European Journal of Social Psychology
Eur. J. Soc. Psychol. **40**, 1120–1135 (2010)
Published online 12 October 2009 in Wiley Online Library
(wileyonlinelibrary.com) **DOI**: 10.1002/ejsp.705

Research article

On the inferential epistemics of trait centrality in impression formation

EDWARD OREHEK[1]*, MARK DECHESNE[1],
AYELET FISHBACH[2], ARIE W. KRUGLANSKI[1]
AND WOO YOUNG CHUN[3]
[1]*University of Maryland, USA*
[2]*University of Chicago, USA*
[3]*Chungnam National University, Korea*

Abstract

We provide a novel, inferential, account of the trait centrality phenomenon. We suggest that a trait possesses the property of "centrality" to the extent that it is subjectively deemed to imply other traits. Five studies explore four central elements of this view. First, trait relations can be stored as unidirectional rules ("if X then Y" but not necessarily "if Y then X"). Second, the strength of individuals' lay inference rules determines the effect of traits on impressions. Third, situationally manipulating the strength of lay inference rules influences the impact of traits on impressions. Fourth, the impact of an inference rule is reduced when it is difficult to discern the inference rule and when processing resources are limited.

Forming impressions of persons constitutes a key aspect of social interaction, as beliefs about others drive to a considerable degree the interpersonal decisions and actions one undertakes. For example, an impression that a politician is principled may prompt one to lend her or him one's vote, a supposition that a lawyer is articulate may lead one to choose her or him as a defender, and an opinion that an athlete is motivated may determine her or his choice as an inductee for a team.

Prior research has suggested that in forming impressions, people organize a large amount of behavioral evidence around a few central traits (Asch, 1946; Goldman, Cowles, & Florez, 1983; Kelley, 1950; Sedikides & Skowronski, 1993; Watkins & Peynircioglu, 1984; Wyer, Bodenhausen, & Srull, 1984). In this sense, Solomon Asch's (1952, p. 208) distinction between ''qualities that furnish the key to a person and those that are subsidiary'' remains critically important for understanding impression formation. Though past research has amply documented that different traits may vary in their degree of centrality, the question has remained what, exactly, contributes to a trait's centrality. Asch (1946) originally posited that trait centrality resides in the traits themselves, representing a property of their specific *contents*. However, follow-up research has challenged this position, and suggested that trait centrality emerges out of the specific relations between co-occurring traits (e.g., principled–trustworthy, articulate–intelligent). This research highlighted the *relational*, rather than essentialist, nature of the phenomenon (Wishner, 1960; Zanna & Hamilton, 1972).

We use principles derived from Lay Epistemic Theory (Kruglanski, 1989) to shed light on the phenomenon of trait centrality, and the representation of trait relations in memory more generally. In brief, we propose that information that a target of evaluation possesses a given trait functions as *evidence* for inferences about this target's other traits and properties. In this sense, a trait is central to the extent it implies other traits, thereby affording generalized social judgment. Three elements of this proposition are of present interest: (1) The *structure* of the inference rule that links a specific trait to

*Correspondence to: Edward Orehek, Department of Psychology, University of Maryland, College Park, MD 20742, USA.
E-mail: eorehek@psyc.umd.edu

Received 16 May 2008
Accepted 20 July 2009

other traits to enable a judgment; (2) the notion of *subjective relevance*, which refers to the strength with which a trait is believed to imply other traits; and (3) the *difficulty* of gleaning the subjectively relevant evidence from the environment. We present six studies that attest to the unidirectional ''if. . . then'' structure of the mental representation of trait relations, and the role of lay inferences in the trait centrality phenomenon. We also show that the ability to use trait information as evidence for a given impression depends on task demands and individuals' available processing resources.

TRAIT CENTRALITY

The notion that some personality traits are more central to impression formation than others has been a classic tenet of research on person perception showcased in social psychology textbooks as a fundamental insight into the way people understand others. In a classic study, Asch (1946) exposed participants to a list of traits. Some participants received a list of traits which included the terms *intelligent, skillful, industrious, warm, determined, practical, and cautious*. Others received a modified list, identical to the first list with the exception that the word *cold* now replaced the word *warm*. Participants' were asked to write a sketch of the target depicted by those traits, and to respond to a checklist of 18 additional positive traits (e.g., *generous, wise, happy, reliable, honest*) by stating the extent to which each of those traits characterized the target. Participants who received the list including the term *warm* wrote more positive sketches, and tended to endorse the positive traits to a greater extent than those who received the list including the term *cold*. To Asch (1952, p. 210), ''These results [justified] the conclusion that a change of one quality produces a fundamental change in the entire impression.''

A mainstay of Asch's analysis was that human traits vary in their centrality to impression formation. Specifically, when the terms ''warm'' and ''cold'' were replaced in the same list by ''polite'' and ''blunt,'' the differences were markedly weaker. Also, in Asch's (1946) studies not all positive qualities were affected by a shift from the ''warm'' to the ''cold'' characterization: ''Reliability,'' ''honesty,'' and ''seriousness'' were completely unaffected by this distinction. The differential impact of the ''warm/cold'' factor on various traits was replicated in Harold Kelley's (1950, p. 435) classic experiment, wherein this variable ''produced differential effects from one rating scale to another.''

A trait's centrality can be understood according to its breadth of implication (Anderson, 1981, p. 217). Asch noted that some traits ''contain implications for wide regions of the person'' (Asch, 1946, p. 277). In addition to the breadth of implication, central traits were thought by Asch to be more likely to exert an effect on the impression of other traits than to be influenced by the presence of those other traits. For example, warm was found to be more likely to influence the perception of politeness than for politeness to influence the perception of warmth. When discussing the reason for the centrality of the dimension of warmth in comparison to the dimension of politeness, Asch (1946, p. 277) notes that ''there is a particular direction of forces.'' Yet, for Asch, the direction of forces emanated from warmth outward. In other words, centrality rested in the traits themselves.

Wishner (1960) went on to clarify the conditions under which some but not all traits would be impacted, and some but not all traits would be central. He demonstrated that ''centrality'' is determined by the correlation between trait items in the stimulus list, and those in the checklist. That is, a high correlation between a trait in the stimulus list and a large number of traits in the checklist renders the trait ''central.'' Similarly, a given checklist trait would be impacted by a stimulus trait to the extent that it was generally correlated with it. Accordingly, Wishner (1960) demonstrated that a seemingly peripheral dimension such as ''blunt–polite'' could be made central by altering the checklist to include traits with which this dimension was correlated.

In a similar vein, Rosenberg, Nelson, and Vivekananthan (1968) investigated the multidimensional structure of personality impressions, based on an analysis of co-occurrences between trait terms. Such analysis yielded evidence for two main dimensions of impressions, namely those of social desirability and of intellectual desirability. Rosenberg et al. (1968) implied that a trait term would affect only trait judgments on one dimension, and not on the alternative dimension. An experiment by Zanna and Hamilton (1972) supported this assertion. Participants in the experiment were presented with a target's socially-desirable or undesirable attribute (*warm* vs. *cold*), both in the context of four positive intellectually desirable traits (*industrious*, *skillful*, *determined*, and *practical)*. Other participants were presented with a target's intellectually-desirable or undesirable attribute (*industrious* vs. *lazy*), both in the context of four socially desirable traits (*warm*, *sociable*, *good natured*, and *humorous)*. Participants then rated the probability that the target (Bob) also possessed further traits, half of which fell on the socially-desirable dimension, and the remaining half on the intellectually-desirable dimension.

DOI: 10.1002/ejsp

In support of the Rosenberg et al. (1968) analysis, Zanna and Hamilton (1972) found that the warm–cold variation only impacted traits on the socially-desirable dimension, and not on the intellectually-desirable dimension, whereas the industrious–lazy variation only impacted traits on the intellectually-desirable but not on the socially-desirable dimension. Furthermore, individual differences in the perceived centrality of social and intellectual traits (e.g., honesty and intelligence) were shown to predict their effect on impression formation (Sedikides & Skowronski, 1993). These latter findings are consistent with Zanna and Hamilton's (1972, p. 354) insight regarding ''the importance of considering different attribute dimensions in understanding patterns of trait inference.''

A GENERAL EPISTEMIC PERSPECTIVE ON TRAIT CENTRALITY

The follow-up studies to Asch's (1946) original work (i.e., Rosenberg et al., 1968; Wishner, 1960; Zanna & Hamilton, 1972) made an important contribution by demonstrating that the centrality of a trait is relative to relations among trait terms (see also Anderson, 1981; Asch & Zukier, 1984; Hamilton, Katz, & Leirer, 1980; Higgins, Rholes, & Jones, 1977). This work provides evidence for an understanding of trait centrality according to the breadth of influence a trait possesses. Nonetheless, in the context of trait centrality research, the structure of the relation between various traits, and sources of variation in trait centrality, have yet to be specified. Although Asch (1946) suggested that the direction of the traits was important, and both Asch (1946) and Anderson (1981) discuss a trait's centrality as arising from the use of implicational inferences, the follow-up research (i.e., Rosenberg et al., 1968; Wishner, 1960; Zanna & Hamilton, 1972) assumed a bidirectional approach that is divorced from implicational reasoning. In addition, Asch (1946) assumed that centrality was located in some traits but not others, whereas our conceptualization suggests that different people and different situations may give rise to varying degrees of centrality for traits. We propose that a general epistemic approach to impression formation (Kruglanski, 1989) may help to clarify these issues.

Essentially, a lay epistemic account of the trait centrality phenomenon holds that information that a target possesses a given trait may function as *evidence* for further judgments, on condition that the particular trait is *subjectively relevant* to those judgments. According to Lay Epistemic Theory, properties of concepts (e.g., traits or social categories) are stored in the form of ''if... then'' linkages in which the knower happens to believe. Each such linkage may be thought of as a *major premise* in a syllogism. The information that a target possesses the trait may be thought of as the *minor premise*. A minor premise provides evidence that the antecedent term of the major premise is instantiated or applies in the given case. For instance, the combination of a major premise which states ''if someone is warm, then s/he is friendly'' and a minor premise which states ''Joe is warm'' would result in judgment that ''Joe is friendly.'' This would render the evidence that someone possesses the characteristic of ''warmth'' more relevant to friendliness ratings than ''studiousness'' that may not be connected in the perceiver's mind to friendliness. On the other hand, a person may believe more that ''if someone is studious (vs. warm), then s/he will do well in school.'' This would render ''studiousness'' versus ''warmth'' more relevant to predictions of academic success, assuming that in the perceiver's mind ''warmth'' does not imply success.

The notion of ''if... then'' rules differs from that of mere correlation between trait terms central to Wishner's (1960) analysis. Correlation (co-occurrence or co-variation) implies a conjoint presence and absence of two categories, whereby if trait X is high then trait Y is high as well, and if X is low than Y is low as well. Thus, the correlational account holds that X is both a sufficient and a necessary condition of Y. In contrast, the notion of ''if... then'' rules assigns X the role of a sufficient condition only. That is, ''if X then Y'' does not imply that ''if not X then not Y,'' for Y may have determinants other than X.

Further, by its very definition, the term *subjective relevance* connotes the possibility of variation across persons, neglected in previous analyses. The major premise (i.e., ''if... then'' rule) that lends subjective relevance to a given informational category (X) with regard to a given inference (Y) may be believed to a greater or lesser extent by different individuals, representing subjective rule strength. For example, some individuals may subscribe strongly to a lay inference rule that links warmth with friendliness, while others may subscribe to this inference rule less strongly, if at all. Under these conditions, someone described as ''warm'' would be particularly likely to be judged as ''friendly,'' more so than would have been the case if the conditional linkage between these two categories were weaker. Similarly, some individuals may subscribe strongly to a lay inference rule linking speed with athletic ability, whereas others' individuals may not (see also Sedikides & Skowronski, 1993). In short, we suggest that the effects of a given trait on the inference of another trait should be determined by the presence of lay inference rules inter-linking those traits in an ''if... then'' fashion.

DOI: 10.1002/ejsp

When the task of processing the evidentiary information is difficult, individuals might not be able to perceive its relevance as readily as in circumstances in which the task is easy (see Chun, Spiegel, & Kruglanski, 2002; Erb, Kruglanski, Chun, Pierro, Mannetti, & Spiegel, 2003). Therefore, the influence of the inference rules may also depend on the difficulty of gleaning the (subjectively) relevant evidence from the environment. For instance, if the trait information is difficult to discern, its ''signal to noise'' ratio being low, participants might not fully appreciate the relevance of the warm–cold information to their friendliness or likeability judgments, whereas if it was easy to discern, the warm–cold information might exert a more powerful influence. However, the difficulty of information processing may be successfully overcome if the knower has sufficient cognitive capacity to cope with the cognitive task at hand.

OVERVIEW OF PRESENT RESEARCH

To the extent that the foregoing account of the trait centrality phenomenon is valid, the following five hypotheses are implied. First, trait relations can be stored as unidirectional rules rather than bidirectional correlations (''if X then Y,'' but not necessarily ''if Y then X''). Second, the strength of individuals' lay inference rules should mediate the effect of a given trait on judgment regarding a target person. Third, situations that induce rule strength should appropriately affect the trait centrality phenomenon. Fourth, the use of a lay inference rule should depend on the difficulty of gleaning the appropriate evidence from the environment. Fifth, the difficulty of gleaning the evidence should impact those who are under cognitive load to a greater extent than those who have available their full processing capacity.

To explore these notions, we carried out five experimental studies. Study 1 examined the asymmetry of trait pair relations. Study 2 tested whether the strength of individuals' lay inference rules mediates the effect of a trait on target impressions. Study 3 manipulated rule strength experimentally and examined the effect of this manipulation on trait effects on impressions. Study 4 investigated whether difficult to detect evidence is less likely to activate an inference rule and to impact impression formation. Finally, Study 5 manipulated cognitive load to determine whether the difficulty of detection affects individuals with limited processing resources to a greater extent than it does individuals with ample processing resources.

STUDY 1: UNIDIRECTIONAL IMPLICATIONS ACROSS 50 TRAIT PAIRS

Study 1 was designed to assess the asymmetry of relations between trait pairs, across 50 pairs. In order to assess the asymmetry of the trait pairs, we assessed the perception of (1) the strength of the rule ''if trait A, then trait B'' and (2) the strength of the rule ''if trait B, then trait A.'' If trait relations can be stored as directional ''if. . . then'' rules, then we should find that one of the implicational rules (e.g., if A, then B) can differ significantly from the converse rule (e.g., if B, then A). However, if trait relations are stored exclusively as correlations, then we should find that A implies B to the same extent that B implies A.

Method

Participants

One hundred twenty-nine University of Maryland undergraduates (83 females and 46 males) participated in the experiment in exchange for course credit.

Procedure

Participants completed two questionnaires as part of the study on ''Social Judgment.'' In order to assess the strength of the ''if. . . then'' rules, the first questionnaire asked participants about their agreement with the statement ''if someone is A, then s/he is B.'' A and B were replaced in the sentence with each of the 50 trait pairs in Table 1. Participants responded on a nine-point Likert scale ranging from 1 (strongly disagree) to 9 (strongly agree). The second questionnaire assessed the

DOI: 10.1002/ejsp

Table 1. Trait pairs and differences between trait orders

	Paired Differences				
	Mean	*SD*	*t*	df	*p*
Sincere–friendly	0.16	1.97	.90	125	.368
Wise–mature	0.52	1.87	3.15	125	.002
Reliable–considerate	0.27	2.19	1.39	124	.167
Helpful–courteous	0.04	2.28	.20	124	.845
Forgiving–polite	0.24	2.11	1.27	124	.205
Bright–alert	0.29	1.82	1.77	124	.079
Talented–spirited	0.51	1.86	3.04	123	.004
Neat–punctual	0.48	2.08	2.60	124	.011
Sympathetic–warm	0.28	1.98	1.58	125	.118
Imaginative–clever	2.29	2.52	10.15	124	.000
Purposeful–eager	0.54	2.05	2.93	124	.004
Decisive–tidy	0.06	1.83	.389	125	.698
Outgoing–easygoing	0.26	2.08	1.39	121	.166
Moral–strong-minded	1.14	2.47	5.08	121	.000
Objective–direct	0.12	1.85	.73	123	.469
Congenial–attentive	0.07	1.52	.54	121	.593
Vigorous–competent	0.13	1.54	.98	125	.328
Inexperienced–naive	0.08	2.12	.42	123	.673
Harsh–stern	0.12	1.89	.71	124	.478
Timid–passive	0.15	2.06	.82	124	.412
Submissive–tense	0.17	1.63	1.16	123	.250
Overconfident–envious	0.54	1.69	3.57	122	.001
Unhappy–sarcastic	0.38	2.09	2.02	124	.046
Lazy–sloppy	0.18	2.34	.85	122	.395
Oversensitive–moody	0.76	2.46	3.43	123	.001
Vain–bossy	0.29	1.86	1.70	118	.091
Nosy–annoying	2.13	2.55	9.34	124	.000
Greedy–untrustworthy	1.54	2.34	7.25	121	.000
Heartless–selfish	1.07	2.38	5.04	124	.000
Conceited–snobbish	0.60	1.76	3.79	121	.000
Cruel–unkind	0.86	2.36	4.03	121	.000
Deceptive–maladjusted	0.33	1.93	1.93	123	.056
Scornful–self-confident	0.01	1.87	.05	122	.962
Crude–impolite	0.72	2.69	2.97	123	.004
Lazy–weak	0.40	1.98	2.24	124	.027
Fearless–hopeful	0.67	2.38	3.14	122	.002
Resentful–grouchy	0.43	2.33	2.06	124	.042
Entertaining–outspoken	0.25	2.36	1.16	123	.248
Stubborn–individualistic	0.32	2.00	1.78	123	.078
Creative–smart	0.58	2.22	2.91	123	.004
Impulsive–blunt	0.06	1.66	.38	120	.702
Overcritical–neglectful	0.03	1.79	.176	123	.861
Funny–charming	0.44	2.21	2.21	121	.029
Original–independent	0.34	1.87	2.03	122	.045
Intelligent–conscientious	0.11	1.85	.63	122	.528
Cooperative–polite	0.34	1.71	2.23	121	.028
Ethical–tolerant	0.49	1.81	3.03	123	.003
Generous–appreciative	0.32	2.00	1.79	123	.623
Versatile–warm	0.08	1.73	.49	123	.623
Independent–intellectual	0.06	1.51	.45	123	.655

Note: The trait relations are presented such that the first trait implies the second trait to a greater extent than the second trait implies the first.

DOI: 10.1002/ejsp

reverse implicational direction, and asked participants to rate their agreement with the statement ''if someone is B, then s/he is A,'' using the same nine-point rating scale. We counterbalanced the order of these questionnaires across participants.

The trait pairs we used were drawn from Anderson's (1968) list of traits, and based on the intuition that the traits within each pair would be perceived as somewhat related to one another. However, our main interest was to assess whether such related traits are invariably linked in a bidirectional manner, or if instead, they may be cognitively linked unidirectionally according to ''if... then'' implications that a given trait may have for another but not *vice versa*.

Results and Discussion

In order to examine asymmetry in trait pair relations, we first calculated within-trait pair association difference scores by subtracting the perceived strength of association between trait A and trait B from the perceived strength of association between trait B and trait A. A one-sample *t*-test revealed that the total of difference scores, with all negative difference scores reversed to positive scores, was significantly greater than 0, at $t(125) = 12.51$, $p < .001$, providing strong evidence for overall asymmetry in the relation of trait pair representation. Separate, pair wise *t*-tests across the 50 traits showed that 23 of the trait pairs showed differences as a function of order at the $p < .05$ level. For four additional trait pairs, trait pair order made a marginally significant difference, at $p < .10$. The results of these *t*-tests for all trait pairs are depicted in Table 1.

Overall, the present findings are consistent with the assertion that trait pair relations can be (and often are) cognitively connected in a unidirectional rather than bidirectional way. The difference found between the implication ''If A... then B'' and ''If B... then A'' seems to be inconsistent with, and raises problems for, a correlational account (Wishner, 1960). This does not mean that all or most of the myriad possible trait relations are asymmetrical. Nonetheless, because many of the trait relations in this study did turn out to be asymmetrical, the possibility of unidirectional rules was corroborated by our data. Moreover, the present framework can account for those pairs which are symmetrical, that is, involve bidirectional implications. Specifically, symmetrical trait relations occur when the knower believes that A implies B to the same extent that B implies A. In contrast, looking merely at the correlations between traits (Wishner, 1960) not only ignores asymmetrical trait relations, but fails to specify the epistemic mechanism of inferring ones' impressions of a target from trait information that serves as evidence in the ''if... then'' syllogistic manner described earlier. This aspect of our analysis is explored more specifically below.

In Study 1, we found a significant difference between the endorsement of the rule ''if A... then B'' and the rule ''if B... then A'' across 50 trait pairs. An alternative interpretation of this finding could be that each of the measurements of the two implicational rules merely measures the correlation between A and B. In this sense, the difference between the two rules could reflect low test–retest reliability in the measurement of the correlation using this method. Such an explanation essentially argues that the reason we found the differences in Study 1 was not because traits are stored as unidirectional implications, but instead because the measurement procedure used in Study 1 was unreliable. If this is the case, then the measurement of the trait relations using this procedure should not be predictive of perceivers judgments based on the presence of warmth/coldness information. Study 2 was designed to test whether measuring the extent to which trait A implies trait B (in this manner) mediates the relationship between the presence of trait A and judgments about trait B.

STUDY 2: RULE STRENGTH EFFECTS ON IMPRESSION FORMATION

Given that trait relations can be represented in the form of unidirectional rules (Study 1), we expect that the *strength* of the stored rule should influence the extent to which trait evidence will lead to its implied judgment. Accordingly, the present experiment tested whether the strength of individuals' inference rules mediate the effect of a ''central'' trait on judgment regarding a target person. Specifically, we assessed the strength of participants' lay inference rules, conditionally relating warmth to friendliness and speed to athletic talent, before evaluating their interest in having a target person as a friend, or recruiting him to their basketball team. It was predicted that the presence of an inference rule (i.e., a major premise to which an individual may subscribe) mediates the effect of trait information on corresponding judgments.

DOI: 10.1002/ejsp

Method

Participants

Seventy-seven University of Chicago undergraduates (41 females and 36 males) participated in the experiment in return for monetary compensation.

Procedure

This study employed a 2: Trait (warm vs. fast) × 2: Role (friend vs. basketball player) between-subject design. Half the participants read about a target person supposedly interested in being their friend, whereas the remaining participants read about a target person interested in joining their basketball team. Participants were presented with a list of traits that this target person possessed. The trait list included the adjectives: Intelligent, skillful, industrious, warm (vs. fast), practical, and cautious.

Next, for each trait on the trait list, participants in the friendship condition rated the *extent* to which the trait indicated that a person is a good friend (i.e., ''given that someone is X, would you assume that s/he is friendly?''), whereas participants in the basketball-player condition rated the extent to which each trait indicated that a person is a good basketball player (i.e., given that someone is X, would you assume that s/he is a good basketball player?''). All ratings were made on seven-point scales (1 = cannot assume, 7 = can assume for sure). Although we were only interested in the perceived diagnostic value of ''warm'' and ''fast,'' in order to conceal the purpose of the study, we repeated this item for each trait on the trait list. Next, participants were asked to evaluate the extent to which they would like to have the target person as their friend or as a player in their basketball team, depending on experimental condition (on a seven-point scale, ranging from 1 = not at all, to 7 = extremely).

Results and Discussion

An ANOVA of target person's evaluations yielded a trait × role interaction, $F(1, 73) = 9.01$, $p < .01$, indicating that ratings of interest in friendship were higher for the warm target ($M = 4.79$) than for the fast target ($M = 4.10$), $t(37) = 2.52$, $p = .01$, whereas ratings of interest in having the person join their basketball team were higher for the fast target ($M = 4.55$) than for the warm target ($M = 3.56$), $t(36) = 2.01$, $p = .05$. No other effects emerged in this analysis.

Next, a series of regression analyses tested whether the observed effect on evaluations is mediated by the strength of lay inference rules, relating warmth and friendship, and athletic skills and fastness. As shown in Figure 1, beginning with the evaluation of basketball potential, this analysis found that in itself, the central trait adjective (fast vs. warm) directly increased interest in hiring the target person ($\beta = .38$, $p = .01$). However, indirectly, fastness (vs. warmth) was rated as a better predictor of basketball talent ($\beta = .71$, $p < .01$), which in turn increased interest in having the target person on their basketball team ($\beta = .64$, $p < .01$). Of greatest interest, controlling for the predictive diagnosticity of fastness, the effect of central trait on interest in having the person on their basketball team diminished ($\beta = .11$, ns). The Sobel test statistic found that diagnosticity significantly mediated the interest in having the person on their basketball team ($z = 3.37$, $p < .001$).

The same analysis, conducted on the friendliness ratings, yielded a similar pattern. In itself, the central trait adjective (warm vs. fast) directly increased interest in having the target person as a friend ($\beta = .54$, $p < .01$). However, indirectly, warmth (vs. fastness) was rated as a better predictor of friendliness ($\beta = .84$, $p < .01$), which in turn increased the interest in friendship ($\beta = .61$, $p < .01$). Controlling for the diagnosticity of warmth, the effect of this trait on interest in friendship diminished ($\beta = .08$, ns). In this case as well, the Sobel statistic attested that diagnosticity significantly mediated the interest in friendship ($z = 2.06$, $p < .05$).

The present results demonstrate the dependence of the trait centrality effect on the strength of participants' lay inference rules. Specifically, if the impact of a stimulus trait on other aspects of person perception (e.g., friendliness or basketball ability) depends on the degree to which individuals believe in an ''if. . . then'' relation between a given trait and a given perception, one should expect to eliminate the effect of trait centrality by controlling for participants' subscription to the inference rule. Accordingly, the degree to which participants believed that warmth (vs. fastness) indicates friendliness mediated the degree to which they viewed themselves likely to adopt a warm person as their friend. Similarly,

DOI: 10.1002/ejsp

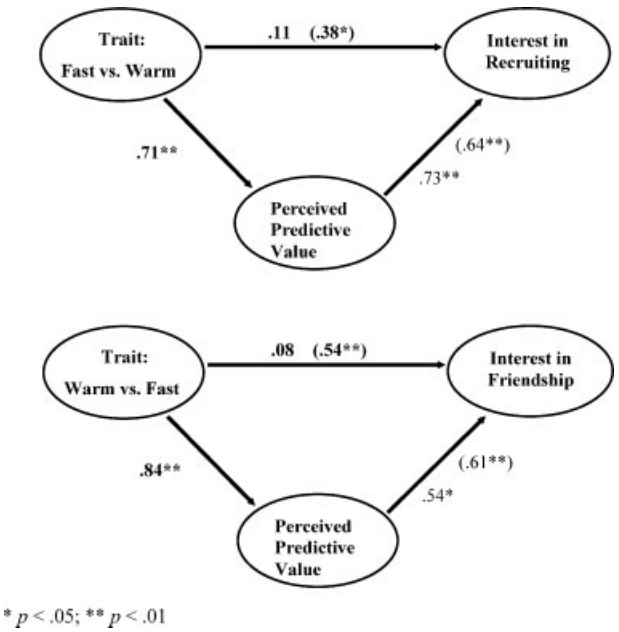

Figure 1. Mediation by lay theories (study 2)

the degree to which participants believed that fastness (vs. warmth) implies athleticism mediated the degree to which they were likely to recruit a fast person to their basketball team. These findings underscore the fact that the phenomenon of trait centrality does not have to do with specific traits as such but rather with the degree to which the traits are seen to constitute *compelling evidence* for given impressions.

In addition, the results from Study 2 provide additional support for our interpretation of the results in Study 1. In Study 1, we found a significant difference between ratings of ''if A. . . then B'' as compared to ''if B. . . then A'' across 50 trait pairs. We suggested that this finding could be interpreted as evidence that trait relations are stored as unidirectional implications. An alternative interpretation of these findings could suggest that the significant difference between the ratings merely indicates a lack of test–retest reliability in the assessment of stored associations. Because the measurement of unidirectional rules in Study 2 was a reliable mediator of the use of given information to make a subsequent judgment, we can be confident that the results from Study 1 were not the result of unreliable measurement of the trait relations. If our procedure for measuring the trait relations was as poor as would be needed to produce significant differences in the ratings over only a few moments, then such ratings should not be able to reliably mediate the trait centrality effect, as found in Study 2.

One limitation of Study 2 is that the measurement of the strength of lay inference rules occurred between the presentation of the information about the target person and the judgments about the target person. To conceal the nature of this measurement, we included in the questionnaire presented to participants all items from the trait checklist. However, this might not have fully removed demands towards consistency. To address these issues, Study 3 systematically manipulated participants' strength of lay inference rules and examined the effect of such manipulation on participants' impressions.

STUDY 3: MANIPULATING THE STRENGTH OF LAY INFERENCE RULES

Study 2 suggests that individuals vary in the strength of the ''if. . . then'' rule that in their mind connect two trait items in an implicational relation. Such differences may stem from the history of association between these terms, which may well

DOI: 10.1002/ejsp

differ substantially across individuals. In the present study, we attempted to create such history experimentally. Study 3 systematically manipulated participants' strength of lay inference rules and examined the effect of such manipulation on participants' impressions. Specifically, we repeatedly paired the terms ''warm'' and ''like'' and ''cold'' and ''dislike'' in the experimental condition and observed the degree to which the ''warm/cold'' trait in the stimulus list affects the inference that the target is ''likeable/unlikable.'' In a control condition, participants went immediately to the judgment task, without any previous task. This allows us to compare individuals with a recent history of association between warmth and likeability to individuals with no recent history of association between these dimensions. We expect that activating the association between warmth and likeability will increase the trait centrality effect.

In order to test for the possibility that the experimental condition simply activates the construct of warmth, or simply activates the concept of likeability (rather than the association between them), we also included a condition in which only the warm–cold dimension is primed, and a condition in which only the likeability dimension is primed. Specifically, we repeatedly presented the terms ''warm'' and ''cold'' in the warmth condition and repeatedly presented the terms ''like'' and dislike'' in the likeability condition.

After the manipulation of the accessibility of the warmth–likeability association, we then measured the degree to which the ''warm/cold'' trait in the stimulus list affects the inference that the target is ''likeable/unlikable.''

Method

Participants

One-hundred-five University of Maryland undergraduates (77 females and 28 males) participated in the experiment in fulfillment of course requirements.

Procedure

Participants were randomly assigned to the 4: Prime (warmth–likeability vs. warmth vs. likeability vs. no prime) $\times$ 2: Trait (warm vs. cold) conditions. The first part of the experiment manipulated the association strength between the warm/cold and likeable/dislikeable terms intended to vary the strength of the ''if warm then likeable'' (''if cold then dislikeable'') rules. That is, participants in the *warmth–likeable condition* responded to eight questions to which the answers were registered on a scale anchored at the ends with the terms ''warm'' and ''cold'' (e.g., What is the weather today? What is the ideal temperature for serving coffee?). Interspersed among these items were eight questions to which the answers were recorded on a scale anchored with ''like'' and ''dislike'' (e.g., What is you opinion of rock music? What is your opinion of George W. Bush?). The ''warm'' and the ''like'' anchors were placed at the same ends of the scale, as were the opposite anchors, ''cold'' and ''dislike.'' *In the warmth condition*, participants responded to 16 questions to which the answers were recorded on a scale anchored with the terms ''warm'' and ''cold'' (e.g., What is the weather today? What is the ideal temperature for serving coffee?). In the *likeability condition*, participants responded to 16 questions to which the answers were recorded on a scale anchored with the terms ''like'' and ''dislike'' (e.g., What is you opinion of rock music? What is your opinion of George W. Bush?). All ratings were recorded on a seven-point scale. Participants in the *no prime condition* did not respond to any questions and commenced directly with the second part of the study.

The second part of the experiment required participants to respond to another, allegedly unrelated, survey said to investigate impression formation. The first part of the experimental survey introduced the impression formation task. Participants were asked to imagine that ''Mike is interested in being your friend.'' Next, participants read about Mike who, ''*according to a reliable source is: Intelligent, skillful, industrious, cold (*or *warm), practical, and cautious*.'' Next, target evaluation scores were obtained by asking participants to rate their readiness to accept Mike as their friend. At that point, we assessed the effect of the manipulated traits on evaluation of friendship related characteristics by asking participants to rate the extent to which they believed Mike to be (1) insightful, (2) in possession of good social skills, (3) invested in his friendships, (4) a good practical advice giver, and someone who (5) worries about his friends. All ratings were made on seven-point scales.

DOI: 10.1002/ejsp

Results and Discussion

Friendship Ratings

A 4: Prime (warmth–likeability vs. warmth vs. likeability vs. no prime) × 2: Trait (warm vs. cold) ANOVA performed on participants' interest ratings in having the target person as a friend yielded a main effect for the trait variable, $F(1, 97) = 46.60$, $p < .001$, indicating that participants preferred the warm over the cold target as their friend (Ms = 5.33 and 4.00, respectively). As shown in Figure 2, this analysis also revealed the predicted prime × trait interaction, $F(1, 97) = 2.40$, $p = .07$. Next, we computed planned comparisons investigating the difference between the friendship ratings between the warm target and the cold target.

Replicating the general trait centrality effect, even in the no-prime condition participants expressed greater interest in establishing friendship with the warm rather than the cold target (Ms = 5.00 and 4.00, respectively) $F(1, 97) = 7.74$, $p < .01$. Similarly, in the warmth condition participants' expressed greater interest in establishing friendship with the warm rather than the cold target (Ms = 5.08 and 4.17, respectively) $F(1, 97) = 5.19$, $p < .05$, as was the case in the likeability condition (Ms = 5.67 and 4.46, respectively) $F(1, 97) = 9.10$, $p < .01$. Participants in the warmth–likeable condition were even more interested in striking a friendship with the warm, as compared to the cold target (Ms = 5.58 and 3.36, respectively) $F(1, 97) = 35.16$, $p < .001$.

Impressions of Target

Next, we collapsed the ratings of the target's friendship-related traits ($\alpha = .69$). A 2: Prime: 4 (warmth-likeability vs. warmth vs. likeability vs. no prime) × trait (warm vs. cold) ANOVA carried out on this measure yielded a main effect for trait, $F(1, 97) = 43.08$, $p < .001$ (Ms = 5.09 and 4.09 for the 'warm' and 'cold' targets, respectively). Of greater interest, we obtained a marginally significant interaction between the prime × trait variables, $F(1, 97) = 2.42$, $p = .07$.

We then computed planned comparisons investigating the difference between the impressions between the warm target and the cold target in each of the prime conditions. In the no-prime condition participants reported a (marginally significant) more positive view of the target's traits when he was portrayed as warm rather than cold (Ms = 4.76 and 4.20, respectively) $F(1, 97) = 3.77$, $p = .06$. In the warmth prime condition participants' impressions of the warm target were significantly more positive than the cold target (Ms = 5.00 and 4.2, respectively) $F(1, 97) = 6.63$, $p < .05$. In the likeability condition participants' reported a preference for establishing friendship with the warm rather than the cold target (Ms = 5.13 and 4.15, respectively) $F(1, 97) = 9.95$, $p < .01$. In the warmth–likeable condition the warm versus cold targets were evaluated more extremely than in the low accessibility conditions (Ms = 5.47 and 3.81, respectively) $F(1, 97) = 29.31$, $p < .001$.

These results affirm that the ''if warm then likeable'' inference rule was generally available to our participants. However, enhancing its strength via our experimental manipulation increased participants' tendency to use it in judgments, resulting in significant interactions regarding the readiness to accept the target as a friend, and regarding the

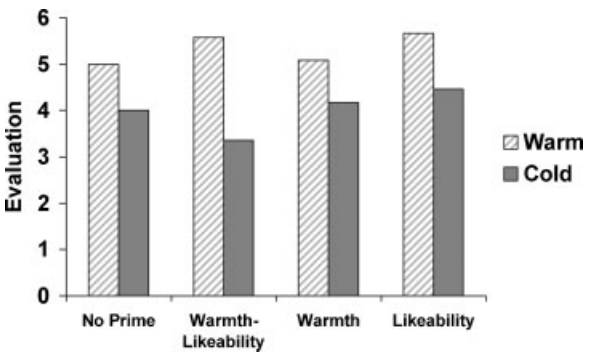

Figure 2. Evaluation of target as a function of association history and trait information (study 3)

DOI: 10.1002/ejsp

target's ratings on friendship-related characteristics. As expected, the association of warmth–likeability resulted in more extreme ratings of the target than in all the three comparison conditions. This suggests that our procedure for developing a history of association between warmth and likeability was not simply influencing the judgments based on mere activation of the warmth of likeability constructs. In fact, the ''no prime'' control condition was not significantly different from the warmth prime condition, nor was it different from the likeability prime condition.

STUDY 4: DIFFICULTY OF INFORMATION PROCESSING

Even when a lay inference rule is held with confidence and is momentarily accessible, it cannot be used unless its relevance to the judgmental task can be appreciated. Our next study investigated the effect of task demands on using trait information in impression formation. Recent work on judgment formation and persuasion (e.g., Kruglanski, Pierro, Mannetti, Erb, Chun, & Sleeth-Keppler, 2007) highlights the critical importance of processing difficulty in individuals' appropriate responsiveness to the information given. As with other information processing tasks, to the extent that the impression formation task is difficult and demanding, individuals may not be able to appreciate the relevance of given trait information to their target impressions. This possibility was explored in the present experiment by presenting the relevant trait information on an experimental survey that either used clear or blurry fonts. The blurry fonts were expected to increase the processing difficulty without affecting the content of the information, and we expected this to attenuate the effect of the trait content on the impressions formed.

Method

Participants

Eighty-one University of Maryland undergraduates (52 females and 29 males) participated in the experiment in fulfillment of course requirements.

Procedure

Participants were randomly assigned to a 2: Processing difficulty (clear vs. blurry fonts) $\times$ 2: Trait (warm vs. control) between-subject design. Participants read about Mike who is interested in being in their friend, and who is ''intelligent, skillful, industrious, cold (vs. warm), practical, and cautious.'' The task difficulty was manipulated via the font clarity. In the *high-difficulty* condition the font used for traits was embossed and hence required greater effort to read. In contrast, the font adopted in the *low-difficulty* condition was presented without any additional effect and hence clear. Pre-testing was conducted to be sure that the high-difficulty font was readable, so that participants were able to ascertain the meaning of the information that was present. Although the font clarity manipulation did make it more difficult to read, it did not make it impossible to read.

Participants were then asked to rate the extent to which they would like to be Mike's friend, and the extent to which they believed Mike possesses friendship-related characteristics, including (1) insightful, (2) in possession of good social skills, (3) invested in his friendships, (4) a good practical advice giver, and someone who (5) worries about his friends. All ratings were made on seven-point scales.

Results and Discussion

An ANOVA performed on participants' interest ratings in having the target person as a friend yielded a main effect for the trait variable, $F(1, 75) = 46.78$, $p < .001$, indicating that participants preferred the warm over the cold target as their friend. More importantly, we obtained the predicted difficulty $\times$ trait interaction, $F(1, 75) = 4.44$, $p < .05$. As shown in Figure 3, the effect of the trait variable was more pronounced when the survey font was clear (Ms = 5.52 and 3.15, for cold vs. warm,

DOI: 10.1002/ejsp

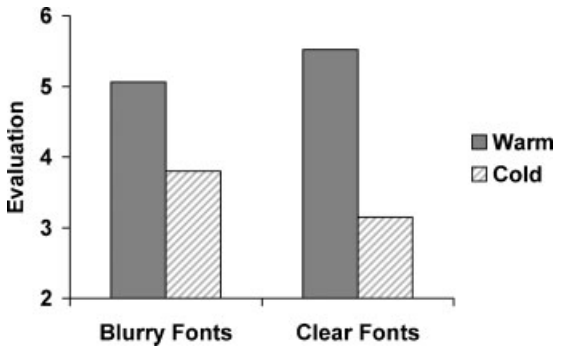

Figure 3. Evaluation as a function of difficulty of information processing and trait information (study 4)

$t(39) = 7.17$, $p < .001$), than when the font was blurry and difficult to read (Ms = 5.06 and 3.80, for warm vs. cold, $t(36) = 2.99$, $p < .01$).

We then analyzed the ratings of the target's friendship-related traits ($\alpha = .65$). An ANOVA yielded a main effect for trait, $F(1, 77) = 33.80$, $p < .001$ and the predicted difficulty × trait interaction, $F(1, 77) = 4.16$, $p < .05$. Again, even in the blurry font condition participants evaluated the warm versus cold target more positively (Ms = 4.49 and 3.83, for warm vs. cold, $t(38) = 2.66$, $p < .05$), but this effect was significantly larger in the clear font condition (Ms = 5.12 and 3.75, warm vs. cold, $t(39) = 5.66$, $p < .001$) than when it was difficult.

These findings suggest that people are less likely to rely on available evidence (i.e., trait information) when it is difficult to process, presumably because they do not appreciate its potential relevance to the judgment. It also follows that when the trait information is difficult to process and one's available cognitive resources are low (e.g., under cognitive load), the effect of the trait information will diminish even more completely. In contrast, having ample cognitive resources should reduce the impact of the difficult to process information. Our next study explored these possibilities by replicating the blurry-font conditions under cognitive load (vs. no load).

STUDY 5: DIFFICULTY OF INFORMATION PROCESSING AND COGNITIVE RESOURCES

As noted above, processing resources may play a particular role when the judgmental task was difficult to carry out for some reason, e.g., due to length or complexity of the information, poor visibility, ambient noise, etc. (for discussion see Chun et al., 2002; Erb et al., 2003; Kruglanski, Fishbach, Erb, Pierro, & Mannetti, 2004; Kruglanski et al., 2007; Kruglanski & Thompson, 1999a,b). Accordingly, we predicted that when the task demands are high (e.g., when the fonts are blurry) limiting processing resources by the imposition of cognitive load would diminish the effect of the trait information as compared to the case where the processing resources are plentiful.

Method

Participants

Ninety-two University of Maryland undergraduates (70 females and 22 males) participated in the experiment in fulfillment of course requirements.

Procedure

Participants were randomly assigned to a 2: Cognitive load (high vs. low) × 2: Trait (warm vs. cold) between-subject design. After receiving the general experimental instructions, participants were told that the study forms part of a research

DOI: 10.1002/ejsp

project on the relationship between font types and readability, which justified the use of blurry fonts in all conditions. Participants in the *high cognitive-load* condition were further told that another interest of the study was to investigate how well people can perform two different tasks simultaneously. Before participants examined the information about Mike, they were asked to recite a nine-digit number silently to themselves until they were later asked to reproduce it (Chun & Kruglanski, 2006; Chun et al., 2002). At that point, participants were allowed to view the number for 25 seconds. Their counterparts in the low cognitive-load condition proceeded without these instructions. Participants were then asked to imagine someone that they trust tells them that a target person is "*intelligent, skillful, industrious, warm (vs. cold), determined, practical, and cautious*." Time allowed for reading the target person's traits was 25 seconds, which according to a pretest was sufficient for a thorough reading the traits. At that point, prior to responding to the dependent measures, participants in the high cognitive-load condition were asked to reproduce the nine-digit number they had been asked to recite. Therefore, none of the participants experienced high cognitive-load while responding to the dependent measures.

After reading the stimulus list, participants were provided with 10 traits (i.e., generous, wise, happy, good-natured, humorous, sociable, popular, humane, altruistic, and imaginative). These traits were shown to be reliably influenced by the central traits manipulation (i.e., warm vs. cold) in Asch's (1946) original experiments and therefore, they enabled us to (a) replicate the original findings and (b) test for moderation, using the exact same evaluative measures as Asch. We specifically asked participants to report the extent to which Mike is a person who has each trait on a nine-point scale (1 = not at all, 9 = extremely).

Results and Discussion

An ANOVA of the target's friendship-related traits yielded a main effect for trait, $F(1, 88) = 9.57$, $p < .01$, with more positive evaluation of the warm target than the cold one. More importantly, it yielded the predicted cognitive load × trait interaction, $F(1, 88) = 4.47$, $p < .05$. As shown in Figure 4, under no load participants evaluated the warm versus cold target more positively (Ms = 6.10 and 4.61, for warm vs. cold), $t(42) = 4.08$, $p < .001$; however, under cognitive load, the warm and cold targets were not significantly different from each other (Ms = 5.77 and 5.50, for warm vs. cold), $t(46) = .65$, ns. We conclude that when the task demands are high (as a result of unclear fonts) and people's processing resources are low, trait information will not affect the impression formed.

GENERAL DISCUSSION

The present epistemic analysis and results suggest that understanding the trait-centrality phenomenon and the variables that affect it may be advanced by considering it as a specific instance of a general judgmental process. Trait centrality varies as a function of individuals' lay inference rules, of momentary and chronic conditions affecting the strength of belief in such inferences (rule strength), and the ability to glean the relevant evidence instantiating the use of such rules.

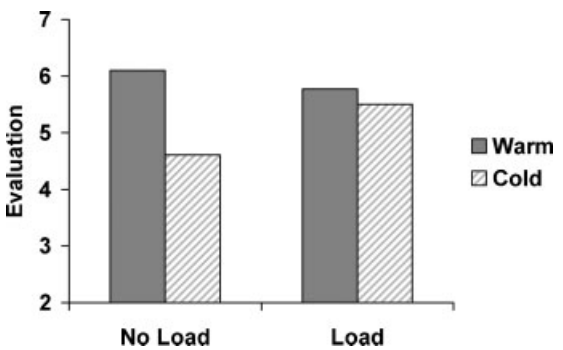

Figure 4. Evaluation as a function of processing resources (study 5)

DOI: 10.1002/ejsp

We assumed that a trait exhibits the property of centrality to the extent that it constitutes subjectively relevant evidence for a social judgment. Rather than some traits possessing a ''central content'' in the absolute sense of the term, trait ''centrality'' may depend on these traits' relation to judgmental dimensions. Contrary to Wishner's (1960) analysis of the trait centrality phenomenon, the present epistemic account of the trait centrality phenomenon suggests that trait centrality does not necessarily require that the trait and the judgment be correlated in the strict sense of this term. A correlation (or conjunctive presence-absence) implies that a trait is both a sufficient and a necessary condition for a judgment. According to the present analysis, however, sufficiency is enough, for the evidential function of traits should depend on their presumed ''if... then'' relation to particular judgments. Consistent with these notions, Study 1 demonstrated that a judgment Y may be inferred from trait X on basis of an ''if X then Y'' relation (a *sufficiency* condition), without assuming also an ''if Y then X'' relation (defining the *necessity* of X to Y).

Locating the phenomenon of trait centrality in the *evidential function* of information, grounded in the concept of subjective relevance (Kruglanski, Erb, Pierro, Mannetti, & Chun, 2006; Kruglanski et al., 2007) also suggests that the trait centrality effect may vary as function of a number of factors beyond the *general* relation across individuals between trait X and judgment Y. An important such factor is the specific presumed relation between the two, depending on the strength of individuals' lay inference rules. This suggests that trait centrality may vary as function of individual differences in the strength of the inference rule or the major premise (the ''if X then Y'' supposition).

Consistent with these notions, Study 2 showed that individuals' lay belief in inference rules, which attest to the perceived relevance that given traits have for given judgmental dimension, mediate trait centrality effects. Controlling for those inference rules eliminated the differences in impressions of the target as a function of the manipulated traits. Studies 3a and 3b demonstrated that situationally varying the strength of lay inferences rules determined the use of the information in judgments of social stimuli.

In order for a particular inference rule to be used, the antecedent term must be available and accessible in the informational array, leading to a judgment consistent with the consequent term. As such, Study 4 found that when processing the relevant antecedent term was made difficult, judgments consistent with the consequent term were less likely. Moreover, Study 5 found that the effect of task demands is particularly strong for those who have limited processing resources available while processing the information.

These results are consistent with the notion that information that a target possesses a given trait may function as evidence for inferences about other traits, on condition that the particular trait is subjectively relevant to those inferences. Trait centrality, then, is assumed to be the extent to which a specific trait implies other traits, and will influence judgment to the extent that the inference rule is believed and the evidence can be gleaned.

Another conceptual approach that has been used to understand impression formation has been the use of mathematical and simulation models, including connectionist models (e.g., Van Overwalle & Labiouse, 2004). Our implicational approach is consistent with this type of modeling to the extent that the linkages between informational nodes are unidirectional. That is, the strength of the linkage moving from node A to node B can have a different weight than the strength of the linkage moving from node B to node A. Indeed, it seems that a connectionist model in which a specific value is assigned to the ''if A... then B'' linkage that can be different from the ''if B... then A'' linkage would have greater predictive power.

Alternative Interpretations

Thus far, we have discussed the importance of our findings for an epistemic approach to human perception and judgment. The main alternative to such an approach is found in associationistic models. The data presented here is problematic for an associationistic account in at least two ways. First, associationistic models assume that traits are stored according to bidirectional linkages. Yet, we found that (a) the strength of believe in trait relations differs depending on the directional implication, and (b) the lay belief in directional trait relations mediates the trait centrality effect. These findings suggest that informational nodes may be connected by directional linkages (as is assumed in an inferential account) rather than bidirectional linkages (as would be assumed in an associationistic account). Second, associationistic models assume that stored associations would be readily used when the judgmental task was difficult or when cognitive resources are low (for a review, see Kruglanski & Orehek, 2007). In fact, associationistic models argue that stored associations are used more as processing demands increase (as compared to when processing demands are lower). Yet, making the information difficult to read and introducing cognitive load disrupted the use of the stored associations.

DOI: 10.1002/ejsp

Taken together, the findings from the six studies make alternative interpretations of the findings cumbersome. In Study 1, we found a significant difference between the endorsement of the rule ''if A... then B'' and the rule ''if B... then A'' across 50 trait pairs. It could be argued that each of the measurements of the two implicational rules merely measures the association between A and B. In this sense, the difference between the two rules could reflect low test–retest reliability in the measurement of the association using this method. However, this account would make it impossible to interpret the results from Study 2. In Study 2 it was found that the measurement of implicational rules (using the same question format as Study 1) mediated the trait centrality effect. If the reliability of measurement was so low in Study 1 as to undermine the importance of the difference in ratings according to the direction of the implication, then the measurement should not have been able to be a reliable predictor in Study 2.

Taken together, the results from these studies corroborate our epistemic approach, and raise significant problems for associationistic accounts. Because of this, the most plausible interpretation of our data seems to be an inferential account. Yet, it is impossible to directly measure the specific way in which information is stored in memory. Future research could profitably explore this topic using different methods to provide further evidence with respect to the two alternative approaches.

ACKNOWLEDGEMENTS

This research was supported by NSF grant 0314291/0313483 to Arie W. Kruglanski and the Korea Science and Engineering Foundation (KOSEF) grant funded by the Korea government (MEST)(NO. R01-2008-000-20802-0) to Woo Young Chun.

REFERENCES

Anderson, N. H. (1968). Likeableness of 555 personality-trait words. *Journal of Personality and Social Psychology*, *9*, 272–279.
Anderson, N. H. (1981). *Foundations of information integration theory.* New York: Academic Press.
Asch, S. (1946). Forming impressions of personality. *Journal of Abnormal and Social Psychology*, *41*, 258–290.
Asch, S. (1952). Knowledge of persons and groups. *Social psychology* (pp. 205–239). Englewood Cliffs, NJ US: Prentice-Hall, Inc.
Asch, S. E., & Zukier, H. (1984). Thinking about persons. *Journal of Personality and Social Psychology*, *46*, 1230–1240.
Chun, W., & Kruglanski, A. (2006). The role of task demands and processing resources in the use of base-rate and individuating information. *Journal of Personality and Social Psychology*, *91*, 205–217.
Chun, W. Y., Spiegel, S., & Kruglanski, A. W. (2002). Assimilative behavior identification can also be resource dependent: The unimodel perspective on personal-attribution phases. *Journal of Personality and Social Psychology*, *83*, 542–555.
Erb, H., Kruglanski, A., Chun, W., Pierro, A., Mannetti, L., & Spiegel, S. (2003). Searching for commonalities in human judgement: The parametric unimodel and its dual mode alternatives. *European review of social psychology* (Vol. 14, pp. 1–47–248). Hove England: Psychology Press/Taylor & Francis (UK).
Goldman, M., Cowles, M. D., & Florez, C. A. (1983). The halo effect of an initial impression upon speaker and audience. *Journal of Social Psychology*, *120*, 197–201.
Hamilton, D. L., Katz, L. B., & Leirer, V. O. (1980). Cognitive representation of personality impressions: Organizational processes in first impression formation. *Journal of Personality & Social Psychology*, *39*, 1050–1063.
Higgins, E. T., Rholes, W. S., & Jones, C. R. (1977). Category accessibility and impression formation. *Journal of Experimental Social Psychology*, *13*, 141–154.
Kelley, H. H. (1950). The warm-cold variable in first impressions of persons. *Journal of Personality*, *18*, 431–439.
Kruglanski, A. W. (1989). *Lay epistemic and human knowledge: Cognitive and motivational bases*. New York: Plenum.
Kruglanski, A., Pierro, A., Mannetti, L., Erb, H., & Chun, W. (2007). On the parameters of human judgment. *Advances in experimental social psychology* (Vol. 39, pp. 255–303). San Diego, CA US: Elsevier Academic Press.
Kruglanski, A., Erb, H., Pierro, A., Mannetti, L., & Chun, W. (2006). On Parametric continuities in the world of binary either ors. *Psychological Inquiry*, *17*, 153–165.
Kruglanski, A. W., Fishbach, A., Erb, H.-P., Pierro, A., & Mannetti, L. (2004). The parametric unimodel as a theory of persuasion. In G. Haddock, & G. Maio R (Eds.), *Contemporary perspectives on the psychology of attitudes* (pp. 399–422). New York: Psychology Press.
Kruglanski, A. W., & Orehek, E. (2007). Partitioning the domain of social inference: Dual mode and systems models and their alternatives. *Annual Review of Psychology*, *58*, 291–316.

DOI: 10.1002/ejsp

Kruglanski, A. W., & Thompson, E. P. (1999a). The illusory second mode or, the cue is the message. *Psychological Inquiry*, *10*, 182–193.

Kruglanski, A. W., & Thompson, E. P. (1999b). Persuasion by a single route: A view from the unimodel. *Psychological Inquiry*, *10*, 83–109.

Rosenberg, S., Nelson, C., & Vivekananthan, P. (1968). A multidimensional approach to the structure of personality impressions. *Journal of Personality and Social Psychology*, *9*, 283–294.

Sedikides, C., & Skowronski, J. J. (1993). The self in impression formation: Trait centrality and social perception. *Journal of Experimental Social Psychology*, *29*, 347–357.

Van Overwalle, F., & Labiouse, C. (2004). A recurrent connectionist model of person impression formation. *Personality and Social Psychology Review*, *8*, 28–61.

Watkins, M. J., & Peynircioglu, Z. F. (1984). Determining perceived meaning during impression formation: Another look at the meaning change hypothesis. *Journal of Personality and Social Psychology*, *46*, 1005–1016.

Wishner, J. (1960). Reanalysis of ''impressions of personality''. *Psychological Review*, *37*, 96–112.

Wyer, R. S., Bodenhausen, G. V., & Srull, T. K. (1984). The cognitive representation of persons and groups and its effect on recall and recognition memory. *Journal of Experimental Social Psychology*, *20*, 445–469.

Zanna, M. P., & Hamilton, D. L. (1972). Attribute dimensions and patterns of trait inferences. *Psychonomic Science*, *27*, 353–354.

DOI: 10.1002/ejsp