Research Article

Does the Name-Race Implicit Association Test Measure Racial Prejudice?

Don van Ravenzwaaij, Han L. J. van der Maas, and Eric-Jan Wagenmakers

University of Amsterdam, The Netherlands

Abstract. Research using the Implicit Association Test (IAT) has shown that names labeled as Caucasian elicit more positive associations than names labeled as non-Caucasian. One interpretation of this result is that the IAT measures latent racial prejudice. An alternative explanation is that the result is due to differences in in-group/out-group membership. In this study, we conducted three different IATs: one with same-race Dutch names versus racially charged Moroccan names; one with same-race Dutch names versus racially neutral Finnish names; and one with Moroccan names versus Finnish names. Results showed equivalent effects for the Dutch-Moroccan and Dutch-Finnish IATs, but no effect for the Finnish-Moroccan IAT. This suggests that the name-race IAT-effect is not due to racial prejudice. A diffusion model decomposition indicated that the IAT-effects were caused by changes in speed of information accumulation, response conservativeness, and non-decision time.

Keywords: Implicit Association Test, diffusion model, racial prejudice, in-group/out-group membership

People aren't perfect. One of our imperfections is that we may be prejudiced, for example against those that do not belong to our age-cohort (e.g., "old people smell bad") or to our race (e.g., "Moroccans are aggressive"). In order to measure such prejudices without the limitations of self-report (e.g., Nosek, Greenwald, & Banaji, 2007), Greenwald and colleagues introduced the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998; for reviews, see Fazio & Olson, 2003; Nosek et al., 2007). In the IAT, participants respond with one of two buttons to one of two types of *attribute concepts* (e.g., positive and negative words) and to one of two types of *target concepts* (e.g., same-race Dutch names and different-race Moroccan names). After a series of practice blocks, participants are confronted with a compatible and an incompatible block of trials. In both blocks, participants have to quickly classify an attribute concept or a target concept. In the *compatible* block, positive attribute concepts and Dutch names require the same button press, whereas in the *incompatible* block, positive attribute concepts and Moroccan names require the same button press. In general, people tend to perform faster and more accurately in the compatible block than in the incompatible block. This so-called *IAT-effect* is thought to reflect the presence of preexisting associations between the attribute and target concepts (e.g., Nosek et al., 2007).

When IAT-effects are obtained with racial stimuli as target concepts, these effects may be interpreted as indicators of racial prejudice (but see Nosek et al., 2007 for cautionary remarks). This is a strong claim with profound societal consequences, and therefore the first goal of this study is to carefully assess an alternative explanation of the racial IAT-effects.

This alternative explanation of the racial IAT-effect holds that participants associate positive stimuli more easily with their in-group, and negative stimuli more easily with out-groups. In this account in-group/out-group membership – not racial prejudice – causes the IAT-effect. This explanation is difficult to investigate, as same-race names tend to be in-group names and different-race names tend to be out-group names. However, there is a crucial difference between racial prejudice and in-group/out-group membership as a cause for IAT-effects. For example, suppose a soccer player shows an IAT-effect with soccer players and tennis players as target concepts. This IAT-effect is probably based on group membership, not prejudice.

In one study on the effect of in-group/out-group membership, Popa-Roch and Delmas (in press) investigated the role of self-inclusion in either the in-group or the out-group on the IAT-effect. The authors administered two IATs. First, an IAT in which the target categories were "French and Me" (the in-group) versus "North African" (the out-group). Second, an IAT in which the target categories were "French" and "North-African and Me." For this second IAT, the standard IAT-effect completely vanished. In a second study, the association with either the in-group or the out-group was established prior to administration of a standard IAT. Once again, associating oneself with the out-group caused the IAT-effect to disappear.

In a different study, Blair, Judd, Havranek, and Steiner (2010) manipulated both in-group and out-group membership. Specifically, the authors administered a White-Black

DOI: 10.1027/1618-3169/a000093

IAT and a White-Latino IAT to three groups of participants: a Caucasian group, an African-American group, and a Latino group. Blair et al. (2010) found that the White-Black IAT-effect was largest for the Caucasian group and smallest for the African-American group, whereas the White-Latino IAT-effect was largest for the Caucasian group and smallest for the Latino group. For the Caucasian group, there was no difference in the White-Black IAT-effect and the White-Latino IAT-effect. However, both African-Americans and Latinos may be racially prejudiced in the US, making it hard to disentangle the effects of racial prejudice and in-group/out-group membership.

Thus, the goal of this study is to disentangle the effects of racial prejudice and in-group/out-group membership on the IAT-effect. In the empirical study presented below, we introduce a new method to test the racial-prejudice account. Specifically, we replace the racially charged (i.e., Moroccan) out-group names with racially neutral (i.e., Finnish) out-group names. If an IAT-effect would still be obtained with Finnish names, this would be evidence against racial prejudice as an exclusive account of the name-race IAT-effect.

Experiment

Three name IATs were used in a between-subjects design: a Dutch-Moroccan (Dut-Mor), a Dutch-Finnish (Dut-Fin), and a Finnish-Moroccan IAT (Fin-Mor, for which the block in which Finnish names and positive words required the same key press was labeled "compatible"). For the participants in our experiment (i.e., Dutch undergraduates), both Moroccan people and Finnish people are national out-groups. In addition, Moroccan people form a racial minority in the Netherlands (2.0%; Centraal Bureau voor de Statistiek, 2008) that are commonly discriminated against (Dolfing & van Tubergen, 2005). There are very few Finnish people living in the Netherlands (0.02%; Centraal Bureau voor de Statistiek, 2008), and they are not commonly discriminated against.[1]

If in-group/out-group membership causes the IAT-effect, effects should be present for both the Dut-Mor and the Dut-Fin IATs, but not for the Fin-Mor IAT. If, on the other hand, implicit racial prejudice causes the IAT-effect, effects should be present for both the Dut-Mor and the Fin-Mor IATs, but not for the Dut-Fin IAT. If the IAT-effect is due to a combination of in-group/out-group membership and racial prejudice, then all three IATs should yield an effect, with the Dut-Mor IAT-effect being the most pronounced.

Method

Participants

Sixty Caucasian Dutch undergraduate students from the University of Amsterdam (39 female) participated for course credit. Participants were randomly assigned to one of the experimental conditions.

Materials

Stimulus materials consisted of 100 positive words, 100 negative words, 100 Dutch names, 100 Moroccan names, and 100 Finnish names (50 female names per language). The set of positive and negative words had been collected for a previous study (Zeelenberg, Wagenmakers, & Rotteveel, 2006), where their valence had been confirmed by ratings. A candidate set of names was constructed to be typical for their country.[2]

Examples of the Moroccan, Finnish, and Dutch names used are "Faiza," "Jarkko," and "Hein," respectively. Examples of positive and negative words used are "engel" (angel) and "crisis" (crisis).

Design

In all three IATs, half of the participants were first confronted with the compatible block, and the other half were first confronted with the incompatible block.

In a between-subjects design, each IAT was administered to 20 participants. All three IATs consisted of five consecutive blocks. Critically, blocks 3 and 5, each containing 320 trials, were the compatible and incompatible blocks, which were counterbalanced. The compatible and incompatible blocks used different stimuli; within each block all stimuli were presented twice. Blocks 1, 2, and 4, each containing 40 trials, were practice blocks (e.g., Greenwald et al., 1998).

Procedure

Participants were instructed to respond as fast as possible without making too many mistakes. Each block started with instructions that described the assignment of the two response keys (i.e., "z" and "m") to the stimulus categories. The response key assignment was continually displayed at the bottom of the screen.

All stimuli were displayed in black letters on a white screen background, vertically and horizontally centered in the display. Stimuli remained on screen until the participant made a response. The response-stimulus interval was 750 ms. After any incorrect response, the word "FOUT!" (Dutch for "ERROR!") immediately replaced the stimulus for 300 ms. No feedback was provided following a correct response. All stimuli, names and attributes, were selected randomly until the entire list of stimuli for a given block had been exhausted. Participants took a short self-paced break between each of the five blocks.

[1] See the online appendix for corroborating evidence, available at http://www.donvanravenzwaaij/com, tab "Research & Codes."
[2] For more information on stimulus selection and a complete stimulus list, see the online appendix.

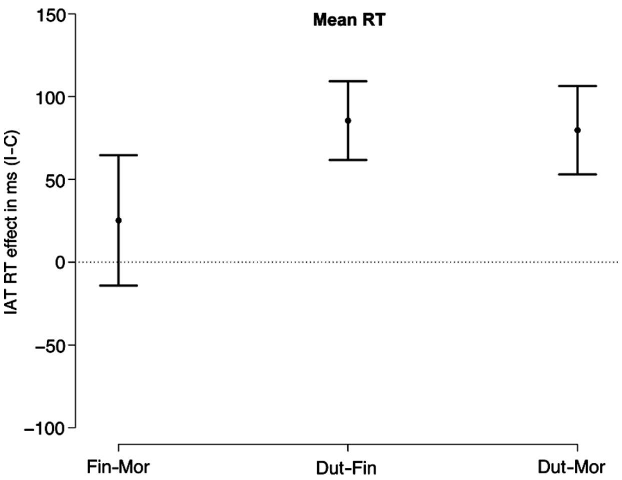

Figure 1. The mean IAT-effects over participants on mean RT for the three different IATs. Error bars represent 95% confidence intervals of the mean.

Results

Four participants were replaced, three for having error rates over 30% and one for having a mean RT over 1 s. For each participant, we removed all RTs below 200 ms. We also removed all RTs slower than 3.5 *SD* from the mean in an iterative procedure. This resulted in a loss of 3.8% of the trials.

We have calculated split-half reliabilities for the three IATs a thousand times using a bootstrap procedure, averaged the results, and found that the Fin-Mor IAT had a reliability coefficient of .92, the Dut-Fin a reliability coefficient of .79, and the Dut-Mor a reliability coefficient of .87.

For the remainder of the statistical analyses, we report Bayesian posterior probabilities in addition to conventional *p* values. When we assume, for fairness, that the null hypothesis and the alternative hypothesis are equally plausible *a priori,* a default Bayesian *t* test (Rouder, Speckman, Sun, Morey, & Iverson, 2009) allows one to determine the *posterior* plausibility of the null hypothesis and the alternative hypothesis. We denote the posterior probability for the null hypothesis as $p_{H_0}^{\text{Bayes}}$. When, for example, $p_{H_0}^{\text{Bayes}} = .9$, this means that the plausibility for the null-hypothesis has increased from .5 to .9. Posterior probabilities avoid the problems that plague *p* values, allow one to directly quantify evidence in favor of the null-hypothesis, and arguably relate more closely to what researchers want to know (e.g., Wagenmakers, 2007).

Figure 1 shows the mean IAT-effects over participants on mean RT for the Fin-Mor, the Dut-Fin, and the Dut-Mor IATs for names and attributes combined. The detailed statistical analyses can be found in Table 1. As the table shows, there is a non-zero IAT effect for the Dut-Fin and the Dut-Mor IATs, but not for the Fin-Mor IAT. Also, both the IAT-effects for the Dut-Fin and the Dut-Mor IAT differ significantly from that of the Fin-Mor IAT, but not from each other.

The mean RT results cast doubt on the racial-prejudice hypothesis, according to which a non-zero IAT-effect for the Fin-Mor IAT and a difference between the Dut-Fin and the Dut-Mor IATs should have been found. However, Greenwald, Nosek, and Banaji (2003) discussed problems with the conventional RT analysis and introduced a new means of calculating the IAT-effect: the *D*-score. The *D*-score is calculated by subtracting the mean RT of the compatible block by the mean RT of the incompatible block (with outliers removed) and dividing the resulting value by the pooled standard deviation. We calculated *D*-score IAT-effects for the three different IATs, based on the data with the same criterion for exclusion of outliers as in the RT analyses. The results are qualitatively identical to the mean RT results.[3]

Interim Conclusion

The results of the presented study are not in line with the racial-prejudice hypothesis; there is no IAT-effect for Finnish versus Moroccan names, but there is an IAT-effect for Dutch versus Finnish and Dutch versus Moroccan names. Crucially, the Dutch versus Finnish and Dutch versus Moroccan IAT-effects do not differ from each other, whereas both do differ from the Finnish-Moroccan IAT-effects.

Although informative, the present analysis is limited in several ways (e.g., Wagenmakers, 2009; for an IAT-specific critique, see Blanton & Jaccard, 2008; Blanton, Jaccard, Gonzales, & Christie, 2006; Klauer, Voss, Schmitz, & Teige-Mocigemba, 2007). For instance, the analysis focuses on mean RT and ignores all other information provided by the RT distributions; the analysis ignores the interplay between RT and error rate; and the analysis is not motivated by any substantive theory, which means that the results do not speak directly to the details of the underlying psychological processes. For example, the fact that participants respond more slowly in the incompatible block than in the compatible block could reflect a lower rate of information processing, or it could reflect an increase in response caution, or some combination of these factors. The standard method of analysis cannot address these fundamental questions about the origin of the IAT-effect.

In order to address the limitations of the traditional analysis, we reanalyze the IAT data with the Ratcliff diffusion model, one of the most successful models for RT and accuracy (Klauer et al., 2007; Ratcliff, 1978; van Ravenzwaaij & Oberauer, 2009).

The Ratcliff Diffusion Model

In the diffusion model for speeded two-choice tasks (Ratcliff, 1978), stimulus processing is conceptualized as a noisy accumulation of evidence over time. A response is initiated when the accumulated evidence reaches a predefined evidence boundary (Figure 2).

[3] For details, see the online appendix.

Table 1. Test statistics for the IAT-effects on mean RT. First three columns: 19 *df*, last three columns: 38 *df*

		Fin-Mor	Dut-Fin	Dut-Mor	F-M vs. D-F	F-M vs. D-M	D-F vs. D-M
Mean RT	t	1.34	7.53	6.25	2.74	2.40	.34
	p	.20	< .05	< .05	< .05	< .05	.74
	$p_{H_0}^{\text{Bayes}}$	.72	.00	.00	.17	.28	.80

The four key components of the diffusion model are (1) the speed of information processing, quantified by drift rate v; (2) response caution, quantified by boundary separation a; (3) a priori bias, quantified by starting point z; and (4) non-decision time, quantified by T_{er}.

The model assumes that the decision process starts at z, after which information is accumulated with a signal-to-noise ratio that is governed by drift rate v. Values of v near zero produce long RTs and high error rates. Boundary separation a determines the speed-accuracy tradeoff; lowering a leads to faster RTs at the cost of a higher error rate. Together, these parameters generate a distribution of decision times *DT*. The observed RT, however, also consists of stimulus-nonspecific components such as response preparation and motor execution, which together make up non-decision time T_{er}. The model assumes that T_{er} simply shifts the distribution of *DT*, such that $RT = DT + T_{er}$ (Luce, 1986). The model specification is completed by including parameters that specify across-trial variabilities in drift rate, starting point, and non-decision time (Ratcliff & Tuerlinckx, 2002).

The diffusion model has been applied to the IAT before (Klauer et al., 2007). In that study, Klauer et al. (2007) reported effects on drift rate, boundary separation, and non-decision time (i.e., in the compatible block, participants processed information more quickly, were less cautious, and had shorter non-decision times). Here we also use a diffusion model analysis to determine if the locus of the IAT-effects found in our experiment is in drift rate, boundary separation, and/or non-decision time.

Diffusion Model Parameter Estimates

We used the fast-dm program (Voss & Voss, 2007, 2008) to fit the diffusion model to the individual data. We fit a total of six different models on the data with RT outliers excluded. The models differed from one another in the parameters that were fixed over conditions. We have selected models for which parameters for all 60 participants were estimated satisfactorily. For models that satisfied this criterion, we chose the simplest one (i.e., the one with most parameters constrained). The best fitting, most parsimonious model was a model that estimated four drift rates and four non-decision times, but two boundary separations. Drift rates and non-decision times were estimated for names and attributes separately, both in the compatible block and in the incompatible block.[4] Boundary separations were estimated for names and attributes combined in the compatible block and in the incompatible block. The starting point

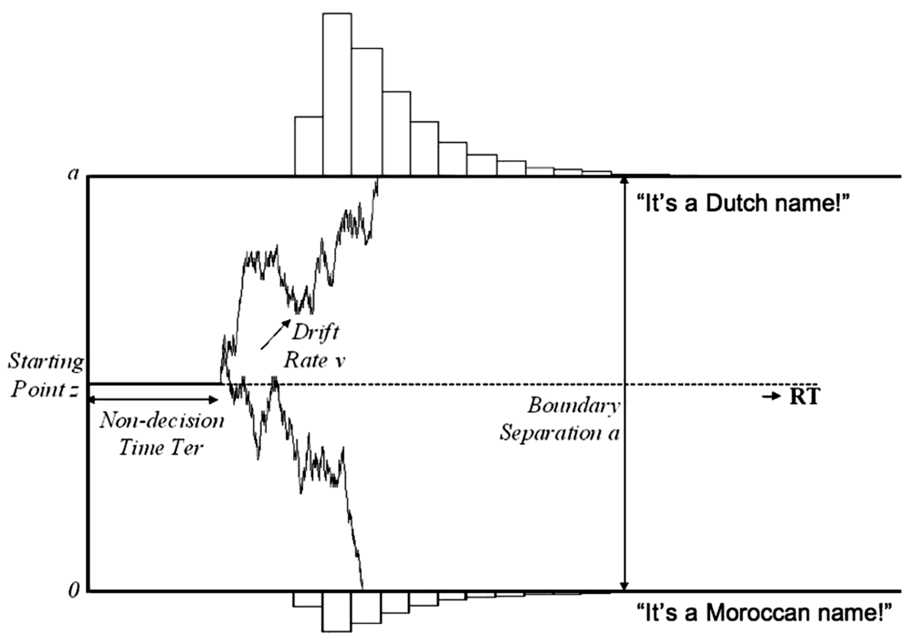

Figure 2. The diffusion model and its parameters. See text for details.

[4] For the analyses, no appreciable differences were found for the names and attribute parameters, so we present averaged parameters for presentational convenience.

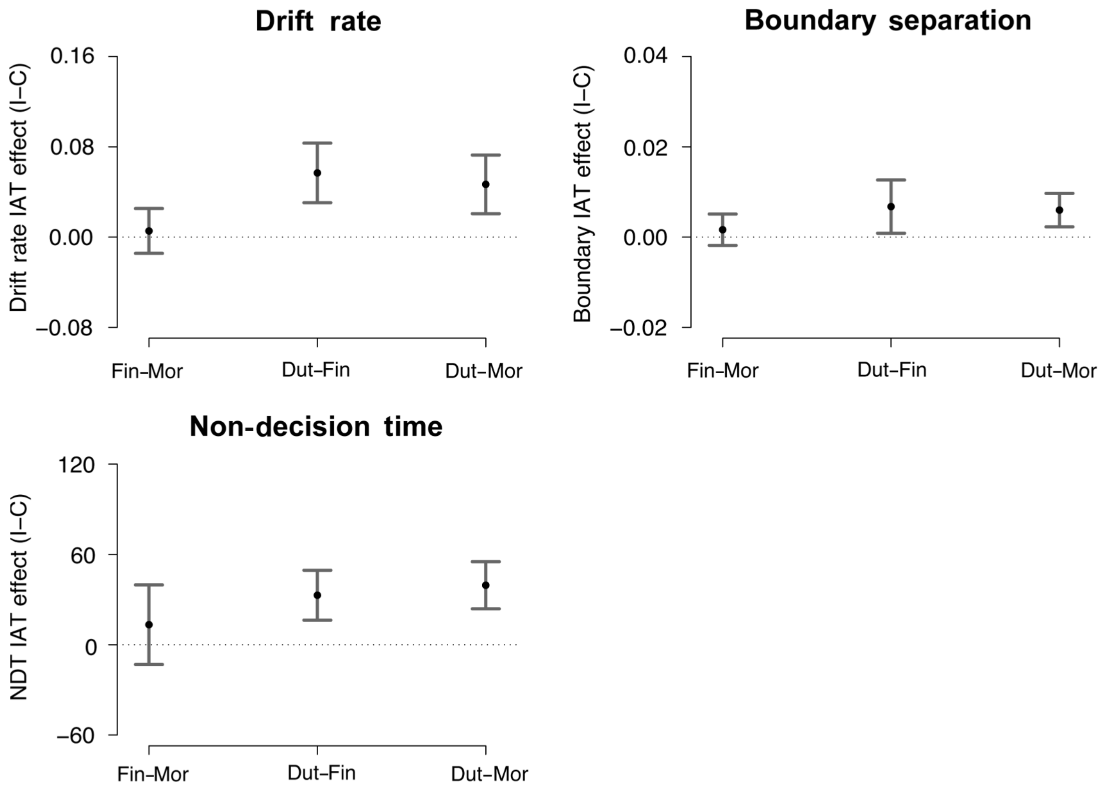

Figure 3. The mean IAT-effects over participants on drift rate (top-left panel), boundary separation (top-right panel), and non-decision time (bottom-left panel) for each of the three different IATs. The figure shows IAT-effects for the names and the attributes combined. Error bars represent 95% confidence intervals of the mean.

parameter was fixed to be half of boundary separation. Our chosen model fit the data well for each of the 60 participants. As a frame of reference, fixing non-decision time over conditions resulted in a poor model fit for 24 of the 60 participants. Other models did not yield qualitatively different results. Graphical summaries of the model fit and estimates for all diffusion model parameters can be found in the sections "Delta Plots" and "Diffusion Model Estimates," presented in the online appendix.

The top-left panel of Figure 3 shows the mean IAT-effects on *drift rate* for the Fin-Mor, the Dut-Fin, and the Dut-Mor IATs. The detailed Statistical analyses can be found in the top three rows of Table 2. As the table shows, there is a non-zero IAT-effect for the Dut-Fin and the Dut-Mor IATs, but not for the Fin-Mor IAT. Also, both the IAT-effects for the Dut-Fin and the Dut-Mor IAT significantly differ from the Fin-Mor IAT, but not from each other.

The top-right panel of Figure 3 shows the mean IAT-effects on *boundary separation* for the three different IATs. The middle three rows of Table 2 show that there is a non-zero IAT-effect for the Dut-Fin and the Dut-Mor IATs, but not for the Fin-Mor version of the IAT. However, none of the three IAT-effects significantly differ from each other. Also, the Bayesian posterior probabilities show that the effects are ambiguous.

The bottom-left panel of Figure 3 shows the mean IAT-effects on *non-decision time* for the three different IATs. The bottom three rows of Table 2 show non-zero IAT-effects for the Dut-Fin and the Dut-Mor IATs, but not for the Fin-Mor IAT. Once again, none of the three IAT-effects significantly differ from each other. The Bayesian posterior probabilities are very clear on the non-zero IAT-effects, but ambiguous on the difference between IAT-effects.

In sum, a diffusion model analysis shows that the IAT-effects originate from a combination of sources. The drift rate effect indicates that people process information faster in the compatible block than in the incompatible block. The non-decision time effect suggests that in the compatible block, people require less time to encode the stimuli or to map their decisions onto the response keys. The tentative boundary separation effect suggests that people are less cautious in the compatible block than in the incompatible block. These effects hold for the Dut-Fin and the Dut-Mor IATs, but not for the Fin-Mor IAT. This confirms the results from the analyses on mean RT.

Discussion

Our results show that in the name-race IAT, the IAT-effect may be mistakenly attributed to the presence of an implicit racial prejudice. In our experimental design, we included same-race Dutch names, racially charged Moroccan names, and racially neutral Finnish names. The results showed no

Table 2. Test statistics for the IAT-effects on drift rate (top three rows), boundary Separation (middle three rows), and non-decision time (bottom three rows). First three columns: 19 *df*, last three columns: 38 *df*

		Fin-Mor	Dut-Fin	Dut-Mor	F-M vs. D-F	F-M vs. D-M	D-F vs. D-M
Drift rate	t	.57	4.51	3.76	3.26	2.64	.57
	p	.57	< .05	< .05	< .05	< .05	.57
	$p_{H_0}^{\text{Bayes}}$	.83	.01	.03	.06	.20	.79
Boundary Separation	t	.98	2.40	3.38	1.57	1.79	.23
	p	.34	< .05	< .05	.13	.08	.82
	$p_{H_0}^{\text{Bayes}}$	.79	.34	.07	.60	.52	.81
Non-decision Time	t	1.05	4.15	5.28	1.31	1.79	.61
	P	.31	< .05	< .05	.20	.08	.55
	$p_{H_0}^{\text{Bayes}}$	.78	.02	.00	.67	.53	.79

effect when Moroccan names were contrasted with Finnish names, and an equivalent effect when Dutch names were contrasted with either Moroccan or with Finnish names. This suggests that the racially charged Moroccan names were processed in a similar fashion as the racially neutral Finnish names.

In an extension of the results by Klauer et al. (2007), we applied a diffusion model decomposition and found that the IAT-effects for the Dutch-Moroccan and the Dutch-Finnish IATs could be attributed to several factors. In the compatible block, participants processed information faster, displayed less response conservativeness, and spent less time on peripheral processing in the compatible block than in the incompatible block. This last non-decision time finding may seem counterintuitive, and it is easy to believe that compatibility effects would reside exclusively in stimulus-specific components such as drift rate. However, Klauer et al. (2007) also found that participants had lower non-decision times in the compatible block than in the incompatible block. They propose that this effect is due to *task-set switching*. According to the task-set switching account, participants can answer intuitively in the compatible block, both for the target and the attribute concepts, as both responding on the basis of category membership and responding on the basis of evaluation lead to the same response. For the incompatible block however, task switching is necessary to answer accurately. For the attribute concepts, responding on the basis of evaluation is appropriate. For the target concepts however, responding should now be on the basis of category membership. The resulting switch in response time should cause a longer non-decision time. Klauer et al. (2007) suggested that the excess time for peripheral processing reflects excess preparatory activity that is required to retrieve the correct response mapping.

The results of this study are also compatible with a name familiarity explanation of the IAT-effect (e.g., Dasgupta, McGhee, Greenwald, & Banaji, 2000; Ottaway, Hayden, & Oakes, 2001; Rothermund & Wentura, 2004; Rudman, Greenwald, Mellott, & Schwartz, 1999): Same-race names are generally more familiar than different-race names, and this can cause same-race names to elicit more positive associations, for example through *mere exposure* (e.g., Zajonc, 1980). We conducted a second study in order to test the alternative name familiarity explanation. Additionally, since our results hinge on a null effect for the Finnish-Moroccan IAT, we sought to replicate this finding with a group of 104 participants. In this study we explicitly manipulated name familiarity of either the Finnish or Moroccan names by means of exposure prior to the experiment.[5] In this study, we again found a null effect on the Finnish-Moroccan IAT. With 104 participants, this effect seems very robust and has very tight confidence intervals. We did not find any effects of name familiarity.

Interpretation of the Finnish-Moroccan IAT-effect is further complicated by the fact that presenting two out-groups in one IAT changes the context such that the out-groups may no longer be viewed as out-groups. More specifically, it could be that Moroccans are only viewed as negative in the direct comparison with Dutch people. Therefore, the absence of an IAT-effect could simply be due to the different context in which the Moroccan category is presented. However, we did find differences in the explicit rating on Finnish and Moroccan people, which were also presented in the same context.[6]

In sum, our study supports the alternative explanation that the IAT-effect is due to in-group/out-group membership: We found that participants responded similarly to the racially charged out-group Moroccan names and the racially neutral out-group Finnish names. When in-group Dutch names were contrasted with either of the two out-group names, there was an IAT-effect. Thus, the present experiment offers no support for the contention that the name-race IAT originates mainly from a prejudice based on race.

Advocates of the IAT might argue that the IAT has been shown to predict prejudice-based behavior. Indeed, the name-race IAT-effect has been shown to correlate with White-Black interracial behavior ($r = .24$) and with other intergroup behavior ($r = .20$; Greenwald, Poehlman, Uhlmann, & Banaji, 2009). Considering the modest size of these correlations, we feel it would be bold to conclude that the IAT provides a reliable indication of racial prejudice.

[5] For details, see section "Study 2" in the online appendix.
[6] See the online appendix.

While the IAT can be used in conjunction with other tools to predict behavior in real-life settings, caution must be used when making claims about the IAT's ability to measure attributes that cause these kinds of behavior.

To conclude, people aren't perfect, and racial discrimination is evident throughout the world. Our research shows, however, that the currently popular name-race IAT may lead researchers to overestimate the degree of people's implicit racial prejudice.

References

Blair, I. V., Judd, C. M., Havranek, E. P., & Steiner, J. F. (2010). Using community data to test the discriminant validity of ethnic/racial group IATs. *Zeitschrift für Psychologie / Journal of Psychology, 218*, 36–43.

Blanton, H., & Jaccard, J. (2008). Unconscious racism: A concept in pursuit of a measure. *Annual Review of Sociology, 34*, 277–297.

Blanton, H., Jaccard, J., Gonzales, P. M., & Christie, C. (2006). Decoding the implicit association test: Implications of conceptual and observed differences scores for criterion prediction. *Journal of Experimental Social Psychology, 42*, 192–212.

Centraal Bureau voor de Statistiek (2008). *Bevolking; herkomstgroepering, generatie, geslacht en leeftijd, 1 januari* [Population; gender, age, sex, and origin, 1 January]. Retrieved from http://statline.cbs.nl/StatWeb/

Dasgupta, N., McGhee, D. E., Greenwald, A. G., & Banaji, M. R. (2000). Automatic preference for white Americans: Eliminating the familiarity explanation. *Journal of Experimental Social Psychology, 36*, 316–328.

Dolfing, M., & van Tubergen, F. (2005). Bensaidi of Veenstra? *Sociologie, 1*, 407–422.

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297–327.

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (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.

Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology, 1*, 17–41.

Klauer, K. C., Voss, A., Schmitz, F., & Teige-Mocigemba, S. (2007). Process components of the implicit association test: A diffusion-model analysis. *Journal of Personality and Social Psychology, 93*, 353–368.

Luce, R. D. (1986). *Response times*. New York, NY: Oxford University Press.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). 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. 265–292). London, UK: Psychology Press.

Ottaway, S. A., Hayden, D. C., & Oakes, M. A. (2001). Implicit attitudes and racism: Effects of word familiarity and frequency on the Implicit Association Test. *Social Cognition, 19*, 97–144.

Popa-Roch, M., & Delmas, F. (2010). Prejudice Implicit Association Test effects. *Zeitschrift für Psychology / Journal of Psychology, 218*, 44–50.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review, 85*, 59–108.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review, 9*, 438–481.

Rothermund, K., & Wentura, D. (2004). Underlying processes in the Implicit Association Test: Dissociating salience from associations. *Journal of Experimental Psychology: General, 133*, 139–165.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t*-tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16*, 225–237.

Rudman, L. A., Greenwald, A. G., Mellott, D. S., & Schwartz, J. L. K. (1999). Measuring the automatic components of prejudice: Flexibility and generality of the Implicit Association Test. *Social Cognition, 17*, 437–465.

van Ravenzwaaij, D., & Oberauer, K. (2009). How to use the diffusion model: Parameter recovery of three methods: EZ, fast-dm, and DMAT. *Journal of Mathematical Psychology, 53*, 463–473.

Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods, 39*, 767–775.

Voss, A., & Voss, J. (2008). A fast numerical algorithm for the estimation of diffusion model parameters. *Journal of Mathematical Psychology, 52*, 1–9.

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of *p*-values. *Psychonomic Bulletin & Review, 14*, 779–804.

Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology, 21*, 641–671.

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist, 35*, 151–175.

Zeelenberg, R., Wagenmakers, E.-J., & Rotteveel, M. (2006). The impact of emotion on perception. *Psychological Science, 17*, 287–291.

Received November 3, 2009
Revision received August 20, 2010
Accepted August 20, 2010
Published online November 24, 2010

Don van Ravenzwaaij

Department of Psychology
University of Amsterdam
Roetersstraat 15
1018 WB Amsterdam
The Netherlands
Tel. +31 20-525-8870
E-mail vanravenzwaaij@uva.nl

