COGNITION AND EMOTION
2008, 22 (4), 553–594

Psychology Press
Taylor & Francis Group

The associations in our heads belong to us: Searching for attitudes and knowledge in implicit evaluation

Brian A. Nosek and Jeffrey J. Hansen
University of Virginia, Charlottesville, VA, USA

Explicitly, humans can distinguish their own attitudes from evaluations possessed by others. Implicitly, the viability of a distinction between attitudes and evaluative knowledge is less clear. We investigated relations between explicit attitudes, cultural knowledge and the Implicit Association Test (IAT). In seven studies (158 samples, $N = 107{,}709$), the IAT was reliably and variably related to explicit attitudes, and explicit attitudes accounted for the relationship between the IAT and cultural knowledge. We suggest that people do not have introspective access to the associations formed via experience in a culture. Ownership of mental associations is established by presence in mind and influence on thinking, feeling and doing. Regardless of origin, associations are influential depending on their availability, accessibility, salience, and applicability. Distinguishing associations as "not mine" is a self-regulatory act and contributes to the distinction between explicit evaluation, where such acts are routine, and implicit evaluation, where they are not.

There is little debate over the source of intentional thoughts and actions. Intended acts are products of the self via psychological mechanisms like goals, attitudes, and beliefs. But what of thoughts and actions that are unintended? In one sense, it is a tautology to say that the activities of a brain and mind belong to the individual taking residence in that brain. In another sense, the question highlights potential confusion over the proper attribution when intentions and actions are dissociated. A committed egalitarian may find her- or himself prejudging a Hispanic job applicant as unqualified. The thoughts may be distressing as inconsistent with her or his self-concept, ideology, and honest attitudes toward Hispanics. What then is the source of the prejudgement—to whom does it belong?

Correspondence should be addressed to: Brian Nosek, Box 400400, 102 Gilmer Hall, Charlottesville, VA 22904–4400, USA. E-mail: nosek@virginia.edu

This research was supported by a grant from the National Institute of Mental Health (R01 MH68447).

Thanks to Mahzarin Banaji, Jerry Clore, Tony Greenwald, Fred Smyth, Bethany Teachman, Tim Wilson and Mark Zanna for helpful comments. Supplementary materials can be found at http://briannosek.com/.

www.psypress.com/cogemotion DOI: 10.1080/02699930701438186

Automaticity has taken hold in psychology with a broad range of mental life now understood to proceed without the encumbrances of awareness, intention, and control (Bargh, 1996; Wegner & Bargh, 1998). The notion of automaticity affirms that the "owner" of actions is the individual, but simultaneously points out that the action need not be a product of conscious will or experienced as coming from the self (Wegner, 2002). The conscious, "experienced" self may be just another observer of the daily activities of the body it inhabits, perhaps having only slightly more privileged access than a self existing in another body (Wilson, 2002). This provides some basis for comprehending the protest "I did it, but it wasn't me". Thoughts and actions may come from my body and brain, but not always with the stamp of self-approval.

Devine (1989), for example, showed that egalitarians and non-egalitarians alike automatically associated Blacks with negative stereotypes. Discrepancies between explicitly intended and automatically assessed evaluations have spurred confusion over how (or to whom) we are to attribute the automatic associations. Should automatic evaluations be considered a reflection of the person, even if they are consciously rejected? Or, do they provide little insight about the person and instead showcase the cultural context in which the person is embedded?

This paper reports an extensive investigation of the relations among cultural knowledge, explicit evaluations, and one popular measure of implicit evaluations—the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998). Cultural knowledge is not a singular experience. Perception of cultural norms varies widely across individuals based on their unique engagement with the cultural context. For example, Figure 1 presents two histogram plots (data from Study 6). One, cultural

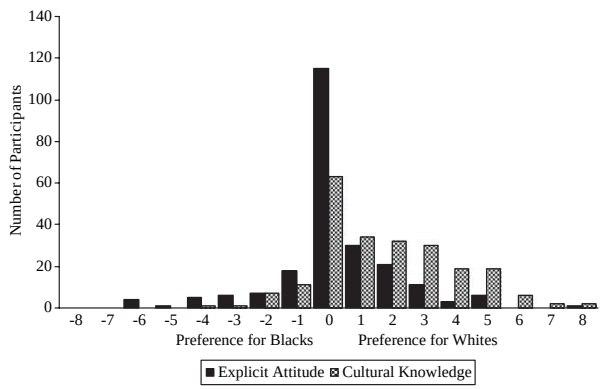

Figure 1. Histogram plots of explicit attitudes and cultural knowledge for evaluations of Blacks compared to Whites. Zero indicates no preference for Blacks compared to Whites.

knowledge, is the difference in response to: "How warmly does the average person feel toward African Americans [White Americans]?" The second is the difference in response to parallel items with oneself as attitude holder: "How warmly do you feel toward African Americans [White Americans]?"

On average, the culture ($M = 1.73$) is perceived to be more pro-White than the average self-rating ($M = 0.25$). However, there is remarkably little consensus in the perceived magnitude of the cultural preference ($SD = 2.14$). In this example, variability of cultural perceptions actually exceeds the variation in individual self-rated preferences ($SD = 1.89$). That is, there is less agreement about the cultural view than among individuals' attitudes. Finally, in this example, explicit attitudes and cultural knowledge are weakly related ($r = .10$, $p = .15$), suggesting two relatively independent classes of evaluative information—evaluations of one's own, and perceived evaluations of the culture.[1] The question that this paper is devoted to answering is: which of these evaluative sources is reflected in performance on the IAT?

Explicit attitudes are known to covary with IAT performance. The strength of that relationship ranges from near zero to strongly positive depending on a variety of features of evaluation such as self-presentation concern, attitude strength, and the degree to which one's evaluation is perceived to be distinctive (Hofmann, Gschwendner, Nosek, & Schmitt, 2005; Nosek, 2005). Also, while weak in the preceding example, explicit attitudes covary occasionally with cultural knowledge either because people use normative beliefs to guide their personal evaluations (Azjen & Fishbein, 2005) or because people use their personal evaluations to estimate what the rest of the culture is like, a phenomenon called naïve realism (Ross, 1996). In terms of clarifying an attitude–knowledge distinction, the former illustrates that there is a substantive relationship between these concepts that is attitude-relevant, i.e., that attitude–knowledge relations are not necessarily threats to accurate attitude measurement. The latter illustrates that people's impression of the culture is perceived through self-tinted glasses.

Theory and evidence concerning the presence of attitudes and knowledge in implicit evaluation

In her seminal demonstration of automatic racial biases, Devine (1989) understood the associations assessed by her evaluative priming procedure to reflect knowledge of a "culturally shared stereotype", whereas explicit measures of racial animus were reflective of personal attitudes. In other

[1] Other ways to represent cultural attitudes are to calculate the mean, median, or mode explicit (or implicit) attitude response from a representative sample of a culture. This approach is a cultural level of analysis whereas, for this paper, we are interested in the influence of cultural knowledge on individual minds. That is, does an individual's cultural knowledge influence their explicit or implicit evaluation?

words, self-report assessed individual differences in racial attitudes, and the priming procedure assessed cultural knowledge that was shared and not reflective of individual differences.

Claiming that a measure administered to persons reflects information about the culture suggests that a cultural construct is represented at a subordinate level of analysis—the minds of individuals living in that culture. This implies that there is no meaningful variability in automatic evaluations across individuals because the variability would show that, by definition, the variation in implicit bias reflects an individual difference, not a cultural constant. As Fazio, Jackson, Dunton, and Williams (1995) explained: "If, as Devine suggested, the shared cultural stereotype is activated in the presence of a minority group, one would expect little meaningful variation in the pattern of facilitation across participants. On the other hand, if it is one's personal evaluation that is activated in the presence of a minority group member, the variation across participants would be more substantial and predictive of race-relevant behaviors" (p. 1015). Indeed, Fazio et al. demonstrated that individual variation in implicit racial evaluations was associated with rated friendliness and interest during an interaction with a Black experimenter. This observation is critical in showing that the meaningful variation across individuals indicates that the implicit evaluation reflects something about persons rather than groups or cultures—even if the origins of these evaluations are steeped in the cultural context, as is all of human experience. Further, as illustrated in Figure 1, the notion of *culturally shared* knowledge is a misnomer as there is little consensus about what the culture thinks. Indeed, across almost 100 topics investigated in this paper, none came close to approximating consensus in cultural knowledge.

In retrospect, data from Devine (1989) also showed variability in perceptions of stereotypes. In the first study, participants reported the cultural stereotype about African Americans. Far from consensus, not a single characteristic was generated by all participants. In fact, most qualities (e.g., low intelligence, uneducated, sexually perverse) were mentioned by between just 20% and 50% of the respondents indicating substantial variability in the perception of cultural stereotypes. Real people aside, even psychologists exhibit variability in perceptions of racial animus in present-day American culture, with some arguing that prejudice is still widespread (e.g., Sears, 2004), and others arguing that it is vastly overestimated by social scientists (e.g., Arkes & Tetlock, 2004). From a cultural psychology perspective, these data come as little surprise because culture is not a singular construct recorded invariantly across minds. Conceiving the influence of culture as a constant makes it too easy to reject the influence of culture on individuals. Individuals have unique, personal experiences of their cultural context and this is reflected in the fact that cultural perceptions vary across individuals (Fiske, Kitayama, Markus, & Nisbett, 1998).

The fact that there is variation in cultural knowledge, suggests that the mental store includes evaluations that are believed (personal attitudes) and evaluations that are known but not believed (cultural knowledge). Explicitly, it is clear that these two sources of information are distinguishable. It is easy for someone to report that he hates green eggs and ham and will not eat them, and simultaneously report that Sam likes them and should eat them himself. With awareness and control, one can opt to use one's own evaluations to guide judgement and behaviour, and choose *not* to use knowledge about others' evaluations. Implicitly, the sensibility of such a distinction is less clear.

A variety of perspectives expect that IAT scores, and implicit attitudes more generally, are steeped in and influenced by the cultural context (Arkes & Tetlock, 2004; Banaji, 2001; Banaji, Nosek, & Greenwald, 2004; Karpinski & Hilton, 2001; Olson & Fazio, 2004). These views differ in the extent to which they assume that the IAT is influenced by cultural experience and whether such influence is an *extrapersonal* contaminant in the assessment of attitudes, or a feature of the implicit attitude construct. These views will be raised again in the general discussion. For the present studies, we investigated a relationship that all views would readily accommodate—that the IAT and cultural knowledge are positively correlated independent of their known relations with explicit attitudes.

Searching for attitudes and knowledge in implicit evaluation

There were three anticipated stages for this investigation. First, we sought to establish whether there is a relationship between cultural knowledge and the IAT. Such a relationship is already known to exist between explicit attitudes and the IAT (Nosek, 2005). Cultural knowledge was operationalised as one's perception of cultural attitudes measured in a variety of ways. Second, we examined whether IAT–knowledge relations exist independently of IAT–attitude relations. This is a necessary condition for the claim that IAT performance is influenced by cultural knowledge that is not part of the explicit attitude construct. And, third, if there is an independent component of IAT performance that is related to cultural knowledge, we would investigate whether this influence is construct-valid, i.e., part of what meaningfully distinguishes implicit attitudes from explicit attitudes, or whether it is a contaminating influence that obscures attitude measurement in the IAT.

We report a broad examination of the IAT–knowledge relationship with 158 samples and a variety of topics. These studies are organised and reported as 5 individual studies (Studies 1–5), and 2 mega-studies (Studies 6–7) the first with samples for 58 topics, and the second with samples for 95 topics. As will be discussed below, we observed a relationship between the IAT and cultural knowledge, and that relationship was accounted for by explicit

attitudes. The IAT and cultural knowledge did not possess an independent relationship to investigate as being either construct-relevant or a confounding influence. In the general discussion, we argue that, implicitly, associations are not coded as belonging to the self or not, and that cultural influences on implicit cognition are not readily accessible to introspection.

STUDY 1

The first study examined the relations among cultural knowledge, explicit attitudes, and the IAT with racial attitudes—a domain that is suggested as a prototypical illustration for the influence of cultural knowledge on the IAT (Arkes & Tetlock, 2004; Gehring, Karpinski, & Hilton, 2003; Karpinski & Hilton, 2001; Olson & Fazio, 2004).

Method

Participants

One hundred forty-nine University of Virginia undergraduates (ages 17–22) participated in the study for partial course credit. Prior to analysis, one participant was removed because of a computer malfunction. Initial analysis removed six participants for high error rates ($>20\%$) in at least one of the response latency tasks. Of the remaining 142 participants, 99 were female and 43 were male; 109 were White, 15 were Asian, 11 were Black, 8 were Hispanic, 8 were a different ethnicity, and 1 did not report ethnicity.

Materials

Stimulus items. Four Black faces and four White faces taken from the IAT demonstration website (https://implicit.harvard.edu/; Nosek, Banaji, & Greenwald, 2002a) served as exemplars for the "Black American" and "White American" categories in the IAT. Stimuli for this and the other studies are available in supplementary materials at http://briannosek.com/.

Implicit Association Test. Participants completed seven blocks of response trials. First, participants sorted evaluative words for 20 trials into categories (Pleasant/Unpleasant) using two response keys on a standard keyboard. Second, using the same response keys participants sorted faces and words associated with Black Americans and White Americans for 20 trials into categories (Black American/White American). Third, participants sorted items for all four categories for 20 trials using the two response keys. One key was used to categorise White American and Pleasant items; the other key was used to categorise Black American and Unpleasant items. Fourth, the same task was repeated for 40 more trials. Fifth, like the 2nd

block, participants sorted Black American and White American items again for 20 trials except that the response mapping was reversed (i.e., if White American items were categorised with the left key before, they were now categorised with the right key).[2] Sixth, again participants sorted items from all four categories for 20 trials except that the response mappings for the category exemplars (Black American/White American) were opposite of the 3rd and 4th blocks. So, in this example, White American and Unpleasant were sorted with one key and Black American and Pleasant were sorted with the other. And, seventh, participants repeated the sorting conditions in the 6th block for 40 more trials.

In blocks with four categories, trials alternated between presenting category (Black American, White American) and attribute (Pleasant, Unpleasant) items. Also, reminder labels appeared at the top of the screen for all blocks reminding participants of the categorisation rules. Further, to emphasise the distinction between the category and attribute dimensions, "Black American/White American" labels and items appeared in white, and "Pleasant/Unpleasant" labels and items appeared in green, all on a black background. Categorisation errors were identified with a red "X" below the stimulus item and participants had to correct the response before continuing to the next trial. An interstimulus delay of 150 milliseconds separated each trial. Finally, the order of the category mapping conditions (Black American with Unpleasant before or after Black American with Pleasant) was counterbalanced between subjects.

IAT analysis followed recommendations of Greenwald, Nosek, and Banaji (2003). The IAT *D* score was coded such that positive values indicated liking for White Americans relative to Black Americans. Reliability calculated on four data parcels was good ($\alpha = .91$).

Explicit measures. Explicit attitudes were assessed by calculating difference of self-reported feelings of warmth and liking of Black and White Americans. These items were standardised ($SD = 1.0$) and averaged to create a single explicit preference index ($\alpha = .77$). Cultural knowledge was assessed by calculating difference of ratings of the historical favourability of society toward Black and White Americans, the favourability of cultural portrayals of the racial groups, and American society's warmth for the racial groups. These items were standardised ($SD = 1.0$) and averaged to create a single cultural knowledge index ($\alpha = .70$). Reported results for this and the other studies are consistent when explicit and cultural items are considered individually as they are as an aggregate. Individual items and descriptive statistics are presented in Table 1.

[2] For Studies 2–7, 40 trials were used in this block based on recommendations from Nosek, Greenwald, and Banaji (2005) to reduce the influence of the order of combined tasks.

TABLE 1
Descriptive statistics for explicit attitude, cultural knowledge, cultural criterion, and IAT measures for Studies 1–7

Type	*Measure*	*Mean*	*SD*
Study 1			
Attitude	Difference between items "How warmly do you feel toward Black Americans [White Americans]?" (−8 strongly pro-Black to +8 strongly pro-White)	0.6	1.6
Attitude	Difference between items "How much do you like Black Americans [White Americans]?" (−6 strongly pro-Black to +6 strongly pro-White)	0.3	1.0
Knowledge	Difference between items "How warmly does American society feel toward Black Americans [White Americans]?" (−8 strongly pro-Black to +8 strongly pro-White)	2.9	2.0
Knowledge	Difference between items "How favourably or unfavourably does American society portray Black Americans [White Americans]?" (−6 strong favours Black to +6 strongly favours White)	2.4	1.6
Knowledge	Difference between items "Historically, how favourably or unfavourably has American society been for Black Americans [White Americans]?" (−6 strongly favoured Black to +6 strongly favoured White)	3.7	1.9
IAT	Implicit preference for White Americans compared to Black Americans (−2 pro-Black to +2 pro-White)	0.47	0.37
Study 2			
Attitude	Difference score between items "How favourable do you feel toward peanuts [shellfish]?" (−100 strongly pro-shellfish to +100 strongly pro-peanut)	15.7	47.4
Attitude	Difference score between items "How much do you like to eat peanuts [shellfish]?" (−5 like to eat shellfish a lot to +5 like to eat peanuts a lot)	0.5	2.5
Attitude	Difference score between items "How much do you like peanuts [shellfish]?" (−10 like shellfish a lot to +10 like peanuts a lot)	1.3	4.6
Knowledge	Difference score between items "How favourable is American culture toward peanuts [shellfish]?" (−5 favourable toward shellfish to +5 favourable toward peanuts)	0.9	1.5
Knowledge	Difference score between items "How much does the average person like peanuts [shellfish]?" (−5 likes shellfish to +5 likes peanuts)	1.0	1.2
IAT	Implicit preference for peanuts compared to shellfish (−2 pro-shellfish to +2 pro-peanut)	0.18	0.45
Study 3			
Attitude	Difference score between "How favourable do you feel towards George Bush [John Kerry]?" (−100 pro-Kerry to +100 pro-Bush)	−23.7	51.5

TABLE 1 (*Continued*)

Type	*Measure*	*Mean*	*SD*
Attitude	Difference score of mean semantic differential ratings: "Who is more intelligent, likeable, qualified, and has stronger character?" (4 items;—6 pro-Kerry to +6 pro-Bush)	−0.9	2.4
Attitude	Difference score between "How much do you like George Bush [John Kerry]?" (−6 pro-Kerry to +6 pro-Bush)	−1.2	2.9
Attitude	If an election involving George Bush and John Kerry for president were held today, for whom would you vote? (1 definitely John Kerry to 7 definitely George Bush)	2.8	2.4
Knowledge	Does the average person like George Bush or John Kerry? (1 like Kerry to 7 like Bush)	3.8	1.6
Knowledge	Do most people like George Bush or John Kerry? (1 like Kerry to 7 like Bush)	3.8	1.5
Knowledge	Difference score between "How much does the average person like or dislike George Bush [John Kerry]?" (−5 pro-Kerry to +5 pro-Bush)	−0.4	1.6
Knowledge	Difference score between "How warm or cold is society to George Bush [John Kerry]?" (−5 pro-Kerry to +5 pro-Bush)	−0.4	1.7
Knowledge	Difference score between "Historically, how favourable or unfavourable has American society been towards George Bush [John Kerry]?" (−5 pro-Kerry to +5 pro-Bush)	0.5	1.7
Knowledge	Difference score between "How favourable or unfavourable does American society portray George Bush [John Kerry]?" (−5 pro-Kerry to +5 pro-Bush)	−0.2	1.9
Criterion	If the vote were held today, who would be elected, John Kerry or George Bush? (−3 John Kerry to +3 George Bush)	−0.3	1.7
Criterion	Who will win the election in November, John Kerry or George Bush? (−3 John Kerry to 3 George Bush)	−0.4	1.8
IAT	Implicit preference for George Bush compared to John Kerry (−2 pro-Kerry to +2 pro-Bush)	−0.18	0.47
Study 4			
Attitude	Difference between items "How warmly do you feel toward Black Americans [White Americans]?" (−10 strongly pro-White to +10 strongly pro-Black)	−0.2	1.8
Attitude	Difference between items "How much do you like Black Americans [White Americans]?" (−6 strongly pro-White to +6 strongly pro-Black)	−0.1	1.1
Knowledge	Difference between items "How warm or cold is society toward Black Americans [White Americans)?" (−5 strongly pro-White to I 5 strongly pro-Black)	−2.1	1.5
Knowledge	Difference between items "How favourably or unfavourably does American society portray Black Americans [White Americans]?" (−5 strongly favours White to +5 strongly favours Black)	−2.4	1.7

TABLE 1 (*Continued*)

Type	*Measure*	*Mean*	*SD*
Knowledge	Difference between items "historically, how favourably or unfavourably has American society been for Black Americans [White Americans]?" (−5 strongly favoured White to +5 strongly favoured Black)	−3.9	1.5
Knowledge	Do most people prefer Black Americans or White Americans? (1 strongly pro-White to 7 strongly pro-Black)	2.9	1.2
Knowledge	Does the average person prefer Black Americans or White Americans? (1 strongly pro-White to 7 strongly pro-Black)	2.8	1.2
Knowledge	Difference between items "How much does the average person like or dislike Black Americans [White Americans]?" (−5 likes White to +5 likes Black)	−0.8	1.2
Criterion	If given the choice, who would most employers choose to hire, a Black American or a White American? (1 definitely White to 7 definitely Black)	3.0	1.1
Criterion	Who is more likely to be a target of discrimination, a Black American or a White American? (1 definitely White to 7 definitely Black)	2.2	1.4
IAT	Implicit preference for White Americans compared to Black Americans (−2 pro-White to +2 pro-Black)	−0.30	0.43
Study 5			
Attitude	Difference score of mean semantic differential ratings for "Candy bars [Apples] are . . . " (−6 apples to +6 candy bars; ugly–beautiful, horrible–wonderful, disgusting–tasty, bad–good, unpleasant–pleasant)	−0.7	1.4
Attitude	Difference score between items "How favourable do you feel toward candy bars [apples]?" (−100 strongly pro-apples to +100 strongly pro-candy bars)	−9.1	32.3
Knowledge	Difference between items "How warm or cold is society to candy bars [apples]?" (−10 strongly pro-apples to +10 strongly pro-candy bars)	−0.4	1.6
Knowledge	Difference between items "How favourably or unfavourably does American society portray candy bars [apples]?" (−6 strongly favours apples to +6 strongly favours candy bars)	−1.0	1.9
Knowledge	Difference between items "Historically, how favourably or unfavourably has American society been toward candy bars [apples]?" (−5 strongly favoured apples to +5 strongly favoured candy bars)	−0.6	1.5
Knowledge	Do most people prefer candy bars or apples? (1 apples to 7 candy bars)	4.3	1.5
Knowledge	Does the average person prefer candy bars or apples? (1 apples to 7 candy bars)	4.4	1.5
Knowledge	Difference between items "How much does the average person like or dislike candy bars [apples]?" (−5 likes apples to +5 likes candy bars)	0.7	1.0
Criterion	If given a choice between an apple and a candy bar, which would most people choose to eat? (−3 candy bars to +3 apples)	1.1	1.2

TABLE 1 (*Continued*)

Type	*Measure*	*Mean*	*SD*
Criterion	Which is purchased more frequently in stores—apples or candy bars? (−3 candy bars to +3 apples)	1.0	1.8
IAT	Implicit preference for candy bars compared to apples (−2 pro-apple to +2 pro-candy bar)	−0.32	0.43
Study 6			
Attitude	Single item preference measure, with A and B representing concepts of interest: I strongly prefer A to B to I strongly prefer B to A with "moderately", "slightly", and "I have no preference between A and B" in between. (−3 strongly pro-B to +3 strongly pro-A)	0.5	3.0
Knowledge	Single item preference measure, with A and B representing concepts of interest: The average person strongly prefers A to B to The average person strongly prefers B to A with "moderately", "slightly", and "The average person has no preference between A and B" in between. (−3 strongly pro-B to +3 strongly pro-A)	0.6	2.6
IAT	Implicit preference for concept A compared to concept B (−2 strongly pro-B to +2 strongly pro-A)	0.27	0.44
Study 7			
Attitude	Single item preference measure, with A and B representing concepts of interest: I strongly prefer A to B to I strongly prefer B to A with "somewhat", "slightly", and "I like A and B equally" in between. (−3 strongly pro-B to I 3 strongly pro-A)	4.5	1.9
Knowledge	Does the average person prefer A or B? (1 strongly pro-B to 7 strongly pro-A)	4.3	1.7
Knowledge	Do most people prefer A or B? (1 strongly pro-B to 7 strongly pro-A)	4.3	1.7
Knowledge	Does the culture you live in prefer A or B? (1 strongly pro-B to 7 strongly pro-A)	4.4	1.9
IAT	Implicit preference for concept A compared to concept B (−2 strongly pro-B to +2 strongly pro-A)	0.29	0.45

Procedure

Participants completed implicit and explicit measures in counterbalanced order. Additional measures that are not relevant for the present purposes in this and the next studies are discussed in Nosek and Hansen (2007).

Analysis strategy

Data analysis for each study followed a two-step process. First, we tested separately whether there was any relationship between IAT, cultural knowledge, and explicit attitude measures. And, second, we conducted regression analyses to see if cultural knowledge related to the IAT after accounting for

explicit attitudes. With correlational data, these studies make no commitment to causation; they speak only to the relations among the constructs.

Results and discussion

Table 1 presents means and standard deviations for IAT, attitude, and knowledge measures. All three suggested a preference for Whites compared to Blacks. And, as noted in the introduction, cultural knowledge evidenced substantial variability.

Relations among the IAT, explicit attitudes, and cultural knowledge

Table 2 presents correlations among measures. Despite substantial variation, knowledge was unrelated to explicit attitudes and the IAT. This contrasts with a reliable, positive relationship between the IAT and explicit attitudes. This pattern of relations obviated the need to conduct regression analyses because knowledge related to neither of the other variables. In any case, the beta coefficients for simultaneous regressions predicting IAT effects are reported in Table 2. Consistent with the correlation evidence, a relationship between explicit attitudes and the IAT observed independently persisted when knowledge was included as a simultaneous predictor. Also, the non-relationship between knowledge and IAT effects was unchanged in the simultaneous model.

In summary, asking participants: "How warmly do you feel . . .?" elicited responses that corresponded with IAT effects. Altering the item to: "How warmly does American society feel . . .?" elicited a distinct response that did

TABLE 2
Correlations among IATs (I), explicit attitudes (E), and cultural knowledge (K), and beta-weights from regressions of attitudes and knowledge predicting IAT effects, independently and simultaneously (Studies 1–5)

			Correlations			*Betas from simultaneous regressions*	
Study	*Topic*	*N*	*I–E*	*I–K*	*E–K*	*E→I*	*K→I*
1	Black American–White American	142	**.31**	.00	−.02	**.308**	.001
2	Peanuts–Shellfish	235	**.51**	.10	**.16**	**.515**	.021
3	John Kerry–George Bush	365	**.63**	**.20**	**.34**	**.632**	−.016
4	Black American–White American	218	**.37**	−.09	−.09	.366	−.054
5	Candy Bar–Apple	371	**.37**	.03	.07	**.365**	.000

Note: Correlations and regression coefficients in bold are significantly different from zero ($p < .05$).

not correspond with IAT effects. This lack of relationship is not attributable to lack of variability in cultural knowledge. Also, the lack of IAT–knowledge relation is stark in contrast to the reliable IAT–attitude relation. This suggests that IAT effects are more a reflection of evaluations that people report as their own than evaluations that people report knowing, but belonging to others.

The potency of this conclusion is hampered by being a single demonstration in a single domain. For Study 2 we changed content domains to attitudes toward foods—a domain in which it is understood that individuals can have their own taste preferences that are distinct from cultural norms.

STUDY 2: FOOD ATTITUDES

In Study 2, inspired by Olson and Fazio's (2004) hypothetical example, we examined attitudes toward peanuts relative to shellfish (two common food allergies), and tested the relationship between different versions of the IAT and self-reported preferences, behaviours, and knowledge of others' preferences. This Internet study recruited participants through random assignment in a large study pool available at the research portion of the Project Implicit website (https://implicit.harvard.edu/; see Nosek et al., 2006b, 2007, for detailed reports on the virtual laboratory).[3]

Method

Participants

A total of 235 people (average age = 27, $SD = 11$; 69% female) completed the study materials. Of these, 187 were White, 10 Asian, 8 Black, 9 Hispanic, 14 a different ethnicity, and 7 did not report ethnicity.

Materials

IAT. Design of the IAT followed the procedures described in Study 1. Four pictures of shellfish and peanuts served as exemplars for those categories. Five IATs were removed from analysis for too many fast responses ($> 10\%$ of responses were less than 300 ms; Greenwald et al., 2003), and eight others because of missing data. The IAT effect was coded such that positive values indicated liking for peanuts relative to shellfish ($\alpha = .92$).

[3] Also, participants with peanut and shellfish allergies were recruited directly at a private website following Olson and Fazio's hypothesis that they should be particularly likely to evidence implicit negativity toward the food domain to which they were allergic. Because only 14 were recruited successfully, these data are not reported separately in the text—see supplementary materials for a report. Substantive results do not differ whether these participants are included or excluded.

Explicit measures. Participants completed a questionnaire assessing attitudes, allergies, eating behaviour and perceptions of others' evaluations for peanuts and shellfish. Participants rated their food attitudes on semantic differential scales (e.g., disgusting–tasty) and rated their liking and eating behaviour. These six items were standardised and aggregated into an explicit preference index ($\alpha = .97$). Likewise, participants rated perceptions of the extent to which American culture or the "average person" favoured or liked shellfish and peanuts. These two items were likewise aggregated ($\alpha = .61$). Measures and descriptive statistics of individual items are presented in Table 1. Positive values indicated greater positivity for peanuts relative to shellfish.

Procedure

After registering to participate in studies at Project Implicit, participants were randomly assigned to one of many dozens of studies, including this one. Those that received this study completed the implicit and explicit measures in a randomised order. Once selected, the registered participant was never again randomly assigned to the study on future visits to the virtual laboratory.

Results and discussion

Table 1 presents individual measures showing that participants implicitly and explicitly preferred peanuts, and believed that cultural preferences were more positive toward peanuts than shellfish. Again, culture knowledge measures were interrelated and diverse suggesting meaningful variability.

Relations among IAT, explicit attitudes, and cultural knowledge

Table 2 presents correlations among aggregate measures. Cultural knowledge was positively related with explicit attitudes ($r = .16$, $p = .02$), and was non-significantly related with IAT effects ($r = .10$, $p = .14$). These weak relations contrast with a substantial positive relationship between the IAT and explicit food attitudes ($r = .51$, $p < .0001$).

Despite the non-significant relation, we conducted regression analyses to determine if the minimal knowledge–IAT relation was accounted for by explicit attitudes. The simultaneous regression of explicit attitudes and cultural knowledge predicting IAT effects (Table 2) showed that explicit attitudes maintained a strong positive relationship with IAT performance and knowledge did not. This reiterates the correlation evidence.

Study 2 replicated Study 1, with the IAT being reliably related to explicit attitudes and not to cultural knowledge. Together these studies illustrate that IAT performance is related to evaluations that people explicitly attributed to themselves and not evaluations that were attributed to others. And, because

no initial IAT–knowledge relationship was observed, the regressions were not additionally informative. These studies are difficult for strong hypotheses about the IAT being related to cultural knowledge as opposed to attitudes. However, the findings suggest that a more in-depth, varied, and high-powered investigation is needed before a confident interpretation is possible.

STUDIES 3, 4, AND 5

In Studies 1 and 2 we observed that cultural knowledge is variable across individuals, not at all related to the IAT, and weakly related with explicit attitudes. Also, in both cases, the IAT and explicit attitudes were reliably correlated, even after accounting for cultural knowledge in the regression analyses. Despite the consistency of these findings, there are still important cautions about drawing a general inference: (a) *power*—even though self-reported attitudes related more strongly to the IAT than did knowledge, perhaps the lack of relationship between knowledge and the IAT was a consequence of insufficient power (however, note that *N*s were relatively large, and power to detect an $r \geq .30$ with a 2-tailed test at $\alpha = .05$ was .955 and .997 respectively); (b) *representation of knowledge*—it is possible that a more diverse assessment of knowledge would show relations to the IAT where the earlier items did not; and (c) *validity of cultural knowledge*—knowledge items used in the first studies were interrelated but no evidence was offered for their criterion validity, so perhaps cultural knowledge was just poorly assessed.

The next studies were designed to address these concerns. Studies 3–5: (a) had large samples (Study 3 $N = 365$; Study 4 $N = 218$; Study 5 $N = 371$); (b) included a wider range of cultural knowledge items; and (c) included criterion validity variables that should be predicted by cultural knowledge. Because Studies 3–5 used similar methods, they are described together with results and discussion following.

Methods for Studies 3, 4, and 5

Participants

Participation occurred through the research website of Project Implicit following the same procedures described in Study 2 (Study 3, $N = 365$; Study 4, $N = 218$; Study 5, $N = 371$).

Materials

While the topics varied, the basic form of the materials was constant across Studies 3, 4, and 5. Study 3 concerned attitudes toward *John Kerry* compared to *George Bush* (Nosek et al., 2007), Study 4 concerned attitudes

toward *Black people* compared to *White people*, and Study 5 concerned attitudes toward *candy bars* compared to *apples*, a comparison for which Karpinski and Hilton (2001) suggested that the IAT was influenced by cultural knowledge.

IAT. IATs were implemented following the procedure described in Study 1. Two additional between-subjects experimental conditions included variations of the IAT design, but the data from those conditions are not presented here (see Nosek & Hansen, 2007).[4] The IAT showed good internal consistency (Study 3: politics, $\alpha = .90$; Study 4: race, $\alpha = .88$; Study 5: food, $\alpha = .86$).

Explicit attitudes, cultural knowledge, and knowledge criterion variables. For each study a collection of explicit attitude (Study 3, $\alpha = .96$; Study 4, $\alpha = .85$; Study 5, $\alpha = .84$), cultural knowledge (Study 3, $\alpha = .84$; Study 4, $\alpha = .74$; Study 5, $\alpha = .67$), and knowledge criterion items were administered (see Table 1). Items were similar to previous studies though additional knowledge questions were administered to broaden representation of that assessment (results are the same when considered individually). Also, criterion variables for knowledge items in each content domain were identified to demonstrate predictive validity of cultural knowledge.

Procedure

The procedure was the same for Studies 3, 4, and 5. After being randomly assigned to the study and giving informed consent, subjects completed an IAT and a short questionnaire in randomised order. Also, item order in the questionnaire was randomised.

Results and discussion for Study 3 (political attitudes)

Relations among IAT, explicit attitudes, and cultural knowledge

Table 1 presents the means and standard deviations of the individual measures. As before, attitude and knowledge items were aggregated after standardising for summary analysis. Table 2 presents relations among aggregate measures. As in the first two studies, the IAT was positively related to explicit attitudes ($r = .63$, $p < .0001$). People who reported stronger preferences for Kerry compared to Bush also showed stronger pro-Kerry

[4] In brief, we also administered the "personalised" IAT introduced by Olson and Fazio (2004) as a corrective IAT design to reduce the influence of extrapersonal (e.g., cultural) knowledge. Nosek and Hansen (2007) found that the procedural changes do not influence the relationship with cultural knowledge, but do encourage a task recoding confound in which participants are more likely to explicitly evaluate the target concepts (e.g., Bush, Kerry) instead of just categorising them as instructed.

implicit preferences (see also Nosek, 2005; Nosek et al., 2007). This time, evaluative knowledge was significantly positively related to both the IAT ($r = .20$, $p = .0003$) and explicit attitudes ($r = .34$, $p < .0001$). Bush supporters compared to Kerry supporters, measured implicitly or explicitly, were more likely to believe that there was a cultural preference for Bush. Notably, cultural knowledge was more related to explicit attitude than to IAT assessments.

Criterion validity of cultural knowledge

One concern with the previous studies was that the weak-to-absent relationship between the IAT and knowledge could have been due to weaknesses in the measurement of cultural knowledge. In Study 3, two criterion validity variables—participants' predictions of who would win "if the election were held today", the day of their participation ($r = .61$, $p < .0001$), and predictions of the upcoming November 2004 election ($r = .49$, $p < .0001$)—were both strongly related to cultural knowledge estimates.[5] Cultural knowledge measures showed interindividual variability, internal consistency, and criterion validity.

Explicit attitudes accounted for the relationship between the IAT and cultural knowledge

Study 3 demonstrated a reliable IAT–knowledge relationship. The next step for evaluating the nature of this relation was to determine whether it existed independently of explicit attitudes. The simultaneous regression analysis is reported in Table 2. While cultural knowledge was correlated with IAT performance, there was no relationship in a simultaneous regression that included explicit attitudes as a predictor too ($\beta = -.016$, $p = .73$). Further, explicit attitudes continued to predict IAT performance in the simultaneous regression ($\beta = .632$, $p < .0001$). In short, the IAT–knowledge relationship was completely accounted for by explicit attitudes suggesting that the IAT–knowledge relation is a consequence of naïve realism—using one's own political attitudes to estimate the cultural preference (Ross, 1996).

Results and discussion for Study 4 (racial attitudes)

Relations among the IAT, explicit attitudes, and cultural knowledge

Table 1 shows means and standard deviations for individual measures. Replicating Study 1, aggregated cultural knowledge was unrelated to explicit racial attitudes ($r = -.09$, $p = .18$) and to implicit racial attitudes, as

[5] This was observed for every one of the six knowledge items for both criterion variables (rs = .20–.62, ps < .0003). Dates of data collection were 28 May 2004 to 10 August 2004.

measured by the IAT ($r = -.09$, $p = .20$). At the same time, implicit and explicit racial attitudes were positively correlated ($r = .37$, $p < .0001$). People who reported stronger pro-White preferences also tended to show stronger implicit pro-White preferences.

As with Study 1, the lack of relationship between cultural knowledge and individual explicit and implicit attitude measures obviated the need to conduct regressions to account for the knowledge–IAT relationship. The regression analyses appear in Table 2 and show that, considered simultaneously, explicit racial attitudes predict IAT effects and cultural knowledge does not.

Criterion validity of cultural knowledge

Cultural knowledge (in aggregate and for individual items) was positively correlated with estimates of employers' preferences for hiring White over Black job candidates, and likelihood estimates of group members being targets of discrimination (see items in Table 1). Those who perceived others to have stronger pro-White preferences predicted more pro-White hiring practices ($r = .51$, $p < .0001$), but explicit racial attitudes were unrelated to perceived hiring practices ($r = -.04$, $p = .59$). Also, those who perceived others to have stronger pro-White preferences predicted greater likelihood of Blacks being discriminated against compared to Whites ($r = .47$, $p < .0001$), but explicit attitudes were unrelated to predicted discrimination rates ($r = -.07$; $p = .34$).

Study 4 reinforces the conclusions from the previous three studies. The IAT maintains a reliable, positive relationship with explicit attitudes—even in socially sensitive domains such as racial attitudes (Nosek, 2007; Nosek et al., 2007) and no unique relationship with parallel assessments of cultural attitudes—even though cultural knowledge perceptions were variable, reliable, and showed criterion validity.

Results and discussion for Study 5 (food attitudes)

Relations among the IAT, explicit attitudes, and cultural knowledge

Table 1 shows means and standard deviations for all measures. Again, aggregated cultural knowledge was unrelated to explicit food attitudes ($r = .07$, $p = .20$) and to implicit food attitudes ($r = .03$, $p = .64$). Also, again, implicit and explicit food attitudes were positively correlated ($r = .37$, $p < .0001$). A subset of two of the cultural knowledge items (average cultural attitude, most people's attitude) was significantly, though weakly, related to explicit attitudes ($r = .13$, $p = .01$), but even this subset was not related to implicit attitudes ($r = .03$, $p = .61$). Again, the non-relationship of knowledge with IAT effects negated the purpose of regressions to account

for a non-existent IAT–knowledge relation. Implicit and explicit attitudes continued to show positive correspondence when explicit attitudes and cultural knowledge were simultaneous predictors of IAT effects (see Table 2) reinforcing the interpretation of IAT effects as indicators of individual attitudes.

Criterion validity of cultural knowledge

Cultural knowledge was significantly correlated with perceptions of consumer purchasing behaviour of apples compared to candy bars, and perceptions of which item the most people would choose if given a choice to eat. Those who perceived the culture to have stronger candy bar preferences predicted more purchasing of candy bars compared to apples in stores ($r = .26$, $p < .0001$) and, demonstrating discriminant validity, individual explicit attitudes were weakly related to those estimates ($r = .11$, $p = .05$). Also, those who perceived others to have stronger candy bar preferences predicted more frequent candy bar selection compared to apple selection by others when given a choice ($r = .43$, $p < .0001$), and self-reported attitudes were not related to such estimates ($r = .09$, $p = .10$). In sum, cultural knowledge showed criterion validity for perceptions and predictions of others' food-related behaviour.

STUDY 6: MEGA-STUDY OF 58 TOPICS

Across five studies the IAT was consistently related to explicit attitudes across content domains, and simultaneously showed weak to absent relations with cultural knowledge. When the IAT did relate to cultural knowledge, the relationship was accounted for by explicit attitudes. This is inconsistent with a conclusion that the IAT captures perceptions of the culture that are independent of explicit, personal attitudes, whether that knowledge would be considered a meaningful part of the implicit attitude construct or not. Further, these effects were found with high-power tests and with simultaneous evidence that cultural knowledge assessments had meaningful variability and criterion validity.

Despite the consistent findings, it is possible that these findings are not general, and instead depend on the selection of attitude topics. One reason to be sceptical of this possibility is the fact that a domain theorised to clearly illustrate cultural knowledge influences on IAT effects is racial attitudes—the topic of both Studies 1 and 4, and Study 4's relation was even non-significantly in the wrong direction. Establishing generality is a critical step for understanding the relationships among implicit attitudes, explicit attitudes, and cultural knowledge. Also, it is dissatisfying to find only one instance (Study 3) in which an IAT–knowledge relationship was reliably

observed (even though that may be an accurate reflection of the actual relations). For Study 6, we conducted a mega-study with participants being randomly assigned to one of 58 topics. This study was similar to the study reported by Nosek (2005) that examined moderators of the relationship between implicit and explicit attitudes. Only the measures that were critical to the present investigation are described.

Method

Participants

A total of 5794 study sessions were completed by 4089 volunteer participants at the Project Implicit research site. Registration and random assignment procedures were the same as described in Study 2. After login, participants were randomly assigned to one of the 58 topics.[6] Participants who logged in again were randomly assigned to a topic that they had not yet seen. Analysing only the participant's first completed study does not alter the conclusions reported here (see also Nosek, 2005).

From the registration demographics survey the sample was: 68% female, 32% male; 0.8% American Indian, 5.7% Asian, 6.0% Black, 5.2% Hispanic, 74.1% White, 1.1% Biracial (Black–White), 3.9% Multiracial, and 3.2% Other; 19.5% Conservative, 30.1% Neutral or Moderate, and 50.4% Liberal; and, on average, 28.7 years old, $SD = 11.7$ years. 1.05% of the IAT data (61 sessions) were excluded because the data suggested careless task performance ($> 10\%$ of response latencies were shorter than 300 ms; see Nosek et al., 2007), and 5 sessions had data transfer problems leaving 5728 usable sessions.

Materials

IAT. Design of IATs followed the procedures described in Study 1. The object pairs and stimulus exemplars were the same as those described in Nosek (2005). IAT scores were calculated such that positive values indicated an implicit preference for the concept implicitly preferred on average. The explicit attitude and evaluative knowledge measures were similarly scaled.

Self-report measures. Explicit attitudes were assessed by the difference in feelings of warmth between the target attitude objects as described in Study 1. Cultural knowledge was assessed with parallel items for rating the attitudes of the average person. See summary in Table 1.

[6] This is one more topic than was examined by Nosek (2005). We also examined attitudes toward Burger King compared to McDonalds.

Procedure

The study was administered via the research website for Project Implicit (https://implicit.harvard.edu) between 13 October 2003 and 17 September 2004. Once randomly assigned to a study, participants completed explicit measures and the IAT in a randomised order. Explicit measures were presented on a single webpage in randomised order.

Analysis strategy

This study was essentially 58 studies of different topics with a common procedure. This design facilitated the use of multilevel analysis so that variability between topics could be distinguished from variability between individuals. Nosek (2005) pursued this strategy with a similar dataset to investigate moderators of implicit–explicit relations. He found that the IAT was positively related to explicit attitudes, and that the strength of this relationship varied across topics. Further, multiple attitude features such as self-presentation concern, attitude strength, attitude distinctiveness, and attitude dimensionality accounted for variation in implicit–explicit correspondence across individuals and helped account for why some topics showed stronger implicit–explicit correspondence than others.

The present study follows the logic of the Nosek (2005) approach but focused specifically on the relations among cultural knowledge, explicit attitudes, and the IAT. The approach is described briefly here, and expounded in the results section. A more detailed summary for this type of application is available in Nosek (2005) and a detailed introduction to multilevel analysis is available in Snijders and Bosker (1999).

Multilevel modelling with the large number of topics enabled a sequenced examination of our core questions. First, does there exist a relationship between IAT effects and cultural knowledge? The previous studies were limited to investigating this question for a single topic. Here, we tested 58 topics simultaneously. Second, does the strength of the IAT–knowledge relationship vary across topics? One speculation from the previous studies is that knowledge might be related to IAT effects for some topics and not for others. Multilevel modelling allows a formal test of this possibility.

Third, do IAT–attitude relations account for IAT–knowledge relations generally, and variation in IAT–knowledge relations across topics? The first part of this question is the universal form of the question tested in the first five studies. In other words, is the IAT–knowledge relationship dependent on their common relation with explicit attitudes? The second part of the question suggests that the known variation in IAT–attitude correspondence across topics (Nosek, 2005) might account for variations in IAT–knowledge correspondence across the same topics. For example, the fact that political attitudes showed a significant IAT–knowledge relation (Study 3) whereas

food and racial attitudes did not (Studies 1, 2, 4, 5), might be explained by the fact that political attitudes elicit stronger IAT–attitude correspondence than those other topics. This would reinforce the conclusion that cultural knowledge has no meaningful relationship with the IAT beyond that accounted for by explicit attitudes. In summary, multilevel modelling enables very high-powered, general tests of relations between attitude and knowledge constructs.

Results and discussion

Correlations among IAT, explicit attitude, and cultural knowledge measures

Our first analysis step was to describe the correlations among cultural knowledge, explicit attitudes, and the IAT. Box-and-whisker plots of correlations for all 58 topics are presented in Figure 2. The average sample size for each topic was 109. As Nosek (2005) observed, implicit and explicit attitudes were positively related and that relationship varied across topics. The median relation was $r = .35$ (minimum $r = -.08$, maximum $r = .74$). Of the 58 topics, 47 (81%) showed a significant positive relationship, 0 of 58 showed a significant negative relationship.

Cultural knowledge was positively and variably related to the IAT (median $r = .08$, minimum $r = -.22$, maximum $r = .42$). Of the 58 topics, 9 (16%) showed a significant positive relationship, 1 of 58 showed a significant negative relationship. This is a generalised confirmation of the findings from earlier studies: IAT–knowledge relations can be observed but they are infrequent and weak.

Finally, cultural knowledge was positively and variably related to explicit attitudes (median $r = .15$, minimum $r = -.24$, maximum $r = .51$). Of the 58 topics, 25 (43%) showed a significant positive relationship, 1 of 58 showed a significant negative relationship. This pattern of relations is consistent with the pattern observed across Studies 1–5: the IAT and explicit attitudes were consistently positively related and, when related, cultural knowledge was more consistently and strongly related to explicit attitudes than to IAT effects.

Explicit attitudes account for the relationship between the IAT and cultural knowledge across 58 topics

Predictions about the relations among constructs were tested by comparing the fit of successive multilevel models. Deviance scores from the chi-square were the basis for comparing models. Higher values indicate greater improvement in fit. A brief summary of models and parameters is presented in Table 3.

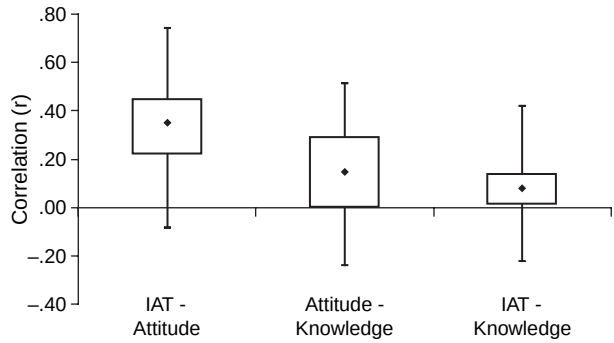

Figure 2. Box-and-whisker plots for attitude–knowledge, IAT–attitude, IAT–knowledge, correlations for 58 topics (Study 6).

Hypothesis 1: Cultural knowledge is related to IAT performance. Model M_0 is a baseline model representing two intercepts that indicate the average IAT effect (fixed effects intercept) and a coefficient showing that the magnitudes of IAT effects vary across topics (random effects intercept). Model M_1 adds cultural knowledge as a fixed effect predictor of IAT effects.

TABLE 3
Multilevel models testing whether cultural knowledge and explicit attitudes predict IAT effects across 58 topics (Study 6)

		Models					
Parameter/statistic		M_0	M_1	M_2	M_3	M_4	M_5
Fixed effects (individual diffs)							
Intercept		**.273**	**.273**	**.273**	**.242**	**.243**	**.243**
Cultural knowledge			**.016**	**.016**	.005	**.0054**	
Explicit attitudes					**.060**	**.056**	**.057**
Random effects							
Intercept variance		**.038**	**.034**	**.034**	**.033**	**.036**	**.037**
Cultural knowledge				**.0003**	.00009	.00002	
Explicit attitudes						**.0004**	**.0004**
Goodness-of-Fit							
−2 Log Likelihood	*–2LL*	6942	6916	6904	6029	5986	5981
Change –2LL	*LRT*		26	38	913	956	961
Change in *df* from M_0			1	2	3	4	2

Note: All models were fit with SAS PROC MIXED with $N = 5769$ and Topics = 58. Boldface MLE parameter estimates are significant at $p < .05$. All variables have rational zero points indicating comparative attitudinal indifference, positive values indicate preference for attitude object implicitly preferred on average.

That is, in general, does cultural knowledge predict IAT performance? Confirming the correlation evidence from the previous section, comparison of Models M_1 and M_0 showed a small, but reliable improvement of model fit $\Delta\chi^2(1) = 26$. Cultural knowledge predicted IAT performance reliably, $B = .016$, $SE_B = .0027$, $t(5669) = 6.00$, $p < .0001$, $d = .16$.

Hypothesis 2: The magnitude of the IAT–knowledge relationship varies across topics. Compared to Model M_1, Model M_2 adds a random-effects parameter for cultural knowledge to test whether the magnitude of the IAT–knowledge relationship varied across topics. The model fit comparison was reliable and weak, $\Delta\chi^2(1) = 12$. There was a slight tendency for some domains to show stronger IAT–knowledge relations than others, $B = .00032$, $SE_B = .00014$, $z = 2.28$, $p = .011$.

Hypothesis 3: Explicit attitudes account for IAT–knowledge relations. Nosek (2005) found that explicit attitudes were reliably related to IAT effects, and that this relationship varied across topics. This study replicates those observations and, more critically, tests whether they can account for the observed relations between cultural knowledge and IAT performance. Compared to Model M_2, Model M_3 adds a fixed effects parameter for explicit attitudes predicting the IAT. If explicit attitudes do account for the relationship, then this parameter should predict IAT effects and eliminate the IAT–knowledge relationship. The model fit comparison was reliable and substantial, $\Delta\chi^2(1) = 875$, showing that explicit attitudes are a good predictor of IAT effects, $B = .060$, $SE_B = .0019$, $t(5611) = 31.05$, $p < .0001$, $d = .83$. And, the addition of explicit attitudes eliminated IAT–knowledge relations overall ($t = 1.81$, $p = .08$, $d = .05$) and variation in IAT–knowledge relations across topics ($z = 1.00$, $p = .15$). This confirms the prediction that, when the IAT and cultural knowledge are related, the correspondence is a consequence of their shared relationship with explicit attitudes.

Model M_4 further illustrated that explicit attitudes were variably related with IAT effects, replicating Nosek (2005), $\Delta\chi^2(1) = 43$ (M_4 compared with M_3). Figure 3 presents the parameter estimates of cultural knowledge and explicit attitudes predicting IAT effects from this model separately for the 58 topics. Topics are ordered from the weakest IAT–attitude relationship estimate at the top, to the strongest estimate at the bottom. Error bars represent their 95% confidence intervals. Evident across topics is the lack of variation in knowledge estimates and the confidence interval overlap with 0 (no IAT–knowledge relationship) in direct contrast with stronger and more variable estimates relating explicit attitude reports and IAT effects.

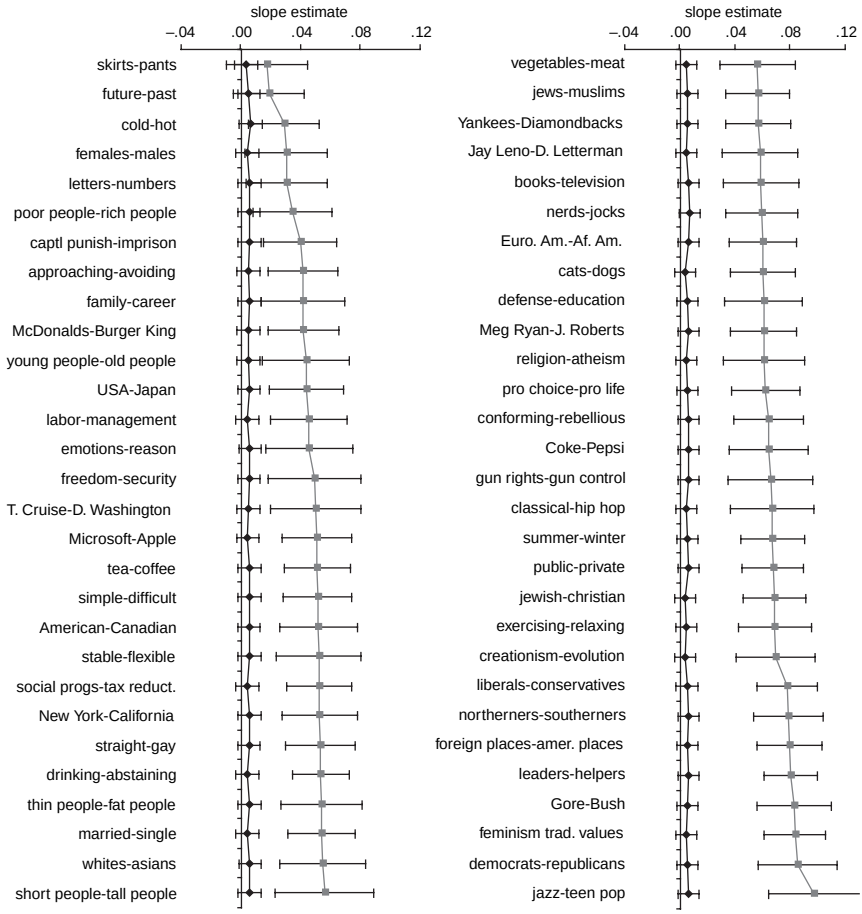

Figure 3. Multilevel modelling parameter estimates of explicit attitudes and cultural knowledge predicting IAT effects for 58 topics with 95% confidence intervals (Study 6, Model M_4).

Finally, Model M_5 removed cultural knowledge from the model. If this model fit the data as well as Model M_4, it suggests that cultural knowledge did not provide independent predictive validity of IAT effects and that the more parsimonious model with just explicit attitudes should be preferred. Comparisons of model fit showed that the simpler model fit equally well, $\Delta\chi^2(2) = -5$ (the even slightly better fit estimates for the simpler model is not meaningful). In summary, Study 6 strongly supported the hypothesis that, even under conditions of a highly reliable IAT–knowledge relationship, explicit attitudes accounted for the relationship and were much more strongly related to IAT effects.

STUDY 7: MEGA-STUDY OF 95 TOPICS

Study 6 provided evidence that the findings from Studies 1–5 were not just a function of the topics investigated. While the findings were general across topics and were highly powered for overall relations, it is possible that there are effects for individual topics that were not detected because the average sample size for any given topic was just over 100. We conducted another mega-study that expanded the variety of topics and dramatically increased the sample size (total $N > 100{,}000$) so that each domain could be tested with very high power and even small relations could be estimated reliably.

Method

Participants

A total of 105,934 sessions included an IAT, explicit attitude, and cultural knowledge measures completed by 66,074 volunteer participants. After login, participants were randomly assigned to one of the 95 topics. The sample was: 62% female, 38% male; 0.8% American Indian, 6.2% Asian, 5.3% Black, 5.1% Hispanic, 74.5% White, 1.0% Biracial (Black–White), 4.0% Multiracial, and 3.1% Other; 20% Conservative, 24% Neutral or Moderate, and 56% Liberal; and, 31.2 years old on average, $SD = 12.2$ years. Approximately 4.6% of the sessions had some missing data for the IAT for any of a variety of reasons (e.g., data transfer errors during high traffic events). For these, IATs were coded as missing. Finally 1.0% of the IAT data was excluded because the data suggested careless task performance ($> 10\%$ of response latencies were shorter than 300 ms, $> 40\%$ errors in a response block; see Nosek et al., 2006b) leaving 100,063 usable sessions.

Materials

IAT. Design of IATs followed the procedures described in Study 1. The object pairs and stimulus exemplars used the same set described in Study 6 with additional topics for a total of 95 (see supplementary materials). IAT scores were calculated such that positive values indicated an implicit preference for the concept implicitly preferred on average. Explicit attitude and cultural knowledge measures were similarly scaled.

Explicit attitudes. Explicit attitudes were assessed with a 7-point single-item preference measures with response options from "I strongly prefer A to B" to "I strongly prefer B to A" with "somewhat", "slightly", and "I like A and B equally" in between. A and B were the same concepts represented in the IAT and knowledge items.

Cultural knowledge. Evaluative knowledge was assessed with three 7-point items, "Does the average person prefer A or B?", "Do most people prefer A or B?", and "Does the culture you live in prefer A or B?" All had response options from "strongly prefer A to B" to "strongly prefer B to A". Each participant received a random subset of four of these plus three other items with the possibility of getting one to three of the items with an average of two of the three. The items were reliably correlated ($\alpha = .85$) and averaged to form a single cultural attitude index. Results do not differ if items are considered individually.

Procedure

The study was administered via the research website for Project Implicit between 17 September 2004 and 17 October 2006. Once randomly assigned to a study, participants completed explicit measures and the IAT in a randomised order. Explicit measures were presented on a webpage in randomised order.

Results and discussion

Correlations among the IAT, explicit attitudes, and cultural knowledge

The analysis strategy was identical to the report for Study 6. Box-and-whisker plots for correlations are presented in Figure 4. The average sample size for each topic was 1053. Consistent with Study 6 and Nosek (2005), the IAT and explicit attitudes were positively related and that relationship varied across topics. The median relation was $r = .36$ (minimum $r = .07$, maximum $r = .70$). All 95 of the 95 topics (100%) showed a significant positive relationship.

Cultural knowledge was positively and variably related to the IAT (median $r = .04$, minimum $r = -.14$, maximum $r = .33$). Of the 95 topics, 39 (41%) showed a significant positive relationship, 9 of 95 showed a significant negative relationship (10%). This replicates earlier studies and shows that the infrequency of relations was a consequence of the effects being very small. Many more significant relations were observed when power was extremely high (most sample Ns > 1000). IAT–knowledge relations are reliable, very weak, and sometimes in the opposite direction from the hypothesis that perceiving stronger cultural preferences for one topic relates to stronger implicit preference for that topic.

Finally, cultural knowledge was positively and variably related to explicit attitudes (median $r = .09$, minimum $r = -.20$, maximum $r = .46$). Of the 95 topics, 58 (61%) showed a significant positive relationship, 8 of 95 showed a significant negative relationship (8%). This pattern replicates

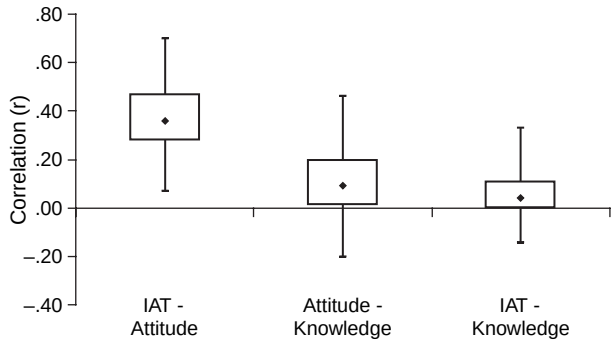

Figure 4. Box-and-whisker plots for attitude–knowledge, IAT–attitude, IAT–knowledge, correlations for 95 topics (Study 7).

earlier studies—implicit and explicit attitudes were consistently related, and cultural knowledge was more reliably related to explicit attitudes than to IAT effects. Notably, none of the significant IAT–knowledge and explicit-knowledge relations had mismatching signs. When cultural knowledge was negatively related to the IAT it was also negatively related to explicit attitudes. This matching is suggestive that the mechanism manufacturing the two relations is the same, i.e., that the IAT–knowledge relation is a function of the attitude–knowledge relation, a question tested next.

Explicit attitudes account for the relationship between IAT effects and cultural knowledge

As in Study 6, predictions about the relations among constructs were tested by comparing the fit of successive models. A brief summary of models and parameters is presented in Table 4. With a sample size of just over 100,000, we would expect that effects of trivial magnitude to be reliably estimated, making statistical significance testing uninteresting. The emphasis, therefore, is on effect magnitude.

Hypothesis 1: Cultural knowledge is related to IAT performance. Compared to baseline Model M_0, Model M_1 adds cultural knowledge as a fixed effect predictor of IAT effects. Consistent with the correlation data, Model M_1 showed a significant improvement of model fit, $\Delta\chi^2(1) = 267$. Cultural knowledge predicted IAT performance reliably and weakly, $B = .018$, $SE_B = .0011$, $t(1.0 \times 10^5) = 16.70$, $p < .0001$, $d = .11$.

Hypothesis 2: The magnitude of the IAT–cultural knowledge relationship varies across topics. Compared to Model M_1, Model M_2 adds a random

TABLE 4
Multilevel models testing whether cultural knowledge and explicit attitudes predict IAT effects across 95 topics (Study 7)

		Models					
Parameter/statistic		M_0	M_1	M_2	M_3	M_4	M_5
Fixed effects (individual diffs)							
Intercept		**.290**	**.212**	**.203**	**−.182**	**−.164**	**−.149**
Cultural knowledge			**.018**	**.020**	**.005**	**.004**	
Explicit attitudes					**.101**	**.098**	**.099**
Random effects							
Intercept variance		**.056**	**.055**	**.088**	**.054**	**.105**	**.099**
Cultural knowledge				**.001**	**.0002**	**.0001**	
Explicit attitudes						**.001**	**.001**
Goodness-of-Fit							
−2 Log Likelihood	*–2LL*	126,383	126,116	125,597	109,508	107,951	108,010
Change –2LL	*LRT*		267	786	16,875	18,432	18,373
Change in *df* from m_0			1	2	3	4	2

Note: All models were fit with SAS PROC MIXED with $N = 100{,}063$ and Topics = 95. Boldface MLE parameter estimates are significant at $p < .05$. All variables have rational zero points indicating comparative attitudinal indifference, positive values indicate preference for attitude object implicitly preferred on average.

effects parameter for cultural knowledge to test whether the magnitude of the IAT–knowledge relationship varied across topics. The comparison of model fit was reliable, $\Delta\chi^2(1) = 519$. Some domains showed stronger IAT–knowledge relations than others, $B = .00081$, $SE_B = .00013$, $z = 6.02$, $p < .0001$.

Hypothesis 3: Explicit attitudes account for IAT–knowledge relations. Compared to Model M_2, Model M_3 includes explicit attitudes to see if it reduces or eliminates the IAT–knowledge relationship. The improvement of model fit was reliable and substantial, $\Delta\chi^2(1) = 16{,}089$, showing that explicit attitudes predict IAT effects, $B = .101$, $SE_B = .00076$, $t(1.0 \times 10^5) = 132.44$, $p < .0001$, $d = .84$. And, the addition of explicit attitudes accounted for IAT–knowledge relations, reducing the parameter estimate by 77% to a trivial magnitude, $B = .0045$, $SE_B = .0018$, $t = 2.50$, $p = .014$, $d = .016$. Also, variation in IAT–knowledge relations across topics was reduced by 74%, $B = .00021$, $SE_B = .000045$, $z = 4.73$, $p < .0001$. This confirms the prediction that, when the IAT and cultural knowledge are related, the correspondence is a consequence of their shared relationship with explicit attitudes.

Replicating Study 6 and Nosek (2005), Model M_4 showed that explicit attitudes were variably related with IAT effects, $\Delta\chi^2(1) = 1557$ (M_4

compared with M_3). Figure 5 presents the parameter estimates of cultural knowledge and explicit attitudes predicting IAT effects simultaneously for the 95 topics. Topics are ordered from the weakest independent IAT–attitude relationship estimate at the top left, to the strongest estimate at the bottom right. Error bars represent 95% confidence intervals for the estimates. Across topics, cultural knowledge varied weakly across topics with sporadic examples of the confidence interval around the knowledge estimate not including 0 (no relationship). Simultaneously, none of the explicit attitude confidence intervals included zero showing that all of them reliably predicted IAT effects, even when considered simultaneously with cultural knowledge. Further, the distributions of explicit attitude and cultural knowledge parameter estimates were almost completely non-overlapping. Every single explicit attitude topic parameter exceeded every single cultural knowledge parameter, except for one (approaching–avoiding). That estimate was smaller than just three of the 95 cultural knowledge parameter estimates (Canadian–American, northerners–southerners, cold–hot). In terms of familiar effect sizes, the two *largest* parameter estimates for cultural knowledge predicting IAT effects independently translated to *r*s of approximately .15 and .13, and across all 95 the estimates translate to a median *r* of approximately .01.

Finally, as in Study 6, Model M_5 removes cultural knowledge from the model. Comparisons of model fit showed a slight decrease in overall fit when cultural knowledge was not included, $\Delta\chi^2(2) = 59$ (M_5 compared with M_4). When that is compared against the improved change in fit of almost 18,000 of models with and without explicit attitudes predicting IAT effects (M_4 compared to M_2), it is clear that cultural knowledge is carrying near zero predictive capacity of IAT effects on its own.

In summary, Study 7 affirmed the evidence offered in Studies 1–6 with a very large and topically varied sample. IAT–knowledge relations can be reliably estimated. However, those relations are weak and are accounted for by explicit attitudes. In short, if one had a goal to predict people's IAT scores, it would be useful to ask them how they feel about the topics (more so for some topics than others), and it would be virtually useless to also ask how other people feel about the topics.

GENERAL DISCUSSION

In seven studies, we investigated the extent to which implicit attitudes, as measured by the IAT, corresponded with explicit attitudes and cultural knowledge. Implicit attitudes were consistently, reliably and variably related to explicit attitudes, and were inconsistently and weakly related to cultural knowledge. More importantly, explicit attitudes accounted for the

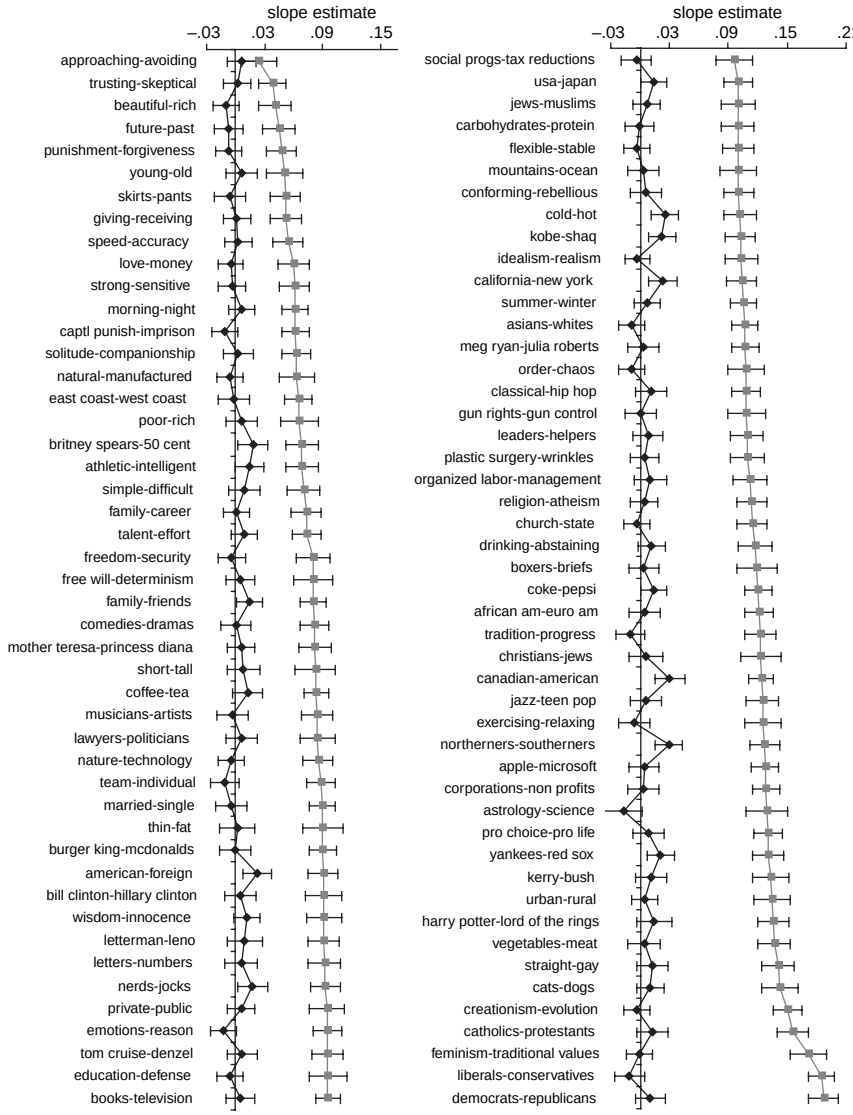

Figure 5. Multilevel modelling parameter estimates of explicit attitudes and cultural knowledge predicting IAT effects for 95 topics with 95% confidence intervals (Study 7, Model M_4).

relationship between implicit attitudes and cultural knowledge. In other words, cultural knowledge had little to no independent relationship with IAT effects. We interpret this result as revealing that the observed relations between cultural knowledge and the IAT are a function of attitudinal

processes such as incorporating cultural norms into one's own attitudes (Azjen & Fishbein, 2005), or naïve realism, using one's own attitudes in the estimate of the culture's attitudes (Ross, 1996). With evidence that cultural knowledge has a minimal independent relationship with the IAT, theoretical opining about whether cultural knowledge is a contaminant to measurement (e.g., Arkes & Tetlock, 2004; Gehring et al., 2003; Karpinski & Hilton, 2001; Olson & Fazio, 2004), or a meaningful component of the implicit attitude construct distinguishing it from explicit evaluation (e.g., Banaji, 2001; Banaji et al., 2004; Nosek, Banaji, & Greenwald, 2002b) might appear to be moot.

Is knowledge dead?

The accumulated evidence reported in this paper advances understanding of key relations among contemporary psychological constructs and measures. The results cannot be dismissed as a function of: (a) lack of meaningful interindividual variability in perceptions of cultural preferences; (b) invalid cultural indicators—for example, cultural knowledge measures predicted judgements of the likelihood of employment based racial discrimination, election predictions, and sales of foods; (c) low power—the study sample sizes were large providing very high powered tests, 100 samples had power in excess of .95 for detecting *r*s of .30 with $\alpha = .05$; (d) null relations—cultural knowledge was significantly related to the IAT in 37% (59/158) of the samples, and significantly related to explicit attitude measures in 59% (93/158) of the samples, and the IAT–knowledge relations were accounted for by explicit attitudes; and (e) narrow investigation of topics and knowledge—99 topics were investigated, and the studies examined a heterogeneous representation of cultural knowledge representing perceptions of media, historical, or societal portrayals of target concepts, and estimates of the average person's or most people's liking, warmth or preferences for the target objects.

These data are a significant challenge to the hypothesis that IAT performance is influenced by cultural knowledge—whether it is considered a contaminating influence or not. And yet, it would be inappropriate to draw a universal conclusion that cultural factors do not influence implicit attitudes. Rather, these data provide a basis of evidence against one form of implicit-culture relations—that cultural knowledge, perceptions of what others think, corresponds to the IAT independently of its relationship with explicit attitudes. Next, we raise three potential avenues for clarifying the relationship between culture and implicit attitudes, and then close with considerations about the meaningfulness of a distinction between *personal* attitudes and *extrapersonal* knowledge in implicit cognition.

Other implicit measures and process accounts

The present studies tested one implicit measure—the Implicit Association Test. It is well-known that implicit measures are a heterogeneous family of tools that have distinct procedures and likely engage a variety of cognitive processes (Bosson, Swann, & Pennebaker, 2000; Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005). It is possible that other implicit measures are sensitive to cultural knowledge where the IAT is not. While this is a logical possibility, it is notable that the IAT has been singled out as being particularly sensitive to cultural knowledge influences (Karpinski & Hilton, 2001; Olson & Fazio, 2004).

Other representations of cultural knowledge

The present data have the strength of a heterogeneous representation of cultural knowledge—and the effects were consistent whether the knowledge measures were considered individually or in aggregate. However, there may be untested forms of cultural or evaluative knowledge that will predict IAT effects. Olson and Fazio (2004), for example, suggested that evaluative knowledge could be from any source that does not contribute to one's evaluation. Knowledge assessment in these studies did not examine every possible knowledge source. If evidence for other sources of evaluative knowledge can be found to influence IAT effects, then it will be an interesting challenge to resolve why those sources *do* relate and cultural knowledge sources examined here *do not*. In other words, what cognitive architecture would selectively include or exclude forms of evaluative knowledge from influencing IAT performance, or implicit cognition more generally? And, why would people's beliefs about historical portrayals, cultural evaluations, or the opinions of most people be in the "excluded" category?

Cultural experience versus cultural knowledge

Given the theoretical positioning of implicit attitudes, it is rather surprising that the IAT–knowledge relation did not exist independent of explicit attitudes. Implicit attitudes are presumed to reflect one's experience in everyday life, and cultural knowledge presumably reflects some of that experience. These data suggest that measuring cultural knowledge is not a good way to capture the experience that is reflected in implicit evaluation.

To us, the most promising development from these data is the opportunity to advance a more specific identification of the implicit attitude construct, and refine the conceptual difference between cultural knowledge and cultural experience. Banaji (2001), for example, proposed that the impact of culture on individual minds may not be introspectively accessible. The associations

that develop between social concepts and attributes may exist distinctly from the explicit, reflective, deliberate processes that comprise conscious experience. A variety of dual-process models distinguish between implicit, impulsive or associative versus explicit, reflective, or propositional mental contents or processes (Chaiken & Trope, 1999; Gawronski & Bodenhausen, 2006; Smith & DeCoster, 1999; Strack & Deutsch, 2004; Wilson, Lindsey, & Schooler, 2000). These models postulate different degrees or means with which implicit and explicit experiences operate independently or interactively, and all provide some basis for understanding how considered assessments of cultural knowledge may be distinct from cultural experience, and the latter may be what manifests in implicit cognition.

Alternate views of distinguishing person and culture

Theoretical conceptions of implicit attitudes suggest that they are introspectively inaccessible reflections of previous experience (Banaji, 2001; Greenwald & Banaji, 1995; Wilson et al., 2000). The presence of correspondence between implicit and explicit measures introduces the possibility that at least some components of the construct are introspectively accessible, or at least introspectively predictable (see Nosek, 2005; Ranganath, Smith, & Nosek, in press). The present studies suggest that variation in IAT effects that is not shared with self-reported attitudes *is not* cultural knowledge, but it is not clear what this unique variation *is*. However, it is this mental stuff that eludes self-report—because we do not know that we have it, or because we are unwilling to say it—that makes implicit measures promising for learning more about the mind than its owner reports (Greenwald & Banaji, 1995).

If explicit attitudes could completely account for effects on implicit measures, then there would be little justification for a theoretical distinction between implicit and explicit attitudes—at least as they are reflected in the present generation of measures. Evidence in the construct development of implicit attitudes suggests that they are distinct but related to self-reported attitudes (Nosek, Greenwald, & Banaji, 2006a; Nosek & Smyth, 2007), there appear to be multiple moderating influences on the relationship between implicit and explicit evaluation (Hofmann et al., 2005; Nosek, 2005), and they appear to have distinct predictive validity (Poehlman, Uhlmann, Greenwald, & Banaji, 2007).

The challenge issued by the present data is to provide an integrative account of the observations that (a) explicit assessment of personal attitude, but not cultural knowledge, is related to implicit attitude measurement, and (b) implicit attitude measures retain unique variation that is independent of both explicit attitudes and knowledge. What is the leftover stuff, if it is

anything more than construct irrelevant method variance? The previous section speculated that it reflects cumulative experience that is either not accessible or not recognised as relevant for attitudinal self-report.

While theories posit that such experience is attitudinal (Banaji, 2001), it is possible that some experience influences implicit measurement but is inert in the everyday behaviour of the individual. If that were the case, then some evaluative associations might not be deemed attitudes because, functionally, they do not influence individual perception, judgement, or action. This would be consistent with hypotheses that there exists a meaningful distinction between personal and extrapersonal associations.

For example, Olson and Fazio (2004) distinguish personal versus extrapersonal associations by defining the latter as "associations that do not contribute to one's evaluation of an attitude object and thus do not become activated when one encounters the object but that are nevertheless available in memory" (Olson and Fazio, p. 653). Cultural knowledge and other extrapersonal associations can be identified in memory as distinct from personal associations, but because the former are not part of one's personal evaluation, they are not automatically activated by the attitude object. That way, they do not influence perception, judgement, and action, and thus should be considered a contaminant for attitude assessment.

The alternate view does not distinguish evaluative associations in memory as being personal or extrapersonal, a priori (Banaji, 2001; Banaji et al., 2004; Gawronski & Bodenhausen, 2006; Greenwald & Banaji, 1995; Nosek et al., 2002b; Strack & Deustch, 2004). Whether associative information is activated does not depend on it being tagged as "one's own" or not. Any association in memory, whatever its source, has the potential to be activated and influence perception, judgement, and action (Higgins, 1996). Whether an association is influential depends on the principles of availability, accessibility, salience, and applicability that have many varieties of causal influence such as repetition or recency of exposure, and may or may not have anything to do with the associations being identified as personal (see Andersen, Moskowitz, Blair, & Nosek, 2007; Higgins, 1996, for reviews).

This difference in theoretical positions is made clear with opposing interpretations of a result reported by Han, Olson, and Fazio (2006). Participants learned information about novel objects—that one was bad and the other was good. Then, some participants observed a video of children talking about the objects and giving "wrong" evaluative feedback—saying that they preferred the bad one. Participants reported that the children were providing silly answers, but their performance on the IAT showed sensitivity to the children's reports. The authors interpreted this effect as evidence that the IAT was sensitive to evaluative knowledge that does not comprise one's personal evaluation and would not influence the individual's evaluation or action related to the attitude

objects.[7] Our perspective interprets this as reflecting an important difference between implicit and explicit evaluation. Explicitly, the participants can recall that some evaluative information that they learned was inaccurate and should be discounted. Implicitly, however, both "accurate" and "inaccurate" exposures would produce associations in memory that can become active and influence subsequent processing.

In sum, the present studies do not eliminate the possibility that implicit measurement contains information that is extrapersonal. The results do suggest, however, that making a distinction between personal and extrapersonal must account for why cultural knowledge is not part of this distinction. Next, we elaborate our view that a taxonomy of associations as personal versus extrapersonal, or "mine" versus "not mine", is not useful for implicit cognition. We argue that identifying associations as "not mine" is a self-regulatory act and contributes to the distinction between explicit (propositional) judgement, where such acts are routine, and implicit (associative) operation, where they are not.

All our associations belong to us (implicitly)

An attitude is defined as the association between a concept and an evaluation that resides in memory (Fazio, Chen, McDonel, & Sherman, 1982). These associations are presumed to form based on experience, direct and indirect, with attitude objects (Eagly & Chaiken, 1998; Greenwald & Banaji, 1995). Once in memory, concept-evaluation associations may influence perception, attention, judgement, and action providing the basis for the prominence of attitudes in social psychological theory and research.

The presence of concept-evaluation associations does not mean that they will influence cognitive processing in all cases. Drawing on Higgins' (1996) distinction between associative information that is available (stored in memory) versus accessible (its activation potential), Eagly and Chaiken (1998) point out that an attitudinal response is dynamic and can draw on different aspects of the available concept-evaluation associations (see also Wilson & Hodges, 1992). What associative information is activated and influential will depend on its availability, accessibility, salience, and applicability (Higgins, 1996).

[7] Han et al. (2006) also found that an alternative form of the IAT was not sensitive to the manipulation and interpreted this as evidence that their alternative version was a more construct valid assessment of attitudes because it was not influenced by the children's statements. We disagree with the theoretical expectation that the statements should not influence implicit attitudes, described next, and we have suggested that the different effects by IAT versions is at least partly a consequence of the alternative being influenced by explicit evaluation processes (see Nosek & Hansen, 2007).

The IAT, like other implicit measures, is thought to measure concept-evaluation associations that have developed from experience through mechanisms such as classical conditioning (Olson & Fazio, 2001). Culturally bound experience is comprised by nationality, state, city, neighbourhood, school, family, birth order, friend, gender, ethnicity, age, social class, spoken language, occupation, and any number of other social categories and contexts. Implicit evaluations are presumed to reflect variations in those experiences. Experience may be culturally bound or culturally independent, but that distinction is irrelevant for implicit evaluation. What is important for implicit evaluation is that experience must happen, associations must form, and those associations must be available.

Where we differ from Karpinski and Hilton (2001) and Arkes and Tetlock (2004) is that we argue that endorsement, especially in the context of implicit cognition, is irrelevant for information to be a measure of individual attitude and predict individual behaviour (Banaji et al., 2004). Associative representations reflect accumulated experience with attitude objects regardless of whether those experiences are accepted or rejected as true (Gawronski & Bodenhausen, 2006; Smith & DeCoster, 1999). Returning to the example that opened the paper, the job interviewer may have had a lifetime of negative cultural exposures to Hispanics that were recorded as associations of Hispanic with bad or incompetent, even if she or he were unaware of its presence, and honestly rejected the association as false. The association does not know if it is true or false, only an explicit, propositional judgement can render such a verdict (Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004).

Where we differ from Olson and Fazio (2004) is that we argue that any evaluative information, no matter how it was learned, is potentially attitudinal and influential for individual thinking, feeling, and acting. Declaring that some information in our own heads is not *personal* may inappropriately focus attention on the source of the information (where we learned it) rather than the consequences of the information (what we do with it). It is in the presence and consequences of information, not the origins, that ownership is established.

Distinguishing myself from my knowledge

The preceding discussion might appear to suggest that humans are slaves to their experience, and that knowing something is akin to believing it. Humans do appear to represent and believe information in a singularly Spinozan process (Gilbert, 1991). But, humans also have the remarkable ability to unbelieve things that they once thought and believed. Distinguishing knowledge that is "mine" from "just the stuff that I know" is where explicit cognition has a decided advantage over implicit cognition.

A luxury of conscious processes is that we get to decide whether we believe the information that bubbles up from memory (Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004). Stereotypes about racial, gender, age, or political groups can come to mind and be accepted or rejected. Also, we can invoke higher-order principles for informing on our judgements and actions toward group members, such as "treat others not by the colour of their skin but the content of their character". These explicit processes provide opportunities to effortfully correct associations present in the culture or our own mind that may conflict with the ways in which we want to perceive, judge, and act toward others, or attitude objects in general.

Implicit or automatic processes that operate outside of conscious awareness or conscious control afford fewer corrective mechanisms. The information available in memory, whatever the source and whether personally accepted or rejected, can influence perception, judgement, and action whenever it becomes actively involved in cognitive processing. Whether certain information becomes influential may be determined by multiple processes such as chronic goals or motivations (Devine, Plant, Amodio, Harmon-Jones, & Vance, 2002; Moskowitz, Gollwitzer, Wasel, & Schaal, 1999) or by the degree to which information is well-learned, situationally relevant, or immediately accessible (Higgins, 1996).

Avoiding the influence of concept-evaluation associations that we would prefer not to claim as our own requires awareness of their presence, capacity to exert control over their expression, and the knowledge or skill to correct for their influence. This may not be simple. Consider stereotype threat in which members of stereotyped groups show performance decrements in the stereotyped domain when the relevant stereotype or social identity is activated (Steele & Aronson, 1995; Steele, Spencer, & Aronson, 2002). The impact of the stereotype knowledge need not be chronically accessible, personally endorsed, or even available to conscious awareness in order to have its insidious impact—it need only be activated (Dijksterhuis, Aarts, Bargh, & van Knippenberg, 2000; Dijksterhuis & van Knippenberg, 1998; Shih, Pittinsky, & Ambady, 1999; Steele et al., 2002). It might be more comfortable to say that those stereotypes are "not mine" because they are explicitly disavowed and a threat to self. Nonetheless, those stereotypes are in mind and influential, making them unavoidably, even undesirably, one's own.

The selves that we are and the selves we intend to be are both *us*, and sometimes they do not agree. One might say that humans are large, containing multitudes. Full recognition of this fact raises serious questions for important issues of responsibility, culpability, and intentionality. When should organisms (even human ones) be held responsible for their actions? What role should intentionality play in drawing the line between the responsible agent and the causal, but not responsible, agent? These issues

reach far beyond the penultimate paragraph of a paper, but are ones that psychologists, ethicists, and legal analysts must continue to scrutinise.

All concept-evaluation associations that are available in memory have the potential to influence processing, perception, judgement, and action—so, all such associations are attitudinal. Efforts to understand when, why, and how various aspects of those attitudes will have influence should keep psychologists busy and ensure that attitudes remain an indispensable construct.

Manuscript received 15 January 2007
Revised manuscript received 12 April 2007
Manuscript accepted 13 April 2007
First published online 6 September 2007

REFERENCES

Andersen, S. M., Moskowitz, G. B., Blair, I. V., & Nosek, B. A. (2007). Automatic thought. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 138–175). New York: Guilford Press.

Arkes, H. R., & Tetlock, P. E. (2004). Attributions of implicit prejudice, or "would Jesse Jackson 'fail' the Implicit Association Test?". *Psychological Inquiry*, *15*, 257–278.

Azjen, I., & Fishbein, M. (2005). The influence of attitudes on behavior. In D. Albarracín, B. Johnson, & M. P. Zanna (Eds.), *The handbook of attitudes* (pp. 173–222). Mahwah, NJ: Lawrence Erlbaum Associates, Inc.

Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger, J. S. Nairne, I. Neath, & A. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117–150). Washington, DC: American Psychological Association.

Banaji, M. R., Nosek, B. A., & Greenwald, A. G. (2004). No place for nostalgia in science: A response to Arkes & Tetlock. *Psychological Inquiry*, *15*, 279–310.

Bargh, J. A. (1996). Automaticity in social psychology. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 169–183). New York: Guilford Press.

Bosson, J. K., Swann, W. B., & Pennebaker, J. W. (2000). Stalking the perfect measure of implicit self-esteem: The blind men and the elephant revisited? *Journal of Personality and Social Psychology*, *79*, 631–643.

Chaiken, S., & Trope, Y. (1999). *Dual-process theories in social psychology.* New York: Guilford Press.

Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, *89*, 469–487.

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5–18.

Devine, P. G., Plant, E. A., Amodio, D. M., Harmon-Jones, E., & Vance, S. (2002). The regulation of explicit and implicit racial bias: The role of motivations to respond without prejudice. *Journal of Personality and Social Psychology*, *82*(5), 835–848.

Dijksterhuis, A., Aarts, H., Bargh, J. A., & van Knippenberg, A. (2000). On the relation between associative strength and automatic behavior. *Journal of Experimental Social Psychology*, *36*, 531–544.

Dijksterhuis, A., & van Knippenberg, A. (1998). The relation between perception and behavior, or how to win a game of Trivial Pursuit. *Journal of Personality and Social Psychology*, *74*, 865–877.

Eagly, A. H., & Chaiken, S. (1998). Attitude structure and function. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 1, pp. 269–322). Boston: McGraw-Hill.

Fazio, R. H., Chen, J., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude–behavior consistency, and the strength of the object-evaluation association. *Journal of Experimental Social Psychology*, *18*, 339–357.

Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027.

Fiske, A. P., Kitayama, S., Markus, H. R., & Nisbett, R. E. (1998). The cultural matrix of social psychology. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 2, pp. 915–981). Boston: McGraw-Hill.

Gawronski, B., & Bodenhausen, G. V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, *132*, 692–731.

Gehring, W. J., Karpinski, A., & Hilton, J. L. (2003). Thinking about interracial interactions. *Nature Neuroscience*, *6*, 1241–1243.

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, *46*, 107–119.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*(1), 4–27.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*, 1464–1480.

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the Implicit Association Test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, *85*, 197–216.

Han, H. A., Olson, M. A., & Fazio, R. H. (2006). The influence of experimentally created extrapersonal associations on the Implicit Association Test. *Journal of Experimental Social Psychology*, *42*, 259–272.

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York: Guilford Press.

Hofmann, W., Gschwendner, T., Nosek, B. A., & Schmitt, M. (2005). What moderates implicit–explicit consistency? *European Review of Social Psychology*, *16*, 335–390.

Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, *81*, 774–788.

Moskowitz, G. B., Gollwitzer, P. M., Wasel, W., & Schaal, B. (1999). Preconscious control of stereotype activation through chronic egalitarian goals. *Journal of Personality and Social Psychology*, *77*(1), 167–184.

Nosek, B. A. (2005). Moderators of the relationships between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, *134*, 565–584.

Nosek, B. A. (2007). Implicit–explicit relations. *Current Directions in Psychological Science*, *16*, 65–69.

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002a). Harvesting intergroup implicit attitudes and beliefs from a demonstration Web site. *Group Dynamics*, *6*(1), 101–115.

Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (2002b). Math = me, me = female, therefore math ≠ me. *Journal of Personality and Social Psychology*, *83*(1), 44–59.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin*, *31*, 166–180.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2006a). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 365–292). New York: Psychology Press.

Nosek, B. A., & Hansen, J. J. (2007). *Testing different versions of the Implicit Association Test*. Unpublished manuscript: University of Virginia.

Nosek, B. A., & Smyth, F. L. (2007). A multitrait-multimethod validation of the Implicit Association Test: Implicit and explicit attitudes are related but distinct constructs. *Experimental Psychology*, *54*, 14–29.

Nosek, B. A., Smyth, F. L., Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., et al. (2007). *Pervasiveness and correlates of implicit attitudes and stereotypes. European Review of Social Psychology, 18,* 36–88.

Nosek, B. A., Sutin, E., Hansen, J. J., Wu, L., Sriram, N., Smyth, F. L., et al. (2006b). *Project implicit: Advancing theory and evidence with technical and methodological innovation*. Unpublished manuscript: University of Virginia.

Olson, M. A., & Fazio, R. H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science*, *12*, 413–417.

Olson, M. A., & Fazio, R. H. (2004). Reducing the influence of extra-personal associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology*, *86*(5), 653–667.

Poehlman, T. A., Uhlmann, E., Greenwald, A. G., & Banaji, M. R. (2007). *Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity.* Unpublished manuscript: Yale University.

Ranganath, K. A., Smith, C. T., & Nosek, B. A. (in press). Distinguishing automatic and controlled components of attitudes from direct and indirect measurement. *Journal of Experimental Social Psychology.*

Ross, L. (1996). Naïve realism in everyday life: Implications for social conflict and misunderstanding. In E. S. Reed, E. Turiel, & T. Brown (Eds.), *Values and knowledge* (pp. 103–135). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Sears, D. O. (2004). A perspective on implicit prejudice from survey research. *Psychological Inquiry*, *15*, 293–297.

Shih, M., Pittinsky, T. L., & Ambady, N. (1999). Stereotype susceptibility: Identity salience and shifts in quantitative performance. *Psychological Science*, *10*, 80–83.

Smith, E. R., & DeCoster, J. (1999). Associative and rule-based processing: A connectionist interpretation of dual-process models. In S. Chaiken & Y. Trope (Eds.), *Dual-process theories in social psychology* (pp. 323–338). New York: Guilford Press.

Snijders, T. A. B., & Bosker, R. J. (1999). *Multilevel analysis.* London: Sage.

Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology*, *69*, 797–811.

Steele, C. M., Spencer, S. J., & Aronson, J. (2002). Contending with group image: The psychology of stereotype and social identity threat. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 34, pp. 379–440). San Diego, CA: Academic Press.

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, *8*, 220–247.

Wegner, D. M. (2002). *The illusion of conscious will*. Cambridge, MA: MIT Press.

Wegner, D. M., & Bargh, J. A. (1998). Control and automaticity in social life. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology* (4th ed., Vol. 1, pp. 446–496). Boston: McGraw-Hill.

Wilson, T. D. (2002). *Strangers to ourselves: Discovering the adaptive unconscious*. Cambridge, MA: Harvard University Press.

Wilson, T. D., & Hodges, S. D. (1992). Attitudes as temporary constructions. In A. Tesser & L. Martin (Eds.), *The construction of social judgment* (pp. 37–65). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, *107*(1), 101–126.