
Implicit Quantification of Personality Traits

David Gidron

Hebrew University

Derek J. Koehler

Amos Tversky

Stanford University

The common usage of personality trait terms in the language includes an implicit quantification that is part of the accepted meaning of the term. This aspect of a trait's meaning is here called its scope. Traits with high scope, such as honest, require a high relative frequency of behavioral manifestation before they are attributed. In contrast, low-scope traits such as dishonest can be attributed on the basis of very few behavioral instances. A number of hypotheses are considered concerning the scope of trait terms within a language and between languages. Speakers of a given language (English or Hebrew) exhibit agreement in their ratings of scope; English and Hebrew speakers also agree with each other on the scope of trait terms even when they disagree about the behavioral manifestations of those traits. These findings are interpreted in terms of an informational view of personality traits: The scope of a trait is set at a level that makes it communicatively useful.

Personality traits are the stuff of lay psychology. They are used to describe people (Adam is warm and charming), explain their behavior (Eve did not speak because she is very shy), and predict future events (that plan won't work because David is too impulsive). Indeed, the study of personality traits has been the focus of several active programs of psychological research devoted to the measurement of traits (Wiggins, 1973), the classification of trait names in the English language (Allport & Odbert, 1936; Goldberg, 1981, 1982; Norman, 1963), the representation of the multitude of traits in terms of a small number of basic factors (Cattell, 1957; Eysenck & Eysenck, 1969; McCrae & Costa, 1987; Peabody & Goldberg, 1989), the temporal stability of traits (Block, 1971; Caspi, 1987; Caspi, Elder, & Bem, 1987, 1988), and their consistency across different situations (Bem & Allen, 1974; Epstein, 1983; Mischel, 1968, 1984; Mischel & Peake, 1982).

Personality traits have also been studied in attribution theory (Jones & Davis, 1965; Kelley, 1967), which seeks to identify inferential principles that govern causal reasoning in general and the inference of personality traits from behavioral evidence (i.e., consistency, distinctiveness, and consensus information) in particular. In contrast to the traditional focus of attribution theory on general principles, some recent work has suggested that different traits give rise to different inferential processes. Reeder and Brewer (1979), for example, distinguished among three trait schemata (partially restrictive, hierarchically restrictive, and fully restrictive) that are governed by different relations between an observed behavior and the dispositional attribution it supports. Their work suggests that traits not only specify a category of relevant behavior but also differ systematically in the way in which behavioral evidence must be combined in attributing the trait to a person.

In the present article we offer a different way of distinguishing among traits, in which traits are compared using a metric of behavioral quantification rather than discrete trait categories. Personality traits vary widely in the relative frequency of behavior they specify, which we will refer to here as the *scope* of a trait. For example, a sincere person is one who is almost always sincere. A person who is sometimes sincere and some-

Authors' Note: Preparation of this article was supported by a National Defense Science and Engineering Graduate Fellowship awarded to the second author. We thank Herb Clark, Tory Higgins, Yaacov Trope, and an anonymous reviewer for their helpful comments. We are especially grateful to Lew Goldberg for his extensive comments and insightful suggestions on earlier versions of this article. Correspondence concerning this article should be addressed to Amos Tversky, Department of Psychology, Stanford University, Stanford, CA 94305-2130.

PSPB, Vol. 19 No. 5, October 1993 594-604

© 1993 by the Society for Personality and Social Psychology, Inc.

times not is unlikely to be called sincere. However, a person can be called original if he or she exhibits originality only on some occasions. Thus, *sincerity* has high scope: It means being sincere (almost) always. In contrast, *originality* has low scope: It means being original in some but not necessarily all contexts.

Rorer and Widiger (1983, pp. 450-451) mentioned this phenomenon as an example of what they call ascription rules, which are the implicit rules governing the attribution of traits to individuals on the basis of behavioral information. They speculated that traits ascription might be based largely on the frequency of observed confirmatory instances of relevant behavior. Along these lines, Rothbart and Park (1986) investigated empirically the notion that traits differ in how easily they can be behaviorally confirmed or refuted. They found that the process of trait confirmation (or disconfirmation) depends on the specific trait in question because traits differ in (a) how much behavioral evidence they require for confirmation or disconfirmation, (b) how frequent are the situations that afford observation of diagnostic behavior, and (c) how easy it is to imagine confirming or disconfirming behavior relevant to the trait. Aspects b and c were highly correlated with each other but relatively independent of aspect a. Essentially, then, Rothbart and Park identified two factors influencing trait attributions. The first is the behavioral frequency required to confirm the trait, which is very similar to what we refer to here as scope. The second factor can be thought of as the availability of relevant behavioral evidence in the world. Together, these two factors can be used to account for the variation among traits in how easily they can be behaviorally confirmed or disconfirmed.

The present article addresses the problem of trait attribution from a slightly different perspective, one that focuses on the semantic rather than on the causal aspect of the attribution process. We are concerned with an observer's willingness to attribute a trait as a function of the shared meaning of the trait itself, rather than with the situational circumstances affecting the ease with which various traits can be detected on the basis of their behavioral manifestations. We argue that the common use of a trait term in the language includes an implicit quantification that is part of the accepted meaning of that term. The lexical meaning of a trait term typically does not specify the quantitative conditions of satisfaction associated with the trait. Instead, the definition of a trait usually conveys qualitative conditions: It specifies a category of behavior associated with the trait. There are other descriptive terms that are understood as either existential or universal quantifiers. For example, a *murderer* is someone who has murdered at least once (i.e., the term acts as an existential quantifier), and a *vegetarian* is someone who never eats meat (i.e., the term acts

as a universal negation). The definitions of personality traits such as *generous* or *conceited*, in contrast, do not specify exactly how often a person must behave in that manner before the trait becomes applicable. Personality traits typically do not carry existential or universal quantification: A single trait-consistent act is rarely sufficient for attributing a trait, and invariable trait-consistent behavior is not usually necessary. Instead, the requisite behavioral rate is some intermediate value that is not derivable from the lexical meaning of the trait. Speakers of a language may nevertheless have a shared idea of how often someone must behave in a certain way before a given trait should be attributed. What we mean by implicit quantification (or scope), then, is the consensus regarding the implied relative behavioral frequency of various trait terms in the language. The failure to recognize variation in the implicit quantification of traits may blur the boundary between the causal and the semantic aspects of the attribution process and confound empirical and conceptual issues. For example, an individual might be led to make a correspondent inference (Jones & Davis, 1965) on the basis of a single behavioral episode either because there are situational forces that would usually inhibit the behavior or because the trait in question has low scope.

In this article we report three experiments conducted to explore the nature of the implicit quantification of personality traits. We hypothesize that speakers of a given language exhibit reasonable agreement on the scope of different traits, and we propose a number of factors determining the scope of a trait. The scope of a trait is affected by the socially normative status of that trait, we argue, because pressure from social norms causes the behavioral manifestation of socially approved traits to be less informative than that of socially undesirable traits. For this reason, we hypothesize that desirable traits have higher scope than undesirable traits. But we suspect that scope within an evaluative domain (i.e., within the set of positive traits or the set of negative traits) depends on other factors essentially unrelated to social desirability. The intensity of a trait, for instance, affects its scope: Stronger traits (either positive or negative) should have higher scope than weaker traits. We further hypothesize that the scope of antonym pairs of traits are negatively correlated: The antonym of a high-scope trait has low scope, and vice versa.

We also investigate scope ratings across languages. Obviously, a specific behavior in one culture may lead to different trait attributions than it would in some other culture. This might come about either because the cultures differ in the behavioral categories specified by trait terms or because they differ in the scope associated with a given trait. We propose that speakers of different languages have reasonable agreement on the scope of trait

terms and that cultural differences in attribution arise primarily from disagreement over the classification of particular instances of behavior. Accordingly, we conjecture that scope ratings are similar across languages even when members of the cultures disagree on the specific behavioral manifestations of the traits under consideration. This last proposition suggests that cross-cultural agreement on the scope of trait terms does not depend directly on specific situational factors mediating the observation of behavioral evidence needed to make trait attributions. Instead, we suggest that the scope of a trait is connected with some deeper aspect of its meaning, which is preserved across languages.

ASSESSMENT OF SCOPE

A list of 60 personality traits was selected from previous studies (e.g., Allport & Odbert, 1936; Anderson, 1968; Goldberg, 1981; Rosenberg & Sedlak, 1972). Equal numbers of positive and negative traits, ranging from the most to the least desirable, were included. The list consists of 24 pairs of opposing traits (e.g., *reliable-unreliable*, *outgoing-shy*) plus 12 other, unpaired traits. Eight traits were obtained by adding the adverb *very* to 4 positive and 4 negative traits in the list (e.g., *very friendly*, *very selfish*). The list includes the most frequent traits used to describe people according to the count of Rosenberg and Sedlak (1972). An effort was made to represent the major domains and personality dimensions discussed in the literature (see Goldberg, 1981). The traits can be divided into three major domains: moral character (20 traits, e.g., *sincere-insincere*), intellect (14 traits, e.g., *brilliant-dull*), and temperament (26 traits, e.g., *warm-cold*). The entire list of traits appears in Table 1. The following instructions and rating scale were used to assess the scope of the selected traits:

Consider the following scale for describing the frequency of exhibiting some personality trait, say, *punctuality*.

0	1	2	3	4	5	6	7	8	9	10
Never	Rarely		Some- times		Often		Usually		Almost Always	Always

What level of frequency is required for calling a person punctual?

Clearly a person who is always (or almost always) on time is called punctual, and a person who is never (or rarely) on time is not called punctual. It seems that each trait has a *minimal frequency level* so that a person is called punctual, say, if his frequency of being on time is *above* that level and he is not called punctual if his frequency of being on time is *below* that level.

In the following pages you are asked to estimate the minimal frequency level of several traits. These questions, of course, do not have precise objective answers, but people tend to agree on the frequency levels required for different traits. Please consider each trait

separately and circle the appropriate number (from 0 to 10) on the scale above. Feel free to use the entire scale, not just the numbers with verbal labels.

Sixty-eight students enrolled in an overseas program at the Hebrew University of Jerusalem, all native speakers of English, assessed the minimal frequency level (MFL) of each trait. The subjects were run in a group; each took about 15 min to complete the questionnaire. The order of the traits was randomized and counterbalanced. A different group of 40 subjects rated the desirability of all traits on an 11-point scale (0 = *extremely undesirable*, 10 = *extremely desirable*).

Figure 1 presents the frequency distributions for a pair of antonyms. The figure shows that *honest* is a high-scope trait: To be considered honest, you have to be honest almost always. In contrast, *dishonest* is a low-scope trait: Even a person who exhibits dishonest behavior on only some occasions is called dishonest. The figure shows that the two distributions are clearly separated. Indeed, all our respondents assigned a higher MFL to *honest* than to *dishonest*. To provide a measure of the strictness of a trait, we will operationally define the *scope* of a trait $S(t)$, as the mean of its frequency distribution (i.e., the mean of the MFL ratings for that trait). Using this measure, *honest* has a scope of 8.6 whereas the scope of *dishonest* is only 3.8. Table 1 presents the estimated scopes (i.e., mean MFL ratings) for all 60 traits, from highest to lowest. The table also presents the mean rated desirability of each trait.

DETERMINANTS OF SCOPE

Table 1 supports our basic conjecture that different traits vary widely in scope; their scope ranges from 9.1 (very honest) to 3.8 (dishonest). These variations reflect large and consistent differences, not merely random fluctuations. (To interpret differences in scope, note that a difference of half a point on the scale is generally significant at the .05 level in a correlated t test.) Although the range of desirability is nearly twice that of scope (as one might expect), the average standard deviation for desirability ($SD = 1.45$) is not much smaller than that for scope ($SD = 1.70$). This observation suggests that there is almost as much agreement among raters regarding the scope of traits as there is regarding their desirability, despite the fact that scope is a novel attribute whereas desirability is perhaps the most basic dimension of meaning (Osgood, Suci, & Tannenbaum, 1957).

These results are consistent with those of Rothbart and Park (1986), who used a different operationalization of what we have here called scope. They asked 10 subjects to rate traits in terms of how many confirming behaviors a person would have to engage in before they would consider the trait to be an accurate description of that




Figure 1 Frequency distributions for a pair of antonyms.

person. The coefficient alpha value measuring between-judge agreement over the 150 traits being considered was .80, indicating a relatively high level of agreement. The coefficient alpha value for favorability ratings obtained from another 10 subjects was .97.

Desirability

The relation between desirability and scope may be illuminated by considering their relations to social or behavioral norms. It is reasonable to assume that the scope (i.e., the minimum relative frequency level required for attribution) of a trait corresponds to a socially significant distinction between individuals—namely, between those to whom the trait applies and those to whom it does not. Hence, the implicit quantification of trait terms may reveal some significant behavioral distinctions. We make the following two assumptions in order to derive a prediction about the relationship between desirability and scope. First, we assume that social norms are generally defined by positive traits (e.g., *honest*, *punctual*, or *wise*) whereas the negative traits represent violations of the norm. (This difference might be reflected in the fact that most of the negative trait adjectives are linguistically marked with respect to their positive counterparts.) Second, we assume that it is more important to distinguish between people who obey the norm and

those who violate it than to distinguish among people who violate the norm or among those who do not violate it. For instance, the distinction between a person who is always reliable and a person who is not appears more significant than the distinction between a person who is always unreliable and a person who is sometimes unreliable. It follows from these assumptions that the scope of positive (i.e., normative) traits such as *reliable* should be considerably higher than the scope of the negative traits such as *unreliable*, so that the positive trait refers only to those who follow the norm almost always and the negative trait refers to those who violate it, whether occasionally or often. Although this analysis implies that positive traits should have higher scope than negative traits, it leads to no predictions about the relative scope of traits within an evaluative domain.

The desirability ratings of the 60 traits used in the current study yielded a bimodal distribution, composed of 29 positive traits (with mean desirability ratings above 7), 4 neutral traits (with mean ratings between 5 and 6), and 27 negative traits (with mean ratings less than or equal to 4). A striking feature of Table 1 is that the scope of the positive traits greatly exceeds that of the negative traits. Only two positive traits (*imaginative* and *original*) have a scope below the median (6.3), and only three negative traits (*passive*, *rigid*, and *very selfish*) have a scope that exceeds the median. Thus, in almost every case, the normative traits have higher scope than the counternormative traits. Rothbart and Park (1986) also report a high correlation ($r = .71$) between judged favorability and the number of required confirming behavioral instances.

There is, however, little or no relation between desirability and scope within the positive or the negative domain. For example, *interesting* is rated as more desirable than *punctual*, but the latter has a higher scope. Similarly, *boring* is rated as less desirable than *selfish*, but the order of their scopes is reversed. The overall product-moment correlation between desirability and scope is .76, whereas the correlations between these ratings within the positive and the negative domains are $-.07$ and $.27$, respectively. Thus, as predicted, desirability provides a near-perfect separation of high-scope from low-scope traits, but it does not account for scope within each domain. Within a domain, there is still wide variation in the scope of trait terms, which must be accounted for by some factor or factors unrelated to desirability.

Complementarity

The preceding discussion suggests a complementary relation between the scope of antonyms. Traits with higher scope (implying a high rate of trait-consistent behavior) should be paired with antonyms of low scope (implying a low rate of trait-inconsistent behavior). If one must behave honestly almost always to be considered an hon-

TABLE 1: Mean Ratings of Scope and Desirability for 60 Personality Traits

Trait	Positive Traits		Trait	Neutral and Negative Traits	
	Scope	Desirability		Scope	Desirability
Very honest	9.08	7.79	Rigid	7.30	2.54
Honest	8.58	9.14	Passive	6.37	3.43
Very friendly	8.49	7.56	Very selfish	6.34	0.79
Very intelligent	8.41	8.14	Very hostile	6.21	0.28
Very generous	8.35	7.35	Cautious	6.21	5.74
Punctual	8.18	7.56	Very dull	6.08	0.17
Brilliant	7.98	7.74	Unimaginative	5.95	2.19
Loyal	7.83	8.92	Boring	5.72	1.12
Reliable	7.83	9.09	Very violent	5.63	0.09
Sincere	7.73	8.98	Aggressive	5.61	5.51
Unselfish	7.67	7.96	Shy	5.61	4.02
Wise	7.42	8.44	Dull	5.57	1.12
Polite	7.33	8.13	Tense	5.52	2.54
Mature	7.29	7.98	Cold	5.52	1.40
Intelligent	7.16	8.70	Immature	5.42	1.65
Loving	7.14	8.84	Selfish	5.40	2.19
Gentle	7.05	8.21	Impulsive	5.24	5.63
Warm	6.97	8.43	Stingy	5.23	1.42
Relaxed	6.89	7.72	Foolish	5.07	1.56
Outgoing	6.80	7.91	Irritable	5.00	1.63
Active	6.79	7.86	Temperamental	4.97	2.61
Friendly	6.79	8.51	Hostile	4.89	1.33
Flexible	6.77	7.93	Risky	4.89	5.86
Patient	6.73	8.54	Unreliable	4.84	0.74
Happy	6.53	8.54	Impatient	4.83	1.67
Generous	6.47	8.02	Insincere	4.81	0.98
Interesting	6.38	8.74	Unfriendly	4.71	1.40
Imaginative	6.14	8.45	Rude	4.64	0.83
Original	5.94	8.00	Cruel	4.36	1.00
			Violent	4.16	1.00
			Dishonest	3.78	0.72

NOTE: Ratings could range from 0 to 10; higher numbers indicate higher scope or greater desirability.

est person, then one need only behave dishonestly occasionally to be considered dishonest. It follows that, in general, if x and y are positive traits and x^* and y^* are their opposites, then

$$S(x) > S(y) \text{ implies } S(x^*) < S(y^*)$$

Complementarity of traits entails a negative correlation between the scopes of antonyms. Indeed, the correlation between the scopes of the 13 antonym pairs rated by our subjects is $-.76$, $p < .05$. Table 2 lists the scope of each member of all 13 antonym pairs, along with the sum and difference in scope across each pair.

The approximate complementarity of opposite traits is not a logical consequence of negation. If each individual has either the trait t or else the trait t^* then $S(t) + S(t^*)$ should approximately equal 10 (thus, there should be no relative frequency range in which neither trait is attributed) at least to the extent that our verbal relative frequency labels designate complementary rates of behavior (e.g., *usually honest* is roughly equivalent to *sometimes dishonest*). The finding that the sum $S(t) + S(t^*)$

exceeds 10 for all 13 antonym pairs ($M = 12.39$) shows that the antonyms were interpreted as contrary, not as merely exclusive. That is, the observed judgments imply a behavioral frequency range in which neither the trait nor its antonym is attributed to the person.

The postulation of a "neutral" frequency is consistent with results obtained by Goldberg and Kilkowski (1985) showing that self-ratings for a given trait are not completely consistent with self-ratings for the antonym of that trait. This is what would be expected if there is a frequency interval in which neither trait is attributed. These researchers' observation that antonym pairs vary widely in the consistency they elicit, taken in combination with our analysis of complementarity of scope ratings, suggests that although the width of the neutral range may be approximately constant, the number of individuals who fall into this range (according to their self-ratings) may vary systematically between antonym pairs. Goldberg and Kilkowski (1985) also found that synonym pairs yielded more consistent self-ratings than antonym pairs. The present results raise the possibility

TABLE 2: Complementarity Analysis of 13 Antonym Pairs Including the Sum and Difference of Their Scores

Trait	Scope	Antonym	Scope	Sum	Difference
Honest	8.58	Dishonest	3.78	12.36	4.80
Brilliant	7.98	Dull	5.57	13.55	2.41
Reliable	7.83	Unreliable	4.84	12.67	2.99
Sincere	7.73	Insincere	4.81	12.54	2.92
Unselfish	7.67	Selfish	5.40	13.07	2.27
Wise	7.42	Foolish	5.07	12.49	2.35
Polite	7.33	Rude	4.64	11.97	2.69
Mature	7.29	Immature	5.42	12.71	1.87
Warm	6.97	Cold	5.52	12.49	1.45
Friendly	6.79	Unfriendly	4.71	11.50	2.08
Patient	6.73	Impatient	4.83	11.56	1.90
Interesting	6.38	Boring	5.72	12.10	0.66
Imaginative	6.14	Unimaginative	5.95	12.09	0.19

NOTE: Ratings could range from 0 to 10; higher numbers indicate higher scope.

that inconsistency in synonym ratings may result from rating error alone whereas the greater degree of inconsistency in antonym ratings may result from the existence of a neutral range in which neither the trait nor its antonym is attributed.

Intensity

Variations in the intensity of a trait are sometimes expressed by different adjectives. For example, *brilliant* is stronger than *bright*, and *hostile* is stronger than *unfriendly*. Alternatively, such variations can be expressed by the use of intensive adverbs (e.g., *very*, *highly*, *extremely*). Cliff (1959) has suggested that the addition of such modifiers increases the intensity of a term in an approximately multiplicative fashion. We propose that greater intensity is associated with higher scope within a set of trait terms that otherwise would appear to refer to the same personality dimension. In short, stronger traits require a higher relative frequency of manifestation. A person who is called highly creative or extremely shy is expected to exhibit creativity and shyness more frequently than a person who is described merely as creative or as shy.

This hypothesis was tested by observing the effect of adding the intensive adverb *very* to four positive traits (*intelligent*, *honest*, *generous*, *friendly*) and to four negative traits (*selfish*, *hostile*, *dull*, *violent*). In every case, the addition of *very* produced a significant increase in scope. Although this result is not very surprising, it is not without interest: The addition of *very* might have changed the kinds of behavior implied by the trait without increasing the implied frequency of these behaviors. The fact that scope was increased may reflect an implicit linguistic principle that a strong claim (e.g., David is very generous) requires stronger behavioral evidence than a weaker claim (David is generous).

Interestingly, the addition of *very* reduced the desirability of both the negative and the positive traits. Thus, *very friendly* was rated as less desirable than *friendly*, and

so on. Evidently, an excess is undesirable, even for a positively valued trait. In other words, the unmodified positive trait terms appear to implicitly specify the optimal rate of normatively appropriate behavior, at least for the traits we have examined here. The fact that intensity increases scope but decreases desirability attenuates the correlations between scope and desirability within the positive and the negative traits.

Comparative Analysis


To test the hypothesis that scope ratings would coincide across languages, we replicated the basic study in Hebrew. Of the 60 traits used in the initial study, we selected 35 that could be translated most directly into Hebrew. A group of 62 students whose major language is Hebrew evaluated the scope of all 35 traits, following the procedure described above. A different group of 27 students rated the desirability of these traits. Despite the fact that the literal translation of a trait does not preserve its exact meaning—let alone its connotations—we found remarkably close agreement ($r = .93$) between the assessed scope of traits in the two languages. The correlation between the ratings of desirability ($r = .96$) was slightly higher.

The relation between desirability and scope in the Hebrew data mirrors the pattern observed in the English data. With only a few exceptions, the desirable traits had higher scope than the undesirable traits, but there was no strong relation between desirability and scope within the positive and negative domains. The correlation between desirability and scope was .83 over the entire set but declined to .38 and .30 within the positive and the negative subsets, respectively.

INTERPRETATION AND ATTRIBUTION

In the preceding sections, we focused on the attribution of a trait to a person as a function of the prevalence

of its manifestations. In the present section, we deal with the relation between the interpretation, or categorization, of specific behaviors and the attribution of traits (see Reeder & Brewer, 1979; Trope, 1986). Consider some specific behavior, such as helping (or refusing to help) a stranger fix a flat tire. This behavior is generally interpreted as helpful (categorization), and a person who exhibits such behavior is regularly called *helpful* (attribution). Schematically, we have



where the first arrow refers to the categorization $C(b)$ of a particular behavior b , and the second arrow refers to the attribution $A[C(b)]$ of a trait to a person on the basis of a characterization of his or her behavior. So far we have focused on the attribution process (i.e., the second arrow) and have demonstrated a close correspondence between the assessed scopes of personality trait terms in English and Hebrew. Cultures clearly differ in the way traits are attributed on the basis of behavioral information, however, and this difference must arise either from differences in the attribution function or from differences in the categorization function (or both). Because we have already found some evidence that—in the two populations we have sampled from, at least—there is relatively strong cross-cultural agreement on scope ratings, we would like to suggest that cultural differences arise primarily from differences in the categorization function.

Accordingly, we next explore the hypothesis that cultures with similar attribution functions may, nevertheless, have very different categorization functions based on different behavioral norms. Members of two different cultures (or speakers of different languages) may agree that to be called *honest*, a person must exhibit honest behavior almost always (i.e., they agree about this scope), but they may differ greatly about the behavioral manifestations of honesty. In other words, they may agree on how often the prototypical honest person acts honestly but at the same time disagree on what constitutes prototypical honest behavior.

To investigate this hypothesis, we constructed a set of 12 behavioral vignettes, each expressing to a greater or lesser extent some personality trait. The entire set of vignettes is presented in Table 3. The first six (V1-V6, called Set A) describe behaviors hypothesized to be about equally common in Israeli and American cultures. Three describe negative behavior (V1-V3) and three describe positive behavior (V4-V6). So, for example, offering one's bus seat to an old man (V6) or fixing an electrical appliance without unplugging it first (V2) was suspected to occur equally often in both cultures.

The other six vignettes (V7-V12, called Set B) describe behaviors hypothesized to be more common in Israeli

TABLE 3: Twelve Behavioral Vignettes With Their Associated Traits

Set A	
V1.	(Forgetful) To arrive at the airport for an international flight and leave one's passport at home.
V2.	(Stupid) To fix an electrical appliance without unplugging it first.
V3.	(Risky) To go on an out-of-town car trip with a defective spare tire.
V4.	(Honest) To return a \$10 bill received by mistake from a bank teller.
V5.	(Helpful) To help a stranger who lost his keys.
V6.	(Polite) To offer a seat in a crowded bus to an old man.
Set B	
V7.	(Impolite) To respond to the arguments of an acquaintance before he finishes them.
V8.	(Annoying) To ask a stranger who is sitting near you on the train for his newspaper just after he finished reading it.
V9.	(Rude) To honk at the car ahead when it does not move immediately as the traffic light turns green.
V10.	(Considerate) To drive two extra miles to bring a hitchhiker to his home on a rainy day.
V11.	(Helpful) To interrupt a conversation between two strangers standing behind you in a long line to tell one of them that the official he wishes to see is not in the office today.
V12.	(Generous) To lend money to a stranger whose wallet was stolen.

than in American culture. Again, three are negative (V7-V9) and three are positive (V10-V12). For example, honking at a car when it does not move immediately after a light turns green (V9) and driving an extra two miles to bring a hitchhiker to his home on a rainy day (V10) were thought to be more frequent in Israeli culture than in American culture. We predicted that although American and Israeli subjects would agree on the scope of all the traits associated with the vignettes, they would only agree on the categorization of the behaviors described in the first six vignettes. Specifically, we expected that the American judges would find the behaviors that are relatively infrequent in American culture to be more diagnostic of the actor's disposition than the Israeli judges would (see also Jones & Davis, 1965).

The vignettes were presented in Hebrew to a group of 68 Israeli students from the Hebrew University in Jerusalem and in English to a group of 49 American students from Stanford University. For each vignette, the students were asked to rate (a) the *frequency* of the described behavior (1 = *very rare*, 5 = *very frequent*), and (b) the *extent* to which the behavior described expresses a given trait (1 = *not at all*, 5 = *to a very large extent*). In addition, the students were asked to rate either the scope or the desirability of a list of 24 traits, which included the 12 traits associated with the vignettes. The mean ratings of frequency and extent for each vignette are presented in Table 4 for both the Israeli and American students.

TABLE 4: Mean Ratings of Frequency and Extent for the 12 Behavioral Vignettes of Table 3, by Israeli and American Students

Vignette	Trait	Frequency		Extent	
		Israeli	American	Israeli	American
<i>Set A</i>					
V1	Forgetful	1.73	1.85	4.14	4.19
V2	Stupid	2.48	2.41	3.94	3.95
V3	Risky	2.93	2.66	3.71	3.45
V4	Honest	3.37	3.09	4.64	4.23
V5	Helpful	2.71	2.80	4.16	4.10
V6	Polite	3.75	3.83	3.95	4.03
Mean		2.83	2.77	4.09	3.99
<i>Set B</i>					
V7	Impolite	4.22	3.72	3.40	3.95
V8	Annoying	3.35	3.17	2.71	3.12
V9	Rude	4.21	3.63	3.83	4.31
V10	Considerate	2.95	2.29	4.43	4.57
V11	Helpful	3.62	2.92	2.19	2.52
V12	Generous	3.08	1.85	4.11	4.32
Mean		3.57	2.93	3.45	3.80

NOTE: Ratings could range from 1 to 5; higher numbers indicate greater frequency or greater extent.

The frequency ratings confirmed our initial classification of the vignettes. In Set A there were no systematic differences in the judgments of frequency made by the American and the Israeli students, $t(115) = 1.2$, n.s. In Set B, however, the average frequency judgment was significantly higher for the Israeli than for the American students, $t(115) = 3.7$, $p < .01$.

We expected that judgments of extent (or diagnosticity) would be inversely related to judged frequency: A given behavior is less diagnostic of the underlying trait when the behavior is common than when it is rare. Indeed, the judgments of extent were significantly higher for the American than for the Israeli students in Set B, $t(115) = 2.8$, $p < .01$. In contrast, there was no significant difference between the judgments in Set A, $t(115) = 1.4$, n.s., where the behaviors were judged to be equally common in the two cultures. Offering a seat in a crowded bus to an old man (V6), for example, was judged about equally frequent and equally polite by both groups. However, honking at the car ahead when it does not move immediately as the traffic light turns green (V9) was considered rarer and ruder by the American students than by the Israeli students. The extent ratings confirm that we have one set of behaviors that are interpreted similarly in the two cultures and another set that yields different interpretations.

The major result of this study is that, in contrast to cultural differences in their categorization of the vignettes, there were no significant differences between American and Israeli ratings of the scope and the desirability of the relevant traits, either in Set A or in Set B. These observations suggest that cultures may vary more in the categorization of behaviors $C(b)$ than in the attribution of traits on the basis of such categorization $A[C(b)]$.

DISCUSSION

The studies reported in this article lend support the hypothesis that the use of personality traits involves an implicit quantification (called *scope*) that serves as part of the conditions of satisfaction for the attribution of traits. Ratings of scope show strong agreement, even though presumably most people have not explicitly considered the scope of various traits before. Furthermore, the distribution of scope for traits in Hebrew closely resembles that of their English counterparts, despite the fact that Israeli and American students differed with respect to the categorization of specific behaviors. Such results suggest that the ratings of scope reflect an implicit quantification that is inherent in the meaning of trait terms themselves rather than a shared cultural stereotype specifying which traits are associated with which behaviors. Our results suggest certain systematic relations between scope and other aspects of the meaning of trait terms: Positive traits have higher scope than negative traits, scope increases with intensity (even as desirability decreases), and antonyms tend to have complementary scopes. In this section we explore some implications of these findings for attribution theory, personality research, and the semantics of trait terms.

The present findings are generally compatible with the spirit of attribution theory. The relation between scope and desirability is consistent with the role of norms: Normative behavior is less diagnostic, and hence it requires higher scope (Jones & Davis, 1965). In addition, the relation between scope and intensity is consistent with the assumption that a stronger attribution requires stronger supporting evidence. The present findings, however, also highlight the semantic aspect of trait

attribution that is not fully captured by a causal analysis. Clearly, if the meanings of various trait terms differ in the amount of evidence they implicitly require for attribution, then an adequate model of the attribution process will need to be sensitive to context effects arising from semantic differences among traits. For example, the work of Goldberg and his colleagues concerning the semantics of personality trait terms, both on the hierarchical relations among trait terms (e.g., Hampson, John, & Goldberg, 1986) and on the traits that have universal linguistic representation (e.g., Goldberg, 1981), carries obvious implications for research on attribution processes. For related work that attempts to extend attribution research to accommodate variations in the schematic representation of different attributes, see Reeder and Brewer (1979) and Rothbart and Park (1986).

The observation that traits exhibit large and systematic variation in scope has some implications for the ongoing debate regarding the cross-situational consistency of personality traits. If some attributes, such as dishonesty, have low scope, then these traits are expected to exhibit a low level of cross-situational consistency. That is, if being dishonest means exhibiting a dishonest act on *some* occasion, then this trait has a limited predictive value owing to the low behavioral base rate it reflects (Meehl & Rosen, 1955). An immediate consequence of the present analysis is that the degree of cross-situational consistency of a trait should be positively correlated with its scope.

In fact, low-scope traits allow for a complete lack of cross-situational consistency. A person who sometimes cheats on exams can legitimately be called dishonest even if he or she never steals, lies, or avoids paying income tax. As a result, two dishonest persons may have no specific behavioral tendencies in common. Indeed, there may be no predictability even within a situational category for a low-scope trait. Someone who commits a single heinous act, for example, will be considered immoral even if he or she never commits another immoral act. Despite their low predictability, concepts of low scope may be useful in both ordinary and scientific discourse. For instance, even if knowing that someone has been labeled dishonest is not predictive of future dishonesty, attribution of this trait may nevertheless be very significant because it serves to mark a clear past violation of a strongly held norm.

The degree of agreement among speakers, both within and across languages, regarding the scope of personality traits is particularly impressive in the face of the enormous vagueness of many trait terms. To say that someone is flexible or immature might bear on any of a great number of different sorts of behaviors in a wide variety of situations. Despite the subtleties of meaning associated with trait terms, however, people show strong agree-

ment on how often the relevant behaviors (whatever they may be) must occur before the trait will be ascribed. The agreement between the Israeli and the American data regarding scope despite the marked disagreement regarding the categorization of behavioral instances suggests that the scope of a trait is tied to some abstract aspect of its meaning.

The relationship between the meaning of a trait and its scope is a curious one. On the one hand, people appear to agree on the scope of various trait terms, suggesting that the shared meaning of these terms includes an aspect of quantification. On the other hand, this quantification is truly implicit in the sense that it is derived neither from the lexical meaning of the term nor from the term's immediate linguistic function. As mentioned earlier, it would not be surprising to find that people agree on the scope of terms with definitions that entail some sort of quantification. But there is nothing in the lexical definition of *generous* that specifies exactly how often one must behave generously to be deemed generous. Furthermore, one might have expected only a limited variation in scope, owing to a spontaneous adjustment in the interpretation of the trait's intensity. If we were to interpret *honest*, for example, as "reasonably honest," and *dishonest* as "less than perfectly honest," the discrepancy between their scope would be greatly reduced.

Trait terms, which have an explicit function that appears quite separate from their implicit quantification, can be contrasted with other terms that have only a quantitative function. Probabilistic terms (e.g., *quite certain*, *highly unlikely*), for instance, are employed specifically to quantify a person's sense of certainty. A number of studies (e.g., Beyth-Marom, 1982; Lichtenstein & Newman, 1967; see Mosteller & Youtz, 1990, for a review) provide evidence of substantial variation in the perceived meaning of probabilistic expressions. Although the methodology of these studies differs from that used here, making direct comparisons difficult, a general reading of the results suggests that the standard deviation of people's numerical interpretations of probabilistic expressions are roughly comparable to those of our subjects evaluating the scope of trait terms. This is a rather surprising observation, given that probabilistic expressions are used with the explicit intention of quantifying one's sense of certainty. Personality trait terms, in contrast, although they may indirectly specify something about the relative frequency of various behaviors, serve many functions in everyday conversation that are essentially unrelated to quantification: On the surface, at least, they seem to specify what sort of behavior—not how much—characterizes the individual.

In summary, then, the current findings imply that there are two relatively independent aspects of a trait

term's meaning. The more explicit, qualitative aspect specifies a category of behavior, which may vary from culture to culture. The implicit, quantitative aspect specifies the relative frequency of behavior necessary to attribute the trait to a person. We suggest that this second aspect does not appear to be specific to a given language or to depend on cultural stereotypes, at least across the two subject populations we have studied here.

Several issues remain unresolved. An immediate question arises from the fact that the meanings of relative frequency terms such as *often* or *rarely* are themselves modified by the term they are being used to modify (Pepper, 1981). To say that it rains very often, for example, implies a higher frequency than to say that earthquakes occur very often. As a result, asking subjects to rate the scope of personality traits in terms of the appropriateness of various relative frequency expressions runs the risk of confounding the subtle factors influencing the implicit quantification of each. We are somewhat reassured of the reliability of scope ratings, though, as the high agreement on scope and the relation between scope and desirability have been documented independently by Rothbart and Park (1986).

It might also be instructive to investigate the scope of traits that are applied to inanimate objects. For example, we commonly apply many trait terms (e.g., *dangerous*, *reliable*, *flexible*, *efficient*) to home appliances and automobiles as well as to humans. We conjecture that the complementarity of scope ratings, and the effects of desirability and intensity, will hold also for these categories. This may be true either because the same attribution principles apply to objects and to people or because we "project" our concept of personality traits to other categories as well. Indeed, it is important to note that implicit quantification is not limited to trait terms but instead is found in numerous linguistic domains. Verbs, for example, may implicitly quantify their objects. Kanouse (1972) found that English speakers generally agree on the scope of verbs and identified two dimensions that appear to partly determine the scope of a verb.

But the most important question still remains: What determines the scope of a given trait? We would like to suggest that the scope of a trait is set at a level that makes the term socially useful. The categories specified by trait terms are informative only if they discriminate among individuals, such that the trait applies to some people but not to others. From this it follows that the scope of a trait should reflect, to some degree, the relative frequency of trait-consistent behavior occurring in the world. Because most people behave honestly most of the time, the term *honest* would not be informative if it meant "reasonably honest" (because then it would apply to almost everyone) rather than "almost always honest." As a conse-

quence, the threshold for applying the term *honest* is set quite high. Likewise, if the term *dishonest* were applied only to those who constantly behave dishonestly, then virtually nobody would be considered dishonest. Note that it is for this reason that social norms help determine the scope of traits: Social norms exert a strong influence on the sorts of behavior that usually takes place. The crucial determinant, however, is the rate of relevant actual behavior (which can be thought of as a *statistical* norm) rather than social norms themselves. Creativity, for example, is a socially approved trait, but there are other factors (having to do with intellectual capacity and even the definition of a creative act) that necessarily limit how often people behave creatively. As a result, creative behavior is relatively rare, and the low scope of the term *creative* reflects this fact. Social norms, then, are one of many factors influencing the statistical frequency of relevant behavior and thus the scope of personality trait terms.

Rothbart and Park (1986) suggested that the differential availability of behavioral evidence in the world leads to differences in how easily traits can be detected. Here, we are offering a different perspective: The pressure to apply trait terms that discriminate among individuals causes differences among these traits in the standards of evidence they require for confirmation. We are not arguing, however, that the scope of a trait term is set to *maximize* discrimination among individuals (e.g., to make equal the numbers of people described by antonym pairs of traits). Other factors undoubtedly affect scope as well. But it is our belief that the scope of a trait is set so as to avoid including almost everyone or virtually no one in the category being specified. The accepted meaning of a trait, then, is selected to make the term informationally useful in describing differences among individuals. In reaching this conclusion, we complete a cycle that began with the observation that the scope of traits affects how they are attributed. We conclude with the suggestion that although the scope of a trait determines the conditions for its attribution at any point in time, the scope of a trait itself is adjusted over time in response to the frequency its attribution.

REFERENCES

- Allport, G. W., & Odbert, H. S. (1936). Trait names: A psycho-lexical study. *Psychological Monographs*, 47(1, Whole No. 211).
- Anderson, N. H. (1968). Likeableness ratings of 555 personality-trait words. *Journal of Personality and Social Psychology*, 9, 272-279.
- Bem, D. J., & Allen, A. (1974). On predicting some of the people some of the time: The search for cross-situational consistencies in behavior. *Psychological Review*, 81, 506-520.
- Beyth-Marom, R. (1982). How probable is probable? Numerical translation of verbal probability expressions. *Journal of Forecasting*, 1, 275-269.
- Block, J. H. (1971). *Lives through time*. Berkeley, CA: Bancroft Books.
- Caspi, A. (1987). Personality in the life course. *Journal of Personality and Social Psychology*, 53, 1203-1213.

- Caspi, A., Elder, G. H., Jr., & Bem, D. J. (1987). Moving against the world: Life-course patterns of explosive children. *Developmental Psychology, 22*, 303-308.
- Caspi, A., Elder, G. H., Jr., & Bem, D. J. (1988). Moving away from the world: Life-course patterns of shy children. *Developmental Psychology, 24*, 824-831.
- Cattell, R. B. (1957). *Personality and motivation structure and measurement*. New York: World Books.
- Cliff, N. (1959). Adverbs as multipliers. *Psychological Review, 66*, 27-44.
- Epstein, S. (1983). The stability of confusion: A reply to Mischel and Peake. *Psychological Review, 90*, 179-184.
- Eysenck, H. J., & Eysenck, S.B.G. (1969). *Personality: Structure and measurement*. San Diego, CA: Knapp.
- Goldberg, L. R. (1981). Language and individual differences: A search for universals in personality lexicons. *Review of Personality and Social Psychology, 2*, 141-165.
- Goldberg, L. R. (1982). From Ace to Zombie: Some explorations in the language of personality. In C. D. Spielberger & J. N. Butcher (Eds.), *Advances in personality assessment* (pp. 203-234). Hillsdale, NJ: Lawrence Erlbaum.
- Goldberg, L. R., & Killowski, J. M. (1985). The prediction of semantic consistency in self-descriptions: Characteristics of persons and of terms that affect the consistency of responses to synonym and antonym pairs. *Journal of Personality and Social Psychology, 48*, 82-98.
- Hampson, S., John, O. P., & Goldberg, L. R. (1986). Category breadth and hierarchical structure in personality: Studies of asymmetries in judgments of trait implications. *Journal of Personality and Social Psychology, 51*, 37-54.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 219-266). New York: Academic Press.
- Kanouse, D. E. (1972). Verbs as implicit quantifiers. *Journal of Verbal Learning and Verbal Behavior, 11*, 141-147.
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska Symposium on Motivation* (Vol. 15, pp. 192-241). Lincoln: University of Nebraska Press.
- Lichtenstein, S., & Newman, J. R. (1967). Empirical scaling of common verbal phrases associated with numerical probabilities. *Psychonomic Science, 9*, 563-564.
- McCrae, R. R., & Costa, P. T., Jr. (1987). Validation of the five-factor model of personality across instruments and observers. *Journal of Personality and Social Psychology, 52*, 81-90.
- Meehl, P. E., & Rosen, A. (1955). Antecedent probability and the efficacy of psychometric signs, patterns or cutting scores. *Psychological Bulletin, 52*, 194-216.
- Mischel, W. (1968). *Personality and assessment*. New York: Wiley.
- Mischel, W. (1984). Convergences and challenges in the search for consistency. *American Psychologist, 39*, 351-364.
- Mischel, W., & Peake, P. K. (1982). Beyond déjà vu in the search for cross-situational consistency. *Psychological Review, 89*, 730-755.
- Mosteller, F., & Youtz, C. (1990). Quantifying probabilistic expressions. *Statistical Science, 5*, 2-34.
- Norman, W. T. (1963). Toward an adequate taxonomy of personality attributes: Replicated factor structure in peer nomination personality ratings. *Journal of Abnormal and Social Psychology, 66*, 574-583.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The measurement of meaning*. Urbana: University of Illinois Press.
- Peabody, D., & Goldberg, L. R. (1989). Some determinants of factor structures from personality-trait descriptors. *Journal of Personality and Social Psychology, 57*, 552-567.
- Pepper, S. (1981). Problems in the quantification of frequency expressions. In D. W. Fiske (Ed.), *Problems with language imprecision* (pp. 25-42). San Francisco: Jossey-Bass.
- Reeder, G. D., & Brewer, M. B. (1979). A schematic model of dispositional attribution in interpersonal perception. *Psychological Review, 86*, 61-79.
- Rorer, L. G., & Widiger, T. A. (1983). Personality structure and assessment. *Annual Review of Psychology, 34*, 431-463.
- Rosenberg, S., & Sedlak, A. (1972). Structural representations of perceived personality trait relationships. In A. K. Romney, R. N. Shepard, & S. B. Nerlove (Eds.), *Multidimensional scaling: Theory and applications* (pp. 133-162). New York: Seminar Press.
- Rothbart, M., & Park, B. (1986). On the confirmability and disconfirmability of trait concepts. *Journal of Personality and Social Psychology, 50*, 131-142.
- Trope, Y. (1986). Identification and inferential processes in dispositional attribution. *Psychological Review, 93*, 239-257.
- Wiggins, J. S. (1973). *Personality and prediction: Principles of personality assessment*. Reading, MA: Addison-Wesley.