

- Sweller, J., Mawer, R.F., & Ward, R.W. (1983). Development of expertise in mathematical problem solving. *Journal of Experimental Psychology: General*, *112*, 639-661.
- Unruh, A., Rosenbloom, P.S., & Laird, J.E. (1987). Dynamic abstraction problem solving in Soar. In *Proceedings of the AOG/AAAI Joint Conference*. Dayton, OH.
- Voss, J., Vesonder, G., & Spillich, G. (1980). Text generation and recall by high-knowledge and low-knowledge individuals. *Journal of Verbal Learning and Verbal Behavior*, *19*, 651-667.
- Wason, P.C. (1966). Reasoning. In B.M. Foss (Ed.), *New horizons in psychology*. Harmondsworth, England: Penguin.

Combining Social Concepts: The Role of Causal Reasoning

ZIVA KUNDA, DALE T. MILLER, AND THERESA CLAIRE

Princeton University

Four studies examined how people combine social concepts that have conflicting implications (e.g., Harvard-educated and carpenter). Several kinds of evidence indicated that such combinations are guided by causal reasoning that draws upon both causal relations contained within the constituent concepts and on broader world knowledge. Open-ended descriptions of members of combinations contained explicit causal descriptors, as well as emergent attributes not used to describe members of constituents. Ratings of the likelihood that combination members possessed various attributes were not fully predicted by comparable ratings of constituents. Causal reasoning appeared to be most pervasive for combinations viewed as more surprising, suggesting that surprise may have triggered the generation of causal accounts.

Upon hearing that a person belongs to a social category for which there is a well-developed representation, or stereotype, it is typical to form a set of expectations about that person. It may be assumed, for example, that a Harvard-educated person will be intelligent and affluent, that a carpenter will be rugged and handy, and so on. But every person belongs to more than one social category, and these different categories will evoke different, sometimes even conflicting, expectations. Consider a person who is both Harvard-educated and a carpenter. Forming a unified impression of such a person requires the combination of not only two concepts, but ones that, for most people, have conflicting implications. The present research seeks to determine whether unified impressions are formed in such cases and, if so, how this is accomplished. Rephrased in terms more familiar to cognitive scientists, our research focuses on conceptual combination as it pertains to social categories.

We propose that the combination of social concepts, especially those with conflicting implications, involves causal reasoning and explanatory

We are grateful to Brian Reiser, Eldar Shafir, and Paul Thagard for comments on an earlier version of this article.

Correspondence and requests for reprints should be sent to Ziva Kunda, Psychology Department, Green Hall, Princeton University, Princeton, NJ 08544-1010.

hypothesis formation. When confronted with a person who belongs to social categories with conflicting implications, people are surprised or puzzled. To resolve this puzzlement, they might ask themselves questions of the form: How could a Harvard-educated person become a carpenter? That is, what might have caused a person belonging to one of the categories to acquire membership in the other? To answer such questions, people will draw upon their knowledge of the categories to be combined, as well as on their broader world knowledge to construct an explanatory causal account of the reasons for the dual category membership. For example, a Harvard-educated person may have become a carpenter because the radical climate of the sixties fostered disillusionment with the materialistic competitive world encountered at Harvard. Guided by such causal accounts, one may include in the combined and unified image of the person some of the attributes associated with each of the original categories as well as some novel or, what we will term, *emergent attributes*. The Harvard-educated carpenter may be intelligent and of an upper-middle-class background like Harvard-educated people, handy and rugged like carpenters, and may also possess some attributes such as nonmaterialistic that emerge from the causal reasoning used to resolve the puzzle of membership in both categories.

Causal relations have been shown to play a central role in a wide variety of cognitive tasks, including the comprehension of utterances (Clark & Clark, 1979), the detection of correlations (Chapman & Chapman, 1967, 1969), and the categorization of objects (Barsalou, 1983). And it has recently been argued that causal relations play an important role in the coherence of concepts (Medin, 1989; Murphy & Medin, 1985). Thus, the representation of a bird includes not only the attributes "has wings" and "can fly," but also the causal relations between these attributes: It is the wings that make flying possible. This latter idea seems especially worthy of further investigation because the way concepts are combined may provide greater insight into the structure of concepts that could be gained through examination of individual concepts themselves (cf. Murphy, 1988).

Surprisingly, research on the conceptual combination of nonsocial concepts has largely ignored the role of causal reasoning and inference. As an example, consider the elegant and explicit modification model developed by Smith and Osherson (1984; Smith, Osherson, Rips, & Keane, 1988). The modification model assumes that the combination of simple adjective-noun pairs such as red apple is achieved by restricting the range of possible values for the noun (apple) on the dimension implied by the adjective (to the color red, in this case) and boosting the importance of that dimension. This model, which does not postulate any causal reasoning, successfully accounted for typicality judgments involving simple adjective-noun pairs (colors and shapes paired with fruits and vegetables). Medin and Shoben (1988) have criticized this model, arguing that it is limited even in the simple domain to which it was applied in that it does not take into account the possibility that

some attributes are correlated with other attributes (e.g., a red apple will also be sweet). And, more important from our point of view, Medin and Shoben point out that the modification model does not recognize that attributes equally true of two concepts may be more central (and hence less mutable) for one concept than for the other. So, although being curved is equally true of bananas and boomerangs, it is more central to the essence of boomerang: A straight banana may still be a banana, but a straight boomerang may no longer be a boomerang. Such differences in centrality seem to result from the causal relations that associate attributes with concepts.

We believe that, in addition to reliance upon intraconcept causal relations, conceptual combination may also entail reliance upon causal relations that are external to the concepts in question and are derived from broader world knowledge. We know of no direct support for this proposition but there is some suggestive evidence. Rokeach and Rothman (1965) and Higgins and Rholes (1976) both showed that people's evaluations of combined concepts (e.g., immoral priest) could not be predicted entirely from their evaluations of the constituent concepts, suggesting that a simple elementaristic model (e.g., Anderson, 1968) cannot fully account for the combination of concepts. Unfortunately, neither of these studies explored in depth the nature of the causal reasoning underlying the evaluations of combined concepts.

In other relevant research Hampton (1987) found that noun-noun combinations sometimes contain novel attributes not found in either of the constituent concepts. For example, birds that are pets are believed to live in cages, a property not true for birds or for pets. One possible account of these novel attributes is that they reflect causal reasoning involving broader world knowledge (e.g., the birds might escape otherwise). Of course, it is also possible that this belief stems from people's familiarity with birds that are pets, in which case the emergent attribute comes from retrieval of stored knowledge about where pet birds live rather than from causal reasoning. Similarly, mere familiarity with the combined concepts may also account for Murphy's (1988) finding that an attribute may be viewed as more typical of a combination than of either of its constituents. Since the combined concepts studied by Hampton and by Murphy seem familiar, it is impossible to say whether the novel attributes that emerged in the combinations resulted from retrieval or from causal inference. The process of retrieval is not always a plausible account for novel attributes, however. Instead of a pet bird, imagine a hippopotamus that is a pet. If you assume, as we do, that this creature is not kept indoors, we doubt that you arrived at this assumption by retrieving existing knowledge about hippopotamuses that are pets; more likely, you did so instead by engaging in causal reasoning based on your knowledge about hippopotamuses, pets, and houses.

In the present studies we attempt to secure more direct evidence for the claim that conceptual combination involves causal reasoning that relies upon world knowledge outside the constituent concepts. We believe that the

investigation of how social concepts are combined is particularly likely to reveal evidence of causal reasoning. Social concepts are rich and complex, and therefore highly flexible and capable of stretching in order to accommodate a wide variety of modifications and contradictions. Indeed, it has been shown that people can easily explain how a person may possess two semantically opposite trait adjectives such as cheerful and gloomy (Asch & Zukier, 1984). Therefore, people should be readily able to construct explanations to account for the combination of broader social stereotypes that have conflicting implications even though they do not stand in direct semantic opposition to each other. Unlike many nonsocial nouns of the sort investigated by Hampton, (1987, 1988), social concepts typically involve few, if any, attributes considered necessary or impossible for category membership. This increases their malleability, and leaves more room for their modification through causal reasoning. More important, the nature of discourse about people seems particularly conducive to causal reasoning—questions about how people might have become what they are, and about the consequences of membership in a given social category—are encountered and dwelled upon quite frequently. It seems much more likely to wonder spontaneously how a person obtained a profession than how an object obtained its color. This complexity should be revealed when exploring how people combine social concepts. We therefore chose to examine the role of causal reasoning in the combination of social concepts.

STUDY 1

Our first study seeks to show that when people encounter a person who belongs to a surprising combination of social categories, they attempt to create a unified image of the person. Furthermore, we propose that in doing this they engage in causal reasoning, drawing inferences based upon causal relations contained within the categories in question as well as on causal relations stemming from knowledge outside these categories. We decided to focus initially on surprising combinations because it has been shown that causal reasoning is facilitated when people encounter surprising or unexpected events or objects (Hastie, 1984; Pyszczynski, & Greenberg, 1981; Wong & Weiner, 1981). We anticipated that the puzzlement aroused by these surprising combinations would be especially likely to trigger causal reasoning as subjects attempt to resolve the puzzle.

For each combination, different groups of subjects provided written descriptions of their expectations about a person who belonged to one or the other of two constituent categories or to both categories. We expected that causal reasoning would be evident in individual descriptions of members of combined categories. We also expected that it would be possible to

infer casual reasoning and reliance upon broad world knowledge from comparisons of the consensual content of combinations to that of constituents: The appearance of novel, emergent attributes in the combinations would imply that their representations were constructed through reliance upon information not contained in the constituents.

Method

Subjects were 85 undergraduates at Gettysburg College who participated voluntarily in a classroom setting.

Each subject received a booklet in which a target person was named on the top of each page. They were instructed to speculate, in as much detail as they could, about what each target person would be like. Subjects were divided into three groups: Those who described members of combined categories (e.g., a person who is Harvard-educated and a carpenter), those who described one set of constituent members (e.g., a person who is Harvard-educated) and those who described the other set of constituent members (e.g., a person who is a carpenter). Half of the subjects described one set of constituents or combinations (leftist businessman, blind marathon runner, and gay construction worker), and half of the subjects described another set (feminist bank teller, blind lawyer, Harvard-educated carpenter, and communist ex-Marine). Within each set, two fixed orders were used. The number of subjects receiving each type of phrase was 9 to 12 for constituents and 20 to 24 for combinations. For the subjects describing combinations, the order of constituents within each combination was counterbalanced. (We obtained larger numbers for the combinations to determine whether the order in which the categories were presented in a given combination would affect the content of descriptions. Since no such order effects were found, this possibility will not be discussed further.)

After completing their descriptions, subjects describing the combinations were asked to rate how surprised they would be to hear about a person who belonged to each combination, on a 7-point scale ranging from "not at all surprised" to "very surprised."

¹ In the absence of order effects in this and the following studies, our findings differ from those obtained by Hampton (1988). Hampton found an asymmetry in the combination of concepts such as "sports which are games" so that greater weight was given to the concept in the relative clause-qualifier position. The difference between the two sets of studies may have resulted from the use of somewhat different methodologies as well as from the use of different kinds of concepts. Unlike Hampton, who used the wording "A *which are* B," our wording, "a person who is A *and* B," did not indicate that one of the concepts should be used to qualify the other. Furthermore, our subjects were asked to describe or rate the attributes of a person belonging to both categories, whereas Hampton's subjects were asked to indicate the extent to which various instances belonged to both categories. Finally, our social categories may have involved different kinds of conceptual combinations than did Hampton's nonsocial categories.

Results and Discussion

Ease of Revolution. Subjects had no trouble describing the members of our incongruous social categories. Of the 156 descriptions of persons belonging to combined categories, there were only four instances in which subjects left a blank page or expressed problems with forming an impression.

Causal Narratives. Clear evidence that at least some subjects engaged in causal reasoning to form impressions of members of combined categories came from the detailed narratives provided by some subjects. For example, one subject wrote of a person who is a lawyer and blind:

...the person is obviously very hard working. I mean, people with no handicaps have a hard enough time making it through law school, let alone someone who is blind.

And another subject wrote of a person who is gay and a construction worker:

This person is most likely sublimating, he is in a position that commands masculinity, however, he himself does not fit the stereotype. Possibly he is hiding his true identity behind the image of a tough and rugged construction worker.

These examples suggest that the attributes contained in the representations of social categories are embedded in a network of causal relations that allow people to make sense of a person's membership in two seemingly incongruous categories.

The narratives tended to be structured as answers to a question of the form: "How could a person belonging to one category (e.g., blind) come to acquire membership in another category (e.g., lawyer)?" Apparently, the surprising combination triggers such questions, and people engage in causal reasoning to answer these questions.

Often, there was more than one way of answering such questions in order to resolve the apparent contradictions stemming from dual category membership. For example, one subject wrote of a Harvard-educated carpenter:

Someone who has inherited a lot of money and is working at what he enjoys, rather than where the money is. A nice person, rather than money grubbing.

Another subject wrote of the same combination:

...high class (status); earns big bucks. Carpentering is his hobby. His life is fulfilled and has nothing more to accomplish in life....

And yet another subject wrote:

Has ambitions to become a master carpenter and design new methods and objects in carpentry.

The fact that there were different ways of forming impressions of members of combined categories limits what can be discovered from examining consensual modes of resolution, across subjects. Nevertheless, interesting patterns emerge from such analyses, as reported below.

Emergent Attributes. More often than not, subjects' descriptions of members of combined categories were in the form of lists of attributes rather than in narrative form. Such attribute listing does not rule out the possibility that causal reasoning was used to combine the concepts, because subjects may have been writing only the end products of their causal reasoning without bothering to specify the causal chains that led them to these end products. One type of evidence that causal reasoning was involved in producing subjects' attribute lists would be the presence of emergent attributes; that is, attributes that were used to describe a member of a combination but were not used to describe either of its constituents. Such emergent attributes would indicate that, in forming an impression of the combination, subjects relied not only upon the information contained within each constituent, but also upon broader world knowledge.

To explore this possibility, we counted the number of times each attribute was mentioned in the descriptions of each constituent and combination. Different attributes used to express the same meaning were counted as the same. Two independent coders sorted subjects' descriptions into attributes, and disagreements were resolved by a third coder. We counted as emergent any attribute that was used by at least three subjects to describe a member of the combination, but that was not used by any subject to describe either of the constituents. Using this criterion, every one of the combinations yielded at least two emergent attributes, as shown in Table 1.

The presence of so many emergent attributes is particularly striking because our criterion for emergent attributes, namely that at least three subjects mention each, requires that several subjects resolve the conflict among constituents in the same way. Since, as described above, there were different ways of combining each pair of categories, it is reasonable to assume that the results presented in Table 1 greatly underestimate both the number of subjects relying upon knowledge outside the constituents to form impressions of the combinations, and the number of ways in which such impressions could be formed. Nevertheless, these findings do provide strong evidence that subjects relied upon world knowledge outside the constituent categories in their combination. Since the combinations seem relatively unfamiliar and unlikely to be stored as existing knowledge structures, the emergent attributes probably resulted from causal inference, and reveal the rich causal network in which the representations of social categories are embedded.

Surprise and the Generation of Causal Antecedents. To obtain a better understanding of the causal reasoning involved in creating images of members of surprising combinations, we coded each subject's descriptions for

TABLE 1
Emergent Attributes Used by at Least 3 Subjects to Describe a Combination,
but Not Used by any Subject to Describe its Constituents

Combination	Emergent Attributes
Blind & Lawyer	Determined (23) ¹ Confident (8) Good lawyer (4)
Blind & Marathon runner	Courageous (10) Confident (8) Enjoys life (4)
Communist & Ex-Marine	Troublemaker (6) Anti-American (5) Independent (3)
Feminist & Bank teller	Hypocritical (9) Anti-Male (5) Unmarried (5)
Gay & Construction worker	Hiding homosexuality (6) Independent (3)
Harvard-educated & Carpenter	Nonconformist (6) Nonmaterialistic (6) Enjoys work (5) Easy going (4) Well rounded (4) Independent (3)
Leftist & Businessman	Difficult (3) Powerful (3) Serious (3)

Note. Numbers in parentheses represent number of subjects mentioning a given attribute in their descriptions of the combination.

the presence of causal antecedents, that is, descriptions that could be construed as accounting for what caused a member of one category to become a member of the other category. For the most part, these causal antecedents were not as detailed as the narratives illustrated above. An example of an antecedent given in a description of a person who is gay and a construction worker was:

Might be compensating for his "unusual tendencies" by taking a job that is considered "manly"

An example written as a description of a blind marathon runner was

Have a lot of faith in God and himself to be able to run without sight.

An example, written as a description of a Communist ex-Marine was

Had bad experiences in the Marine Corps. Now doesn't like country.

To be counted as a causal antecedent, a description had to contain explicit mention of causal relations. Thus, we did not count as an antecedent a mere listing of "courageous" as a descriptor of a person who is blind and a marathon runner, even though such a descriptor probably results from causal reasoning. Rather, we required that descriptions contain additional information making the causal relations clear, such as

... must have much courage and ability to overcome their handicap

We employed such stringent criteria to ensure that coders did not read causal reasoning that was not intended by subjects into subjects' descriptions, but our resulting coding probably underestimates the pervasiveness of antecedents. One coder coded all descriptions for presence of antecedents. A second, independent coder, coded a random subset of 40 descriptions, and agreed with the first coder on 90% of these. Using these strict criteria, we found numerous examples of causal antecedents. Seventy-five percent of the subjects who described the combinations used at least one antecedent in their descriptions. This suggests that causal reasoning was pervasive in subjects' attempts to combine these conflicting social categories.

There is some indication that causal reasoning was triggered by surprise. The combinations we used differed in the extent to which subjects found them surprising. The mean surprise ratings ranged from 2.29 for a person who was feminist and a bank teller to 5.25 for a person who was blind and a marathon runner. We counted, for each combination, the number of subjects who used antecedents in their descriptions, and correlated this measure with the average surprise rating given for each combination. The moderately high correlation obtained, $r = .46$, suggests that people were more likely to search for and construct explanations for why a member of one category became a member of the other when they found the joint membership more surprising. This method is clearly rather crude, and the resulting correlation could be subject to more than one interpretation. Nevertheless, the finding that causal antecedents may be generated in response to surprise is important, because it lends credence to the notion that the representations of members of combinations were constructed on line, through causal reasoning triggered by surprise, rather than retrieved from stored exemplars of persons belonging to these combinations.

STUDY 2

Study 1 suggests that when people form impressions of persons who belong to social categories that have conflicting implications for what a person might be like, they rely upon the causal relations within the categories and upon broader world knowledge to form unified impressions. Perhaps the

