains if an average version of $SIM_s(k_1 \ldots k_n; g)$ is used.

The Model

Our model is formulated with the help of the following notation. Given a list $k_1 \ldots k_m$ of categories, we denote by $[k_1 \ldots k_m]$ the lowest level category K such that each of $k_1 \ldots k_m$ is a subset of K. For example:

1. [TROUT, SHARK] = FISH;
2. [RABBIT, ELEPHANT] = MAMMAL;
3. [LION, SALMON] = ANIMAL;
4. [PORCUPINE, MAMMAL] = MAMMAL;
5. [HORNET, COW, ANIMAL] = ANIMAL.

The similarity–coverage model of argument strength: For every person S there is a positive constant $\alpha \in (0, 1)$ such that for all arguments $A = P_1 \ldots P_n/C$, the strength of A for S is given by

$\alpha SIM_s(CAT(P_1) \ldots CAT(P_n); CAT(C))$

$+ (1 - \alpha) SIM_s(CAT(P_1) \ldots CAT(P_n); [CAT(P_1) \ldots CAT(P_n), CAT(C)])$.

Thus, the model allows for individual differences in the relative importance attributed to similarity and coverage in argument strength. Such differences are represented by the parameter α. In contrast, for any given subject it is assumed that a single value of α applies to all arguments evaluated in a given context.

To illustrate the model, consider the argument

> Beavers require oxydilic acid for good digestion.
> Raccoons require oxydilic acid for good digestion.
> ───
> Bears require oxydilic acid for good digestion.   (23)

According to the model, the strength of Argument 23 for a given subject S is a weighted sum of terms a and b, where

a = SIM_s(BEAVER, RACCOON; BEAR) and

b = SIM_s(BEAVER, RACCOON; [BEAVER, RACCOON, BEAR]).

Term a is the greater of SIM_s(BEAVER; BEAR) and SIM_s(RACCOON; BEAR). By the definition of the bracket notation (i.e., the lowest level category that includes the bracketed categories),

---

[3] Our present definition of $SIM_s(k_1 \ldots k_n; G)$ does not correctly apply to cases like SIM_s(EAST-MEXICAN-CHIHUAHUA; MAMMAL) because hardly any members of MAMMAL are at the same level as the very specific category EAST-MEXICAN-CHIHUAHUA. Such cases can be handled by a slight reformulation of our definitions. We do not pause for the details, however, because they are not relevant to the arguments considered in this article.

Table 2
Summary of the Similarity–Coverage Model

Theoretical concept	Explanation
$SIM_s(k; g)$, where k is at the same hierarchical level as g	Similarity according to S of k to g
$SIM_s(k_1 \cdots k_n; g)$, where $k_1 \cdots k_n$ are at the same hierarchical level as g	Maximum $\{SIM_s(k_i; g) \mid i \leq n\}$
$SIM_s(k_1 \cdots k_n; G)$, where $k_1 \cdots k_n$ are at the same hierarchical level and G is at a higher level	Average $\{SIM_s(k_1 \cdots k_n; g) \mid g \in G\}$ (= average $\{maximum\{SIM_s(k_i; g) \mid i \leq n\} \mid g \in G\}$)
$[k_1 \cdots k_n]$	Lowest level category G such that each of $k_1 \cdots k_n$ is a subset of G
$CAT(P)$, $CAT(C)$	Category terms figuring in premise or conclusion
Strength of $P_1 \cdots P_n / C = \alpha SIM_s(CAT(P_1) \cdots CAT(P_n); CAT(C))$	
$+ (1 - \alpha)SIM_s(CAT(P_1) \cdots CAT(P_n); [CAT(P_1) \cdots CAT(P_n), CAT(C)])$	

[BEAVER, RACCOON, BEAR] = MAMMAL, so Term b amounts to SIM_s(BEAVER, RACCOON; MAMMAL). This term represents the coverage of MAMMAL by BEAVER and RACCOON, that is, the average of MAX{SIM_s(BEAVER; m), SIM_s(RACCOON; m)} across all mammals m known to S.

Now consider the related argument

> Beavers require oxydilic acid for good digestion.
> Raccoons require oxydilic acid for good digestion.
> ───────────────────────────────────────
> All mammals require oxydilic acid for good digestion.   (24)

According to the model, the strength of Argument 24 for a given subject S is given by the sum of a and b, where

$$a = \alpha SIM_s(BEAVER, RACCOON; MAMMAL) \text{ and}$$

$$b = (1 - \alpha)SIM_s(BEAVER, RACCOON; [BEAVER,$$
$$RACCOON, MAMMAL]).$$

By the definition of the bracket notation, [BEAVER, RACCOON, MAMMAL] = MAMMAL, so b may be rewritten as

$$b' = (1 - \alpha)SIM_s(BEAVER, RACCOON; MAMMAL).$$

The sum of a and b′ is SIM_s(BEAVER, RACCOON; MAMMAL), regardless of the value of the parameter α. Consequently, according to the model, the strength of Argument 24 for S depends only on the coverage of MAMMAL by {BEAVER, RACCOON}.

The last example motivates the additive form of our model. It is intuitively plausible that the strength of Argument 24 depends only on the sole variable of coverage. This dependency is deduced by adding terms of the form αX and $(1 - \alpha)X$, where X is the coverage variable. We note as well that the additive combination of similarity and coverage is the simplest hypothesis for a model that invokes both variables. Support for the additive form of the model thus provides better confirmation for its underlying idea (stated in the *Two Variables in Confirmation* section above) than would support for a version of the model that relies on more complicated mechanisms.

Table 2 summarizes the concepts figuring in the similarity–coverage model.

The Phenomena Revisited

Given plausible assumptions about the SIM_s function, the similarity–coverage model predicts the 13 phenomena discussed earlier. In this sense, the phenomena provide qualitative support for the model. For each phenomenon we repeat its description, and then apply the similarity–coverage model to the contrasting arguments that were used as illustration. See Table 1 for a synopsis of relevant arguments.

Phenomena Concerning General Arguments

Phenomenon 1 (premise typicality). The more representative or typical $CAT(P_1) \ldots CAT(P_n)$ are of $CAT(C)$, the more $P_1 \ldots P_n$ confirm C. According to the model, for a given person S, the strengths of Arguments 5a and 5b are given by

$$\alpha SIM_s(ROBIN; BIRD) + (1 - \alpha)SIM_s(ROBIN; [ROBIN, BIRD])$$

and

$$\alpha SIM_s(PENGUIN; BIRD) + (1 - \alpha)SIM_s(PENGUIN; [PENGUIN, BIRD]),$$

respectively. Because [ROBIN, BIRD] = [PENGUIN, BIRD] = BIRD, the foregoing expressions reduce to SIM_s(ROBIN; BIRD) and SIM_s(PENGUIN; BIRD), respectively. SIM_s(ROBIN; BIRD) equals the average similarity of robins to other birds, whereas SIM_s (PENGUIN; BIRD) equals the average similarity of penguins to other birds. It is reasonable to suppose that for a majority of subjects, the former value is greater than the latter, which yields the greater strength of Argument 5a compared with Argument 5b. More generally, the average similarity of an instance to the members of a given category is known as the *typicality* of that instance in the given category (Smith & Medin, 1981; Tversky, 1977), and the model thus predicts greater strength for general arguments whose premises are typical rather than atypical (other factors held constant). This generalization captures Phenomenon 1.[4]

─────────────
[4] A spatial interpretation of similarity helps in understanding why central birds such as robins have greater average similarity to other birds than do peripheral birds such as penguins. Consider again the linearly arranged category Γ from the Extended Similarity Functions section.

Phenomenon 2 (premise diversity). The less similar CAT(P_1) ... CAT(P_n) are among themselves, the more P_1 ... P_n confirm C. According to the model, for a given person S, the strengths of Arguments 6a and 6b are given by

αSIM$_s$(HIPPO, HAMSTER; MAMMAL)

+ (1 − α)SIM$_s$(HIPPO, HAMSTER; [HIPPO, HAMSTER, MAMMAL])

and

αSIM$_s$(HIPPO, RHINO; MAMMAL)

+ (1 − α)SIM$_s$(HIPPO, RHINO; [HIPPO, RHINO, MAMMAL]),

respectively. Because [HIPPO, HAMSTER, MAMMAL] = [HIPPO, RHINO, MAMMAL] = MAMMAL, these expressions reduce to SIM$_s$(HIPPO, HAMSTER; MAMMAL) and SIM$_s$(HIPPO, RHINO; MAMMAL), respectively. To see that for most persons S, SIM$_s$-(HIPPO, HAMSTER; MAMMAL) is likely to be greater than SIM$_s$-(HIPPO, RHINO; MAMMAL), it suffices to observe that

1. For many $k \in$ MAMMAL (e.g., LION, ELEPHANT, HORSE), SIM$_s$-(HIPPO, HAMSTER; k) $\approx$ SIM$_s$(HIPPO, RHINO; k) because most everything that resembles rhinoceroses resembles hippopotamuses as well. (The use of the MAX interpretation of SIM$_s$ is crucial here.)

2. For no $k \in$ MAMMAL does SIM$_s$(HIPPO, RHINO; k) exceed SIM$_s$-(HIPPO, HAMSTER; k) by much because no mammal resembles rhinoceroses much more than it resembles hippopotamuses; and

3. For some $k \in$ MAMMAL (e.g., MOUSE, SQUIRREL, CHIPMUNK), SIM$_s$(HIPPO, HAMSTER; k) appreciably exceeds SIM$_s$(HIPPO, RHINO; k) because these mammals resemble hamsters more than they resemble rhinoceroses.

These facts yield the greater strength of Argument 6a compared with Argument 6b, in conformity with Phenomenon 2.

Phenomenon 3 (conclusion specificity). The more specific is CAT(C), the more C is confirmed by P_1 ... P_n. According to the model, the strengths of Arguments 7a and 7b reduce to SIM$_s$-(BLUEJAY, FALCON; BIRD) and SIM$_s$(BLUEJAY, FALCON; ANIMAL), respectively. The greater homogeneity of BIRD compared to ANIMAL implies that {BLUEJAY, FALCON} covers the former better than the latter. This implies that Argument 7a is stronger than Argument 7b, in conformity with Phenomenon 3.

Phenomenon 4 (premise monotonicity). For general arguments, more-inclusive sets of premises yield more strength than less inclusive sets. Similarly to before, the model implies that the strengths of 8a and 8b boil down to SIM$_s$(HAWK, SPARROW, EAGLE; BIRD) and SIM$_s$(SPARROW, EAGLE; BIRD), respectively. Obviously, {HAWK, SPARROW, EAGLE} covers BIRD better than {SPARROW, EAGLE} does. This implies that Argument 8a is stronger than Argument 8b, as required by Phenomenon 4.

Phenomena Concerning Specific Arguments

Phenomenon 5 (premise-conclusion similarity). The more similar CAT(P_1) ... CAT(P_n) are to CAT(C), the more P_1 ... P_n

confirm C. According to the model, for a given person S, the strengths of Arguments 9a and 9b are given by

αSIM$_s$(ROBIN, BLUEJAY; SPARROW)

+ (1 − α)SIM$_s$(ROBIN, BLUEJAY; [ROBIN, BLUEJAY, SPARROW])

and

αSIM$_s$(ROBIN, BLUEJAY; GOOSE)

+ (1 − α)SIM$_s$(ROBIN, BLUEJAY; [ROBIN, BLUEJAY, GOOSE]),

respectively. Because [ROBIN, BLUEJAY, SPARROW] = [ROBIN, BLUEJAY, GOOSE] = BIRD, these expressions reduce to

αSIM$_s$(ROBIN, BLUEJAY; SPARROW)

+ (1 − α)SIM$_s$(ROBIN, BLUEJAY; BIRD)

and

αSIM$_s$(ROBIN, BLUEJAY; GOOSE)

+ (1 − α)SIM$_s$(ROBIN, BLUEJAY; BIRD),

respectively. Because the two (1 − α) terms are identical, Argument 9a is predicted to be stronger than Argument 9b if SIM$_s$-(ROBIN, BLUEJAY; SPARROW) > SIM$_s$(ROBIN, BLUEJAY; GOOSE). Surely this is the case for most subjects. Phenomenon 5 is thereby captured.[5]

Phenomenon 6 (premise diversity). The less similar CAT(P_1) ... CAT(P_n) are among themselves, the more P_1 ... P_n confirm C. According to the model, for a given person S, the strengths of Arguments 10a and 10b are given by

αSIM$_s$(LION, GIRAFFE; RABBIT)

+ (1 − α)SIM$_s$(LION, GIRAFFE; [LION, GIRAFFE, RABBIT])

and

αSIM$_s$(LION, TIGER; RABBIT)

+ (1 − α)SIM$_s$(LION, TIGER; [LION, TIGER, RABBIT]),

respectively. Because [LION, GIRAFFE, RABBIT] = [LION, TIGER, RABBIT] = MAMMAL, these expressions reduce to

αSIM$_s$(LION, GIRAFFE; RABBIT)

+ (1 − α)SIM$_s$(LION, GIRAFFE; MAMMAL)

and

αSIM$_s$(LION, TIGER; RABBIT)]

+ (1 − α)SIM$_s$(LION, TIGER; MAMMAL),

respectively. Because SIM$_s$(LION; RABBIT) is likely to be no smaller than either SIM$_s$(TIGER; RABBIT) or SIM$_s$(GIRAFFE; RABBIT), it follows (via the MAX interpretation of SIM$_s$) that SIM$_s$(LION, GIRAFFE; RABBIT) is no smaller than SIM$_s$(LION, TIGER; RABBIT) for most persons S. On the other hand, it is clear that {LION, GIRAFFE} covers MAMMAL better than {LION, TIG-

---

The reader can verify that the average distance between the peripheral member B and the rest of Γ is greater than the average distance between the central member D and the rest of Γ.

[5] For some subjects, [ROBIN, BLUEJAY, SPARROW] may equal SONGBIRD rather than BIRD. Because {ROBIN, BLUEJAY} covers SONGBIRD even better than it covers BIRD, the model predicts such subjects to prefer Argument 9a and 9b even more strongly than subjects for which [ROBIN, BLUEJAY, SPARROW] = BIRD.

ER} does. Under the assumptions of the model, these facts imply that Argument 10a is stronger than 10b, in conformity with Phenomenon 6.

Phenomenon 7 (premise monotonicity). More inclusive sets of premises yield more strength than less inclusive sets, provided that the new premise is drawn from the lowest level category that includes the old premises and conclusion. According to the model, for a given person S, the strengths of Arguments 11a and 11b are given by

$$\alpha SIM_s(FOX, PIG, WOLF; GORILLA)$$

$$+ (1 - \alpha)SIM_s(FOX, PIG, WOLF; [FOX, PIG, WOLF, GORILLA])$$

and

$$\alpha SIM_s(PIG, WOLF; GORILLA)$$

$$+ (1 - \alpha)SIM_s(PIG, WOLF; [PIG, WOLF, GORILLA]),$$

respectively. Because [FOX, PIG, WOLF, GORILLA] = [PIG, WOLF, GORILLA] = MAMMAL, these expressions reduce to

$$\alpha SIM_s(FOX, PIG, WOLF; GORILLA)$$

$$+ (1 - \alpha)SIM_s(FOX, PIG, WOLF; MAMMAL)$$

and

$$\alpha SIM_s(PIG, WOLF; GORILLA)$$

$$+ (1 - \alpha)SIM_s(PIG, WOLF; MAMMAL),$$

respectively. By the MAX interpretation of SIM_s, $SIM_s(FOX, PIG, WOLF; GORILLA)$ is at least as great as $SIM_s(PIG, WOLF; GORILLA)$. Also by MAX, {FOX, PIG, WOLF} covers MAMMAL better than {PIG, WOLF} does. Argument 11a is thereby predicted to be stronger than Argument 11b, in conformity with Phenomenon 7.

Phenomenon 8 (premise-conclusion asymmetry). Single-premise arguments are not symmetric, in the sense that P/C may not have the same strength as C/P. According to the model, for a given person S, the strengths of Arguments 12a and 12b are given by

$$\alpha SIM_s(MOUSE; BAT) + (1 - \alpha)SIM_s(MOUSE; [MOUSE, BAT])$$

and

$$\alpha SIM_s(BAT; MOUSE) + (1 - \alpha)SIM_s(BAT; [BAT, MOUSE]),$$

respectively. Because [MOUSE, BAT] = [BAT, MOUSE] = MAMMAL, these expressions reduce to

$$\alpha SIM_s(MOUSE; BAT) + (1 - \alpha)SIM_s(MOUSE; MAMMAL)$$

and

$$\alpha SIM_s(BAT; MOUSE) + (1 - \alpha)SIM_s(BAT; MAMMAL),$$

respectively. It may be assumed that $SIM_s(BAT; MOUSE)$ is roughly equal to $SIM_s(MOUSE; BAT)$. On the other hand, the average similarity of mice to other mammals is greater than that of bats to other mammals. Hence, $SIM_s(MOUSE; MAMMAL) > SIM_s(BAT; MAMMAL)$. Putting these facts together yields greater predicted strength for Argument 12a than for 12b, in line with Phenomenon 8. The foregoing derivation also reveals the following prediction of the similarity–coverage model: For a spe-

cific argument P/C to exhibit asymmetry, CAT(P) and CAT(C) must differ in typicality.

Phenomena Concerning Mixed Arguments

Phenomenon 9 (nonmonotonicity-general). Some general arguments can be made weaker by adding a premise that converts them into mixed arguments. According to the model, for a given person S, the strengths of Arguments 13a and 13b are given by

$$\alpha SIM_s(CROW, PEACOCK; BIRD)$$

$$+ (1 - \alpha)SIM_s(CROW, PEACOCK; [CROW, PEACOCK, BIRD])$$

and

$$\alpha SIM_s(CROW, PEACOCK, RABBIT; BIRD)$$

$$+ (1 - \alpha)SIM_s(CROW, PEACOCK, RABBIT;$$

$$[CROW, PEACOCK, RABBIT, BIRD]),$$

respectively. Because [CROW, PEACOCK, BIRD] = BIRD and [CROW, PEACOCK, RABBIT, BIRD] = ANIMAL, these expressions reduce to

$$\alpha SIM_s(CROW, PEACOCK; BIRD)$$

$$+ (1 - \alpha)SIM_s(CROW, PEACOCK; BIRD)$$

and

$$\alpha SIM_s(CROW, PEACOCK, RABBIT; BIRD)$$

$$+ (1 - \alpha)SIM_s(CROW, PEACOCK, RABBIT; ANIMAL),$$

respectively. Regarding α-terms, {CROW, PEACOCK} probably covers BIRD as well as {CROW, PEACOCK, RABBIT} does. Regarding $(1 - \alpha)$-terms, {CROW, PEACOCK} covers BIRD better than {CROW, PEACOCK, RABBIT} covers ANIMAL, in view of the greater variability among animals compared to the subset birds. Under the assumptions of the model, these facts imply that Argument 13a is stronger than 13b, as specified by Phenomenon 9.

Phenomenon 10 (nonmonotonicity-specific). Some specific arguments can be made weaker by adding a premise that converts them into mixed arguments. According to the model, for a given person S, the strengths of Arguments 14a and 14b are given by

$$\alpha SIM_s(FLY; BEE) + (1 - \alpha)SIM_s(FLY; [FLY, BEE])$$

and

$$\alpha SIM_s(FLY, ORANGUTAN; BEE)$$

$$+ (1 - \alpha)SIM_s(FLY, ORANGUTAN; [FLY, ORANGUTAN, BEE]),$$

respectively. Because [FLY, BEE] = INSECT and [FLY, ORANGUTAN, BEE] = ANIMAL, these expressions reduce to

$$\alpha SIM_s(FLY; BEE) + (1 - \alpha)SIM_s(FLY; INSECT)$$

and

$$\alpha SIM_s(FLY, ORANGUTAN; BEE)$$

$$+ (1 - \alpha)SIM_s(FLY, ORANGUTAN; ANIMAL),$$

respectively. Regarding α-terms, $SIM_s(FLY; BEE) > SIM_s$

(ORANGUTAN; BEE) and, consequently, $SIM_s(FLY; BEE) = SIM_s$ (FLY, ORANGUTAN; BEE). Regarding $(1 - \alpha)$-terms, {FLY} covers INSECT better than {FLY, ORANGUTAN} covers the varied category ANIMAL. The model thus implies that Argument 14a is stronger than Argument 14b, in conformity with Phenomenon 10.[6]

A Phenomenon Involving Both General and Specific Arguments

Phenomenon 11 (inclusion fallacy). A specific argument can sometimes be made stronger by increasing the generality of its conclusion. The model implies that for a given person S, the strengths of Arguments 15a and 15b are given by

$$\alpha SIM_s(ROBIN; BIRD) + (1 - \alpha)SIM_s(ROBIN; [ROBIN, BIRD])$$

and

$$\alpha SIM_s(ROBIN; OSTRICH) + (1 - \alpha)SIM_s(ROBIN; [ROBIN, OSTRICH]),$$

respectively. Because [ROBIN, BIRD] = [ROBIN, OSTRICH] = BIRD, these expressions reduce to

$$\alpha SIM_s(ROBIN; BIRD) + (1 - \alpha)SIM_s(ROBIN; BIRD)$$

and

$$\alpha SIM_s(ROBIN; OSTRICH) + (1 - \alpha)SIM_s(ROBIN; BIRD),$$

respectively. The $1 - \alpha$ terms are identical. Regarding the α-terms, $SIM_s(ROBIN; BIRD)$ represents the average similarity of robins to other birds, including songbirds like sparrows, cardinals, and orioles. Because this average is partially weighted by the similar songbirds, $SIM_s(ROBIN; BIRD)$ exceeds $SIM_s(ROBIN; OSTRICH)$, since ostriches are highly dissimilar to robins. Phenomenon 11 follows. The foregoing derivation also reveals the following prediction of the similarity–coverage model: Arguments P/C and P/C'—with $CAT(C') \in CAT(C)$—can give rise to the inclusion-fallacy phenomenon only if $CAT(C')$ is an atypical member of $CAT(C)$.

Two Limiting-Case Phenomena

Phenomenon 12 (premise–conclusion identity). Any argument of the form Q/Q is perfectly strong. According to the model, for a given person S, the strength of Argument 16 is given by

$$\alpha SIM_s(PELICAN; PELICAN)$$

$$+ (1 - \alpha)SIM_s(PELICAN; [PELICAN, PELICAN]).$$

Because [PELICAN, PELICAN] = PELICAN, this expression reduces to $SIM_s(PELICAN; PELICAN)$. It is safe to assume that subjects perceive the similarity of pelicans to themselves to be extremely high, thereby accounting for the extreme strength of Argument 16. If we assume that $SIM_s(PELICAN; PELICAN)$ is in fact the maximal value 1, then Argument 16 is predicted to be perfectly strong.

Phenomenon 13 (premise–conclusion inclusion). Suppose that statements P and C are such that the conclusion category is included in the premise category. Then the argument P/C is perfectly strong. We must explain why Argument 17 is at least as strong as any other argument. Let $k_1 \ldots k_n$ be all the animals

known to S. Then, according to the model, $SIM_s(ANIMAL; BIRD)$ equals the average of

$$\{SIM_s(k_1 \ldots k_n; g) \mid S \text{ believes that } g \text{ is a bird}\}$$

Because birds are animals, this expression is the average of terms of the form $SIM_s(x; x)$. Such an average may be assumed to equal 1, and Phenomenon 13 is explained thereby.[7]

A Related Finding

Gelman and Markman (1986) documented a pattern of inference in young children and adults that may be illustrated as follows. Subjects were told that a pictured flamingo had a right aortic arch, whereas a pictured bat had a left aortic arch. They were then shown a pictured blackbird that resembled the bat in appearance more than it did the flamingo. Subjects nonetheless attributed the flamingolike, right aortic arch to blackbirds rather than the batlike, left aortic arch. Gelman and Markman concluded that category membership rather than similarity governs these kinds of inferences in both young children and adults.

We may represent the Gelman–Markman finding in terms of the strengths of the following arguments.

Flamingos have a right aortic arch.	
Blackbirds have a right aortic arch.	(25a)

Bats have a left aortic arch.	
Blackbirds have a left aortic arch.	(25b)

According to the similarity–coverage model, for a given person S, the strengths of Arguments 25a and 25b are given by

$$\alpha SIM_s(FLAMINGO; BLACKBIRD)$$

$$+ (1 - \alpha)SIM_s(FLAMINGO; [FLAMINGO, BLACKBIRD])$$

and

$$\alpha SIM_s(BAT; BLACKBIRD) + (1 - \alpha)SIM_s(BAT; [BAT, BLACKBIRD]),$$

---

[6] It has been suggested to us that Argument 14b is weaker than Argument 14a because the former contains a pragmatic violation. Specifically, the violation is said to consist in the fact that the orangutan premise of Argument 14b appears irrelevant to the conclusion, inasmuch as orangutans and bees belong to such different categories. Against this interpretation we may report that 64 out of 100 Chilean undergraduates judged Argument 14c, below, to be stronger than Argument 14a, even though more of its premises violate the alleged pragmatic constraint.

Flies require trace amounts of magnesium for reproduction.
Orangutans require trace amounts of magnesium for reproduction.
Salmon require trace amounts of magnesium for reproduction.
Hawks require trace amounts of magnesium for reproduction.
Jellyfish require trace amounts of magnesium for reproduction.
Rattlesnakes require trace amounts of magnesium for reproduction.
Bees require trace amounts of magnesium for reproduction.

(14c)

We leave it to the reader to deduce from the Similarity–Coverage Model the greater strength of Argument 14c compared with Argument 14a.

[7] We note that this derivation rests on a questionable assumption, namely that a category such as ANIMAL can be mentally construed as a set of instances.

respectively. Because [FLAMINGO, BLACKBIRD] = BIRD and [BAT, BLACKBIRD] = ANIMAL, these expressions reduce to

αSIM$_s$(FLAMINGO; BLACKBIRD) + (1 − α)SIM$_s$(FLAMINGO; BIRD)

and

αSIM$_s$(BAT; BLACKBIRD) + (1 − α)SIM$_s$(BAT; ANIMAL]),

respectively. Regarding the α-terms, SIM$_s$(FLAMINGO; BLACK-BIRD) < SIM$_s$(BAT; BLACKBIRD). Regarding the (1 − α)-terms, because {FLAMINGO} covers BIRD better than {BAT} covers the varied category ANIMAL, SIM$_s$(FLAMINGO; BIRD) > SIM$_s$(BAT; ANIMAL). As a consequence, Argument 25a will be judged stronger than Argument 25b if (a) α is not too large, and (b) the coverage advantage of Argument 25a is not greatly outweighed by the similarity advantage of Argument 25b. We find these latter two assumptions reasonable, and thus believe that the Gelman–Markman finding is explainable in the context of the similarity–coverage model.

Quantitative Test of the Model

We performed 12 experiments designed to obtain quantitative data bearing on the similarity–coverage model. In an initial study, subjects rated the similarity of pairs of mammals in order for us to empirically estimate the SIM function underlying the model. From the approximated SIM function, predictions were derived about the relative strength of an extensive set of arguments. The predictions were then tested against ratings of argument strength provided by an independent group of subjects. The following is a condensed description of these experiments. A full report is available in Smith, Wilkie, López, and Osherson (1989).

Initial Similarity Study

Seven of the experiments were based on the category MAMMAL and the following base set of instances:

HORSE, COW, CHIMP, GORILLA, MOUSE, SQUIRREL,   (26)

DOLPHIN, SEAL, ELEPHANT, RHINO.

An initial study was performed to obtain similarity judgments for all 45 pairs of distinct mammals drawn from this set. Each pair was printed on a separate card, and 40 subjects rank ordered all 45 cards in terms of "how similar the mammals appearing on each card are" (no ties allowed). The mean rank of each pair was divided by 45 to obtain a similarity scale between 0 and 1 (1 for perfect similarity, 0 for perfect dissimilarity). In addition, each identity pair (e.g., (HORSE, HORSE)) was assigned a score of 1. Table 3 records these similarity scores.

We used these pairwise similarity scores to approximate SIM$_s$(k; g) for each subject S and each pair of instances k, g drawn from the base set shown earlier. Averaging over subjects yields a composite similarity function defined over pairs of instances. This composite function will be denoted by SIM (without subscript), and represents the similarity intuitions of the average subject.

SIM may be extended via the MAX principle discussed earlier so that SIM($k_1 \ldots k_n$; g) is defined for any choice of instances $k_1 \ldots k_n$, g. However, we have no direct estimate of

Table 3
Similarity Scores for Pairs of Mammals

Mammals		Score	Mammals		Score
HORSE	COW	.93	CHIMP	RHINO	.48
HORSE	CHIMP	.60	GORILLA	MOUSE	.37
HORSE	GORILLA	.62	GORILLA	SQUIRREL	.48
HORSE	MOUSE	.50	GORILLA	DOLPHIN	.39
HORSE	SQUIRREL	.54	GORILLA	SEAL	.34
HORSE	DOLPHIN	.33	GORILLA	ELEPHANT	.65
HORSE	SEAL	.37	GORILLA	RHINO	.65
HORSE	ELEPHANT	.80	MOUSE	SQUIRREL	.94
HORSE	RHINO	.74	MOUSE	DOLPHIN	.17
COW	CHIMP	.55	MOUSE	SEAL	.25
COW	GORILLA	.59	MOUSE	ELEPHANT	.35
COW	MOUSE	.48	MOUSE	RHINO	.36
COW	SQUIRREL	.49	SQUIRREL	DOLPHIN	.18
COW	DOLPHIN	.26	SQUIRREL	SEAL	.27
COW	SEAL	.38	SQUIRREL	ELEPHANT	.41
COW	ELEPHANT	.79	SQUIRREL	RHINO	.35
COW	RHINO	.79	DOLPHIN	SEAL	.92
CHIMP	GORILLA	.97	DOLPHIN	ELEPHANT	.29
CHIMP	MOUSE	.51	DOLPHIN	RHINO	.26
CHIMP	SQUIRREL	.56	SEAL	ELEPHANT	.36
CHIMP	DOLPHIN	.50	SEAL	RHINO	.32
CHIMP	SEAL	.45	ELEPHANT	RHINO	.92
CHIMP	ELEPHANT	.53			

SIM($k_1 \ldots k_n$; G), where G is the given, natural kind category and $k_1 \ldots k_n$ are instances of G drawn from the base set. This is because calculation of SIM($k_1 \ldots k_n$; G) presupposes the value of SIM($k_1 \ldots k_n$; g) for all $g \in G$ that are retrieved by S, and not all of these g figure among the base set of instances (e.g., MOOSE is an instance of MAMMAL retrievable by most subjects, but does not figure in our base set for MAMMAL). An approximation to SIM($k_1 \ldots k_n$; G) is therefore necessary. For this purpose, we have replaced G by its base set of instances. For example, to compute SIM(SQUIRREL, HORSE; MAMMAL), we computed SIM(SQUIRREL, HORSE; G'), where G' = {HORSE, COW, CHIMPANZEE, GORILLA, MOUSE, SQUIRREL, DOLPHIN, SEAL, ELEPHANT, RHINO}. This approximation is crude, but it represents in straightforward fashion the larger set of computations entailed by the model.

Finally, we consider the exact form of the predictions to be tested in the experiments. For a given argument A, the predictor variable of the model has the form $\alpha X_A + (1 − \alpha)Y_A$, where X_A is the model's similarity term for A and Y_A is its coverage term for A. Both of these terms are empirically estimated from ratings of similarity. Likewise, the predicted variable is estimated from ratings of argument strength. The two rating procedures cannot, however, be relied on to provide identical scales for the two types of judgment. As a result, we take the model to be supported by any observed linear relation between predictor and predicted variables; that is, we test the prediction that for some choice of the parameter α and constants c, d, and for all arguments A, the empirically determined strength of A equals $c[\alpha X_A + (1 − \alpha)Y_A] + d$. This latter predictor has the form $aX_A + bY_A + d$, so the model predicts a high, multiple correlation between (a) the empirically obtained estimates of argument strength, and (b) approximations to the similarity and coverage variables figuring in the model. Since the Similarity–

Table 4
Confirmation Scores for Three-Premise, General Arguments

Mammals			Score	Mammals			Score
HORSE	COW	MOUSE	.33	COW	SEAL	ELEPHANT	.47
HORSE	COW	SEAL	.39	COW	ELEPHANT	RHINO	.14
HORSE	COW	RHINO	.17	CHIMP	GORILLA	SQUIRREL	.30
HORSE	CHIMP	SQUIRREL	.55	CHIMP	GORILLA	DOLPHIN	.31
HORSE	CHIMP	SEAL	.75	CHIMP	GORILLA	SEAL	.30
HORSE	GORILLA	SQUIRREL	.64	CHIMP	SQUIRREL	DOLPHIN	.80
HORSE	GORILLA	DOLPHIN	.73	CHIMP	SQUIRREL	ELEPHANT	.62
HORSE	MOUSE	SQUIRREL	.28	CHIMP	SQUIRREL	RHINO	.61
HORSE	MOUSE	SEAL	.69	CHIMP	DOLPHIN	ELEPHANT	.72
HORSE	MOUSE	RHINO	.42	GORILLA	MOUSE	SEAL	.82
HORSE	SQUIRREL	SEAL	.63	GORILLA	MOUSE	ELEPHANT	.58
HORSE	SQUIRREL	ELEPHANT	.47	GORILLA	SQUIRREL	DOLPHIN	.80
HORSE	DOLPHIN	SEAL	.27	GORILLA	SEAL	ELEPHANT	.60
HORSE	DOLPHIN	ELEPHANT	.49	GORILLA	ELEPHANT	RHINO	.26
COW	CHIMP	DOLPHIN	.76	MOUSE	SQUIRREL	SEAL	.35
COW	CHIMP	SEAL	.70	MOUSE	DOLPHIN	SEAL	.32
COW	CHIMP	ELEPHANT	.40	MOUSE	SEAL	ELEPHANT	.70
COW	MOUSE	SEAL	.68	MOUSE	SEAL	RHINO	.65
COW	MOUSE	RHINO	.40	MOUSE	ELEPHANT	RHINO	.31
COW	SQUIRREL	DOLPHIN	.76	SQUIRREL	DOLPHIN	SEAL	.30
COW	SQUIRREL	RHINO	.36	SQUIRREL	DOLPHIN	RHINO	.68
COW	DOLPHIN	ELEPHANT	.48	SQUIRREL	SEAL	RHINO	.62
COW	DOLPHIN	RHINO	.49				

Coverage Model makes no claims about the average value of the parameter a in the sample of subjects participating in our studies, we leave the a, b, d coefficients as free parameters.

Confirmation Studies

Separate groups of 20 subjects ranked the strength of arguments based on the instances in the base set. For example, one group ranked 45 arguments of the form

X requires biotin for hemoglobin synthesis.
Y requires biotin for hemoglobin synthesis.
Z requires biotin for hemoglobin synthesis.
All mammals require biotin for hemoglobin synthesis.

where X, Y, and Z are distinct mammals drawn from the base set, and different arguments contain distinct trios of mammals in their premises. Together, there are 120 such premise-triples, and 45 were randomly generated to create the 45 arguments. These premise-triples are presented in Table 4.

Four sets of 45 cards were prepared, corresponding to the 45 arguments generated for the experiment. The names of the three mammals figuring in the premises were printed near the top of each card. The four sets differed in the order in which the mammals on a card appeared; four different random patterns were used. The following instructions were used:

We are frequently called upon to make judgments of the likelihood of something being true on the basis of limited information. Consider the following statement:

All mammals require biotin for hemoglobin synthesis.

How likely would you think that this statement is true if you knew, say, that all coyotes required biotin for hemoglobin synthesis? Would your opinion change if, instead of coyotes, you knew the statement to be true of moles, or anteaters?

In this task you will be helping us to find out more about this type of reasoning. You will be handed a set of 45 cards. On each card will be written the name of the three mammals. For each card, you are to accept it as given that the mammals listed require biotin for hemoglobin synthesis. On the basis of this evidence, you are to determine how likely it is that all mammals require biotin for hemoglobin synthesis. Each card is to be evaluated entirely independently of the others.

Some of the mammals may seem to provide stronger evidence than others. Your task is to arrange the 45 cards in order of increasing strength of evidence.

The mechanics of a ranking procedure were then explained, and it was made explicit that no ties in the ranking were permitted.

The ranks assigned by the 20 subjects were averaged and divided by 45. Each argument thus received an "obtained confirmation score," namely, a number between 0 and 1, where 1 represents *high* assessed confirmation and 0 represents *low* assessed confirmation. These obtained confirmation scores are presented in Table 4.

Consider now the predicted confirmation scores. According to the similarity–coverage model, the strength of each of the arguments is given by

$$a\text{SIM}(X, Y, Z; \text{MAMMAL})$$
$$+ (1 - a)\text{SIM}(X, Y, Z; [X, Y, Z, \text{MAMMAL}]).$$

Because X, Y, and Z are mammals, $[X, Y, Z, \text{MAMMAL}] = \text{MAMMAL}$, so the foregoing expression reduces to $\text{SIM}(X, Y, Z; \text{MAMMAL})$. For each triple X, Y, Z of mammals figuring in the experiment, an approximation to $\text{SIM}(X, Y, Z; \text{MAMMAL})$ was computed by first determining the maximum similarity of each mammal in the base set to X, Y, Z, and then taking the average of these maximum similarities. The correlation between pre-

Table 5
Confirmation Scores for Two-Premise Specific Arguments (Horse, Experiment 4)

Mammals		Score	Mammals		Score
COW	CHIMP	.79	GORILLA	SEAL	.41
COW	GORILLA	.75	GORILLA	ELEPHANT	.61
COW	MOUSE	.74	GORILLA	RHINO	.63
COW	SQUIRREL	.72	MOUSE	SQUIRREL	.17
COW	DOLPHIN	.73	MOUSE	DOLPHIN	.28
COW	SEAL	.73	MOUSE	SEAL	.25
COW	ELEPHANT	.75	MOUSE	ELEPHANT	.58
COW	RHINO	.77	MOUSE	RHINO	.62
CHIMP	GORILLA	.23	SQUIRREL	DOLPHIN	.32
CHIMP	MOUSE	.42	SQUIRREL	SEAL	.26
CHIMP	SQUIRREL	.40	SQUIRREL	ELEPHANT	.54
CHIMP	DOLPHIN	.40	SQUIRREL	RHINO	.61
CHIMP	SEAL	.43	DOLPHIN	SEAL	.06
CHIMP	ELEPHANT	.59	DOLPHIN	ELEPHANT	.54
CHIMP	RHINO	.64	DOLPHIN	RHINO	.54
GORILLA	MOUSE	.48	SEAL	ELEPHANT	.51
GORILLA	SQUIRREL	.47	SEAL	RHINO	.56
GORILLA	DOLPHIN	.38	ELEPHANT	RHINO	.57

dicted confirmation scores and obtained confirmation scores is .87 ($N = 45, p < .01$).

A replication of the previous study was performed with new subjects using all 45 arguments based on 2 distinct mammals from the base set. The resulting correlation between obtained and predicted confirmation scores was .63 ($N = 45, p < .01$). Another replication used all one-premise arguments derived from the base set and gave a correlation of .75 ($N = 10, p < .01$).

The foregoing experiments provide evidence for the predictive value of the coverage variable of the similarity–coverage model. To evaluate the role of the similarity variable, a second series of studies was performed with specific conclusions. For example, 20 new subjects rated all 36 possible arguments of the form

> X requires biotin for hemoglobin synthesis.
> Y requires biotin for hemoglobin synthesis.
> Horses require biotin for hemoglobin synthesis.

where X and Y are distinct mammals drawn from the base set, neither of them HORSE, and different arguments contain distinct pairs of mammals in their premises. As before, the ranks assigned by the 20 subjects to the 36 arguments were averaged and divided by 36. Table 5 presents these mean ranks.

According to the Similarity–Coverage Model, the strength of each argument is given by

$$\alpha \text{SIM}(X, Y; \text{HORSE}) + (1 - \alpha)\text{SIM}(X, Y; [X, Y, \text{HORSE}]).$$

Because X and Y are mammals, $[X, Y, \text{HORSE}] = \text{MAMMAL}$, so the foregoing expression reduces to

$$\alpha \text{SIM}(X, Y; \text{HORSE}) + (1 - \alpha)\text{SIM}(X, Y; \text{MAMMAL}).$$

For each pair X, Y of mammals figuring in the experiment, the value of SIM(X, Y; HORSE) was taken directly from the data of the initial similarity study (using the MAX interpretation of SIM(X, Y; HORSE)). Regarding the second term, an approximation to SIM(X, Y; MAMMAL) was computed as described above.

The Similarity–Coverage Model implies that these two predictor variables should predict the obtained confirmation scores up to linearity. In fact, the multiple correlation coefficient between the latter two variables and the obtained confirmation scores is .96 ($N = 45, p < .01$).

Can the data be used to provide evidence for both similarity and coverage variables in the strength of specific arguments, that is, is there evidence for both SIM(X, Y; HORSE) and SIM(X, Y; MAMMAL) in an argument of form X, Y/HORSE? A natural way to test for the effect of these variables would be to compute the partial correlation between each predictor variable and the obtained confirmation score with the effects of the other predictor variable partialled out. Unfortunately, the interpretation of such an analysis is clouded by the fact that the similarity and coverage variables rely on overlapping facts about SIM. In particular, a high value of SIM(X, Y; HORSE) increases the value of SIM(X, Y; MAMMAL). For this reason, instead of partial correlations, we have computed nonpartial, Pearson coefficients between obtained confirmation and each predictor variable taken alone. The correlation between obtained confirmation scores and the similarity variables SIM(X, Y; HORSE) is .95 ($N = 45, p < .01$). The correlation between obtained confirmation scores and the coverage variable SIM(X, Y; MAMMAL) is .67 ($N = 45, p < .01$). These two coefficients are significantly different ($p < .01$).

The foregoing results suggest that maximum similarity to HORSE is sufficient to account for the obtained confirmation scores. This fact should not be taken to support the view that the strength of specific arguments depends only on the similarity of the premise categories to the conclusion category. Such a hypothesis is contradicted by qualitative phenomena discussed earlier (e.g., premise diversity for specific arguments, see Phenomenon 6).

The foregoing study was replicated three times using different mammals for the conclusion category, and different numbers of premises. The obtained correlations between predicted and observed confirmation scores were all .94 or better.

Other Replications

As a check on the robustness of the preceding findings, five additional studies were performed. Each study involved one or more of the following changes compared to the seven original studies. First, instead of ranking arguments, subjects rated the probability of an argument's conclusion assuming the truth of its premises; in addition, different blank properties were used for every argument. Second, subjects were native French or Spanish speakers, working with translated materials. Third, the category INSECT was used in place of MAMMAL. All correlations in these studies between predicted and observed confirmation scores were significant at the .01 level, with a median correlation of .88. See Smith et al. (1989) for details.

Discussion

The conjunction of qualitative and quantitative evidence discussed in previous sections provides reason to believe that the two terms of the similarity–coverage model reflect genuine psychological processes that are central to confirmation. The

model nonetheless remains underdetermined by the data considered in this article. In this section, we take up several proposals for theoretical refinement or amendment.

Weighting of Instances by Availability

Different members of the same category are often differentially available to a person S, even if S can recognize all of them as members of the category. For example, robins may be more accessible than turtledoves in S's memory as members of the category BIRD. To represent such differential availability, the present version of the similarity–coverage model requires refinement. One way to achieve this is to incorporate relative availability in the computation of coverage. Specifically, for members $k_1 \ldots k_n$ of category G, SIM($k_1 \ldots k_n$; G) can be redefined so that the maximum similarity of $k_1 \ldots k_n$ to $g \in G$ is weighted by the availability of g to S.

Although weighting by availability does not affect the model's ability to deduce the 13 qualitative phenomena, it could conceivably improve predictive accuracy in the experiments reported. Accordingly, we examined several principled bases—all derived from rated typicality—for assigning relative availability to the mammals figuring in the experiments. None of these revisions of the model resulted in better overall predictive accuracy. This lack of improvement is probably due to low variability in availability among the mammals used; all were highly typical, and no doubt became even more available by virtue of their continued use in a given experiment.

MAX Versus SUM in SIM($k_1 \ldots k_n$; g)

Consider a specific argument P_1, P_2/C. In computing the overall similarity of CAT(P_1), CAT(P_2) to CAT(C), the similarity–coverage model employs a MAX function over the similarities of CAT(P_1) to CAT(C) and CAT(P_2) to CAT(C). Use of maximization was motivated by the observation that the strength of P_1, P_2/C seems not to be the sum of the strengths of P_1/C and P_2/C.

MAX is an extreme example of a nonadditive function, and it is possible that subjects use a function somewhere between MAX and SUM. Indeed, a more elaborate form of the similarity–coverage model might be equipped with another parameter that reflects an individual subject's position in the MAX to SUM continuum. Given this new parameter value β, the model would define SIM$_s$($k_1 \ldots k_n$; g) to be

$$\beta \, \text{MAX}\{\text{SIM}_s(k_1, g) \ldots \text{SIM}_s(k_n, g)\}$$
$$+ (1 - \beta)\text{SUM}\{\text{SIM}_s(k_1, g) \ldots \text{SIM}_s(k_n, g)\}.$$

This parameterized SIM function could then be incorporated into the similarity–coverage model as before, the resulting model having two parameters instead of one.

Diversity Versus Coverage

Philosophers of science underscore the usefulness of diversified data in testing scientific theories. Intuitively, there are fewer plausible alternatives to a theory that predicts phenomena of different sorts compared to one whose predictions are always of the same kind. (For discussion, see Horwich, 1982.) Given

subject S and argument $P_1 \ldots P_n/C$, the closest variable to diversity in the present model is the coverage by $\{\text{CAT}(P_1) \ldots \text{CAT}(P_n)\}$ of the lowest level category that includes the premise and conclusion categories.

What is the relation between diversity in the philosopher's sense, and coverage in the present sense? To answer this question, it is necessary to assign a precise meaning to the diversity concept. Given subject S and set K of instances, we define DIV$_s$(K)—the diversity of K (for S)—to be

$$\text{SUM}\{1 - \text{SIM}_s(k_1; k_2) | k_1, k_2 \in K\},$$

that is, the sum of the dissimilarities between members of K. Given sets K and G, we define the "diversity of K compared to G (for S)" to be DIV$_s$(K) divided by DIV$_s$(G). This latter quotient is denoted by DIV$_s$(K; G). Observe that if either K or G have less than two members, then DIV$_s$(K; G) is not defined.

Now let a subject S, a category G, and two instances k_1, k_2 be given. Then, SIM$_s$(k_1, k_2; G) is the coverage of $\{k_1, k_2\}$ in G, and DIV$_s$($\{k_1, k_2\}$; G) is the diversity of $\{k_1, k_2\}$ compared with G. These two terms may differ considerably. For example, if k_1, k_2 represent highly dissimilar but very eccentric instances of G (e.g., whales and bats in the category MAMMAL), then DIV$_s$($\{k_1, k_2\}$; G) may be comparatively high but SIM$_s$(k_1, k_2; G) may be comparatively low. We may use the similarity data of Table 3 to empirically contrast coverage and diversity. Taking G to be the base set of mammals, we calculated SIM(k_1, k_2; G) and DIV($\{k_1, k_2\}$; G) over all 45 pairs k_1, k_2 of instances drawn from this set. The correlation of these two variables is only .55.

In terms of the DIV function, the philosopher's intuition about diversity may be stated as the following theory about general arguments involving more than one premise.

The diversity model for general arguments: For every person S and every general argument $A = P_1 \ldots P_n/C$ in which $n \geq 2$, the strength of A for S is given by

$$\text{DIV}_s(\{\text{CAT}(P_1) \ldots \text{CAT}(P_n)\}; \text{CAT}(C))$$

The reader can verify that the diversity model is compatible with those qualitative phenomena that bear on general arguments having more than one premise, namely, premise diversity, conclusion specificity, and premise monotonicity (Phenomena 2–4). We do not believe, however, that the diversity model is an accurate portrayal of the strength of general arguments. First, the model provides no account of single-premise, general arguments (because individual premises manifest no diversity). As a consequence, incorporating diversity into a more complete theory of category-based induction requires positing separate psychological mechanisms for single- versus multiple-premise arguments. Special provision would also be necessary for mixed arguments, for in such arguments confirmation can decrease rather than increase as premises become more dissimilar (cf. Phenomenon 9, nonmonotonicity–general). In contrast, the similarity–coverage model relies on a single (albeit extended) similarity function, and it applies uniformly to arguments of any number of premises, be they specific, mixed, or general. It may turn out that multiple, independent mechanisms underlie human inductive judgment, even for the restricted class of arguments at issue in this article. But complex models should not be favored over simple ones until required by recalcitrant data.

Also, there is a datum that seems to favor the coverage approach to general arguments over the diversity approach. There are general arguments P_1, P_2/C and Q_1, Q_2/C such that the diversity of $\{Q_1, Q_2\}$ exceeds that of $\{P_1, P_2\}$, but P_1, P_2/C is stronger than Q_1, Q_2/C. To provide an example of such a pair of arguments, we note that {PELICAN, ALBATROSS} is at least as diverse as {ROBIN, SPARROW}. The same experimental procedure used to verify the 11 qualitative phenomena also yields the following contrast.

Robins have a choroid membrane in their eyes.
Sparrows have a choroid membrane in their eyes.
_____   (27a)[28]
All birds have a choroid membrane in their eyes.

Pelicans have a choroid membrane in their eyes.
Albatrosses have a choroid membrane in their eyes.
_____   (27b)[12]
All birds have a choroid membrane in their eyes.

The diversity model is incompatible with this result. In contrast, the similarity–coverage model provides the following explanation for it. According to the model, for a given person S, the strengths of Arguments 27a and 27b are given by:

αSIM$_s$(ROBIN, SPARROW; BIRD)

$+ (1 - \alpha)$SIM$_s$(ROBIN, SPARROW; [ROBIN, SPARROW, BIRD])

and

αSIM$_s$(PELICAN, ALBATROSS; BIRD)

$+ (1 - \alpha)$SIM$_s$(PELICAN, ALBATROSS; [PELICAN, ALBATROSS, BIRD]),

respectively. Because [ROBIN, SPARROW, BIRD] = [PELICAN, ALBATROSS, BIRD] = BIRD, these expressions reduce to SIM$_s$(ROBIN, SPARROW; BIRD) and SIM$_s$(PELICAN, ALBATROSS; BIRD), respectively. These latter terms represent the coverage by {ROBIN, SPARROW} of BIRD, and the coverage by {PELICAN, ALBATROSS} of BIRD, respectively. SIM$_s$(ROBIN, SPARROW; BIRD) equals the average of MAX{SIM$_s$(ROBIN; b), SIM$_s$(SPARROW; b)} over all birds b known to S. SIM$_s$(PELICAN, ALBATROSS; BIRD) equals the average of MAX{SIM$_s$(PELICAN; b), SIM$_s$(ALBATROSS; b)} over all such birds b. It is obvious that for most subjects S, the former average is greater than the latter, because most birds known to S are small, sing, and so forth. The greater strength of Arguments 27a compared to 27b is thereby deduced.

The Multiplicity of Categories in Confirmation Judgment

The similarity–coverage model assumes the existence of a preestablished hierarchy of categories that classify the instances figuring in an argument. The success of the model in predicting

the qualitative phenomena discussed earlier testifies to the approximate soundness of the model's assumption. Greater predictive accuracy nonetheless requires supplementary principles to describe the variety of categories that subjects may create "on line" when reasoning about argument strength (cf. Barsalou, 1983; Kahneman & Miller, 1986). Thus, the following argument may give rise to the covering category SMALL ANIMAL in the minds of many subjects, despite the absence of this category from their prestored list of animal classes.

Hummingbirds require Vitamin L for carbohydrate breakdown.
Minnows require Vitamin L for carbohydrate breakdown.
_____
Titmice require Vitamin L for carbohydrate breakdown.

The mental origin of such categories in reasoning about argument strength remains a central problem in the study of confirmation.

References

Barsalou, L. (1983). Ad hoc categories. *Memory & Cognition, 11,* 211–227.

Carey, S. (1985). *Conceptual change in childhood.* Cambridge, MA: MIT Press.

Collins, A., & Michalski, R. (1989). The logic of plausible reasoning: A core theory. *Cognitive Science, 13,* 1–50.

Gelman, S. (1988). The development of induction within natural kind and artefact categories. *Cognitive Psychology, 20,* 65–95.

Gelman, S., & Markman, E. (1986). Categories and induction in young children. *Cognition, 23,* 183–209.

Horwich, P. (1982). *Probability and evidence.* Cambridge, England: Cambridge University Press.

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

Osherson, D., Smith, E., & Shafir, E. (1986). Some origins of belief. *Cognition, 24,* 197–224.

Rips, L. (1975). Inductive judgments about natural categories. *Journal of Verbal Learning and Verbal Behavior, 14,* 665–681.

Rothbart, S., & Lewis, P. (1988). Inferring category attributes from exemplar attributes. *Journal of Personality and Social Psychology, 55,* 861–872.

Shafir, E., Smith, E. E., & Osherson, D. (in press). Typicality and reasoning fallacies. *Memory & Cognition.*

Smith, E. E., & Medin, D. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Smith, E., Wilke, A., López, A., & Osherson, D. (1989). *Test of a confirmation model* (Occasional Paper). Cambridge: Massachusetts Institute of Technology.

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