Title
Toward a Process-Level Understanding of Correspondence Among Indirect Measures of Racial Threat Stereotypes

Journal
Social Cognition, 44(4)

ISSN
0278-016X

Authors
Laukenmann, Ruben
Calanchini, Jimmy

Publication Date
2026-08-01

DOI
10.1521/soco.2026.44.4.426

THIS MANUSCRIPT IS ACCEPTED FOR PUBLICATION IN *SOCIAL COGNITION*

Towards a Process-level Understanding of Correspondence among Indirect Measures of Racial Threat Stereotypes

Ruben Laukenmann[1,2,3] & Jimmy Calanchini[2]

[1] University of Mannheim

[2] University of St. Gallen

[3] University of California, Riverside

Author Note

Ruben Laukenmann https://orcid.org/0000-0002-4780-4845

Jimmy Calanchini https://orcid.org/0000-0003-1959-143X

Manuscript preparation was supported by the Research Training Group “Statistical Modeling in Psychology funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – GRK 2277 – Projektnummer 310365261“, by a Doctoral Research Fellowship from the Stiftung der Deutschen Wirtschaft (sdw) gGmbH, funded by the German Federal Ministry of Education and Research, and by a one-year research grant for doctoral candidates funded by the German Academic Exchange Service (DAAD) awarded to Ruben Laukenmann.

The authors are very grateful to C. E. Phills and A. M. Rivers for sharing the stimuli material we used for our studies.

Correspondence concerning this article should be addressed to Ruben Laukenmann, Institute of Behavioral Science and Technology, University of St. Gallen, CH-9000 St. Gallen, Switzerland (email: ruben.laukenmann@unisg.ch) or Jimmy Calanchini, Department of Psychology, University of California Riverside, Riverside, CA 92521 (email: jimmy.calanchini@ucr.edu). Manuscript preparation by Ruben Laukenmann was conducted during institution changes between the University of Mannheim to the University of St. Gallen and during a research visit at the University of California, Riverside.

Online supplementary material is available at https://osf.io/2whze/

This research was approved by the Institutional Review Board of the University of California, Riverside.

We have no conflict of interest to declare.

This manuscript is a Research Article.

Towards a Process-level Understanding of Correspondence among Indirect Measures of Racial Threat Stereotypes

Abstract

Previous research often reveals relatively poor correspondence among three prominent indirect measures configured to assess racial stereotypes: the Weapon Identification Task (WIT), the First-Person Shooter Task (FPST), and the Implicit Association Test (IAT). Importantly, these measures differ on a variety of procedural dimensions (e.g., task instructions, stimulus presentation, response time limits), which confounds explanations of low correspondence. In a fully within-participants design, *N* = 372 participants completed versions of the WIT, FPST, and IAT that were aligned on several procedural dimensions. Process modeling revealed high correlations among measures for model parameters reflecting controlled and automatic processes. However, model-based racial bias estimates correlated only moderately between the WIT and FPST, and not at all with the IAT. Our findings suggest that all three measures share some correspondence when procedurally aligned but each reflects a specific aspect of racial threat stereotypes.

Abstract: 139 words

Keywords: weapon identification task, first person shooter task, implicit association test, multinomial processing tree modeling, process model, racial bias

Word count: 9902 (excluding references [1946], footnotes [706], tables, and figures)

Introduction

Indirect measures[1] were initially developed to assess the influence of racial biases on behavior by minimizing the influence of other mental processes that would otherwise constrain the expression of bias (Gawronski et al., 2020). Measures such as the Weapon Identification Task (WIT; Payne, 2001), First-Person Shooter Task (FPST; Correll et al., 2002), and the Implicit Association Test (IAT; Greenwald et al., 1998) are often configured to assess racial bias operationalized as stereotypes (i.e., traits) associated with different racial groups (e.g., Black men seen as threatening). However, previous research has revealed relatively weak to no correlations between racial bias estimates assessed by different indirect measures of racial bias[2] (Bar-Anan & Nosek, 2014; Cunningham et al., 2001; Glaser & Knowles, 2008; Ito et al., 2015; Payne, 2005; Volpert-Esmond et al., 2020), which is surprising for measures that are configured to assess a common construct.

[1] The measures we examine in the present research are sometimes colloquially referred to as "implicit measures" (Gawronski et al., 2020; Greenwald & Lai, 2020; Payne & Correll, 2020). However, the term "implicit" is used inconsistently in (and beyond) the social cognitive literature to refer to characteristics of the measurement instrument, outcome measure, behavioral response, or underlying mental representation (Corneille & Hütter, 2020; Gawronski et al., 2020). Instead, here we adopt the term "indirect" to refer to a measure that quantifies responses in terms of their latency or accuracy, and in contrast to measures in which participants are directly asked to report a response (e.g., "How much do you associate Black men with threat?").

[2] Bar-Anan and Vianello (2018) found relatively close correspondence across indirect measures of racial bias by accounting for measurement error. However, Schimmack (2001) criticized this study for unrealistic assumptions about method variance in their structural equation modeling approach.

Small correlations among conceptually analogous indirect measures may call into question the validity of their racial bias estimates (Schimmack, 2021). Nevertheless, even when indirect measures are configured to assess the same construct, a wide variety of procedural differences among those measures remain, such as differences in stimuli, response time limits, number of trials, task instructions, and stimuli presentation procedures. Task procedures necessarily determine which processes influence responses (Gawronski et al., 2010), so procedural differences may attenuate correlations across measures that otherwise assess the same construct. Thus, in the present research, we assess correspondence across three indirect measures configured to assess threat stereotypes – the WIT, FPST, and IAT– that are equated across several procedural dimensions. In doing so, we remove procedural differences that otherwise confound comparisons across measures (i.e., stimuli material, response time limit, number of trials) while retaining features that are inherent to each indirect measure (i.e., task structure).

Moreover, responses on indirect measures of racial bias are traditionally operationalized in terms of summary statistics (e.g., Bar-Anan & Nosek, 2014; Bar-Anan & Vianello, 2018; Cunningham et al., 2001; Glaser & Knowles, 2008; Greenwald et al., 2003), which provide little insight into the cognitive process(es) that contribute to responses (Calanchini, 2020). For example, similar performance across indirect measures could be due to participants' cognitive control abilities in responding, their biased mental representations, or both (Ito et al., 2015; Klauer et al., 2010; Payne, 2005). Consequently, in the present research we apply the Process Dissociation Procedure (Payne, 2001) as a formal mathematical modeling technique to disentangle the influence of controlled and automatic processes on performance and investigate correspondence among those processes across indirect measures.

Process modeling and indirect measures of racial threat stereotypes

We are not the first to recognize that procedural differences across indirect measures affect correspondence across racial bias estimates (e.g., Olson & Fazio, 2003), nor are we the first to use formal mathematical modeling to assess correspondence across measures. Indeed, Ito and colleagues (2015) investigated the contributions of two qualitatively distinct types of cognitive processes on the WIT, FPST, and IAT. They found that estimates of controlled processing correlated moderately-to-strongly across measures (r = .29 – .61) but estimates of racial bias based on automatic processing correlated weakly across measures (r = .02 – .16).

To the extent that indirect measures of racial bias were developed to primarily assess automatically activated mental representations (Fazio et al., 1995; Gawronski et al., 2020; Greenwald et al., 1998), the small magnitude of correlations among racial bias estimates calls into question whether these measures configured to assess threat stereotypes reflect a common construct. However, and importantly, Ito and colleagues (2015) relied on the standard versions of each indirect measure as they are typically configured. This approach certainly has ecological validity, in terms of correspondence with how these measures are traditionally used by researchers. But as we review below, these three indirect measures differ from one another on many dimensions – which in turn confounds any strong interpretations of the small correlations in racial bias estimates between measures.

Overview of standard versions of indirect measures of racial bias

Weapon Identification Task

The WIT is a sequential priming paradigm developed by Payne (2001). Over a series of trials, participants view a prime image (i.e., Black or White male face) quickly followed by a target image (i.e., gun or tool). On each trial, participants' task is to identify quickly and accurately the target while disregarding the prime (Payne, 2001; Rivers, 2017).

First-Person Shooter Task

The FPST is a simplified videogame simulation developed by Correll et al. (2002). Over a series of trials, participants view a naturalistic scene (e.g., a park) in which a person (i.e., Black or White man) appears who is either armed (e.g., holding a gun) or unarmed (e.g., holding a cell phone). On each trial, participants' task is to quickly and accurately decide whether to "shoot" or "don't shoot". Target persons are presented as full body photographs with varying postures against different background scenes, with changing target onset times and screen positions (Correll et al., 2002; 2007; 2015).

Implicit Association Test

The IAT is a dual-categorization task developed by Greenwald et al. (1998). Over a series of trials, participants view two stimulus types – typically target groups (e.g., pictures of Black and White male faces) and attributes (e.g., words referring to threat or safety) – presented one at a time. In some blocks of trials, one target type shares a response key with one attribute type (e.g., Black/threat) and the other target type shares a response key with the other attribute type (e.g., White/safety). However, in other blocks of trials, the key pairings are reversed (e.g., Black/safety, White/threat). On each trial, participants' task is to categorize the presented stimulus quickly and accurately.

Similarities and differences among standard versions of indirect measures of racial bias

All three of these measures – the WIT, FPST, and the IAT – have been used to assess racial bias operationalized in terms of threat stereotypes. Moreover, all three measures share a common feature in that they require participants to make a categorization decision (e.g., whether the stimulus is a gun or tool). However, as Table 1 illustrates, these three measures diverge on several procedural characteristics, including: the number of critical trials, the presence and length of the response time limit, and stimulus material. Not only do these measures differ from one another on these procedural dimensions, but versions of each measure have also been implemented with a variety of task procedures. For example, the WIT has been used with additional neutral face outlines and has been configured with as few as 80 and as many as 1100 critical trials (for a review, see Rivers, 2017). The FPST exists in several different adaptations that present full body images of people holding the target object in their hand (Correll et al., 2002; Payne & Correll, 2020), target objects superimposed on the forehead of a face (Plant et al., 2005), and target objects presented on the left or right side of a face (Unkelbach et al., 2008, 2009). Furthermore, IATs have been developed that implement response deadlines, different numbers of trials, and different block structures (Calanchini et al., 2021; Meissner & Rothermund, 2013; Sriram & Greenwald, 2009).

Table 1

Procedural details of the WIT, FPST, and IAT in their originally published versions.

	Weapon Identification Task (WIT; Payne, 2001)	First-person Shooter Task (FPST; Correll et al. 2002)	Implicit Association Test (IAT; Greenwald et al., 1998; Study 3)
response time limit	none (Study 1); 500 ms (Study 2)	850ms (Study 1 & 3); 630ms (Study 2)	none
number of practice trials	48 trials	no trials	150 trials
number of experimental trials	192 trials	80 trials	200 trials
stimulus material: target group / category	cropped face images of Black and White male faces as primes	full body image of Black and White men	first names judged to be more likely to belong to Black or White persons
stimulus material: target object / attribute	target images of weapons and tools	target weapon object or an innocuous object (e.g., a cell phone, a wallet) held by Black and White men	unpleasant and pleasant words as attributes
presentation order of stimulus material	sequential: prime face followed by target objects presented in the center of the screen	concurrent: person holding target object presented at a random position on screen	serial: target category or attribute stimuli presented in the center of the screen in a random order between target and attribute stimuli
instruction	identifying target object while ignoring face prime	“shoot” armed person, “don’t shoot” unarmed person	correctly categorize the target category and attribute stimuli
analysis	correct response latencies, error rates, process dissociation procedure	correct response latencies, error rates, signal detection modeling	correct response latencies

Procedural differences necessarily introduce variance between measures, which in turn may obscure relationships among estimates of racial bias (Fazio et al., 1995; Greenwald & Lai, 2020; Mekawi

& Bresin, 2015; Payne, 2005; Olson & Fazio, 2003; Volpert-Esmond et al., 2020). Differences in the stimuli presentation traditionally used in each indirect measure provide a straightforward illustration of this point. For example, the WIT typically presents the target object clearly visible in the center of the screen without any distracting cues, and at approximately the same size as the face primes are presented. In contrast, the FPST typically presents the target object held in the hand of a person who is pictured in a variety of scenes and thus represents only a small feature in a larger context. The IAT differs from both of these paradigms and typically does not use images of target objects but, instead, represents attributes like words indicating safety versus danger (Ito et al., 2015). Stimuli that are intended to reflect different operationalizations of the same construct (e.g., words versus pictures) can nevertheless activate different sets (or subsets) of mental representations (Foroni & Bel-Bahar, 2010; Rosch, 1975). Thus, low correspondence among these measures may reflect different kinds of stimuli activating different subsets of threat stereotypes.

Though all three indirect measures differ on a variety of procedural dimensions, we propose that task structure is the most crucial difference. The WIT presents stimuli sequentially and explicitly instructs participants to respond to target objects and disregard face primes. The FPST presents stimuli concurrently, and participants are instructed to "shoot" an armed person and "don't shoot" an unarmed person.[3] The IAT presents stimuli serially, such that participants must attend and categorize both types of stimuli. In the present research, we argue that task structure is the defining feature of each measure. Thus, we aligned measures on procedural characteristics that have varied in previous research but held constant what we view as the defining feature of each measure: task structure.

---

[3] A recent version of the FPST also allows for a dynamic presentation of target stimuli by showing a video clip of a person reaching into their pocket and pulling out the target object, resulting in a sequential presentation of the person first and the target object second (Frenken et al., 2022).

Indirect measures of racial bias reflect the contributions of multiple cognitive processes

Despite procedural differences among the WIT, FPST, and IAT, they were all designed with the same goal of assessing racial stereotypes by minimizing the contributions of other processes that would constrain their expression. To be sure, compared to analogous direct measures (e.g., feeling thermometers; questionnaires), these indirect measures minimize the contributions of motivations and biases that would modify responses in a socially desirable direction. However, responses on indirect measures do not solely reflect the contributions of stereotype information. Instead, multinomial processing tree (MPT; Riefer & Batchelder, 1988) models can disentangle the joint contributions of multiple cognitive processes to responses on indirect measures. MPT models are tailored to specific experimental paradigms that provide frequency data (e.g., number of correct and incorrect responses), and specify the number, nature, and composition of cognitive processes thought to be involved in the paradigm (for reviews see Calanchini, 2020; Erdfelder et al., 2009; Hütter & Klauer, 2016).

In the present research, we applied the Process Dissociation Procedure (PDP; Payne, 2001) to all three indirect measures of racial bias. The PDP is an MPT that has been applied to each of these measures in previous research and provides a robust way to disentangle the contribution of controlled and automatic processes on response behavior (Bishara & Payne, 2009; Payne, 2009; Ito et al., 2015).[4]

[4] For full transparency, one of the original goals of the present research was to compare different MPT models across the three indirect measures. A variety of other models, such as Lindsay and Jacoby's (1994) Stroop model and Conrey et al.'s (2005) quad model, have been previously compared in the context of indirect measures like the WIT (Bishara & Payne, 2009). We applied different MPT models to the data from each indirect measure, and the PDP consistently provided best fit to all three measures. Consequently, we focus on results from the PDP in the main text and report the results from the other MPT models in the supplement.

Below we describe the structure and the logic of the PDP in the context of the WIT, which also apply to the FPST and IAT.

Process Dissociation Procedure

The PDP (Jacoby, 1991; Payne, 2001) is depicted in Figure 1. This model consists of a Controlled (*C*) and an Automatic (*A*) process parameter.[5] The Controlled process parameter represents any process(es) that result(s) in a correct response, which include but are not limited to general accuracy in responding based on successful target discrimination, and conflict monitoring/resolution when automatic processes would produce a response that conflicts with the correct response (Klauer & Voss, 2008; Laukenmann et al., 2023). The Automatic process parameter represents response tendencies towards a "gun" versus "tool" response, which may include the influence of threat stereotypes by prime faces, handedness, guessing, and recoding (Klauer & Voss, 2008; Meissner & Rothermund, 2013). Vitally, in the PDP, Controlled processing dominates automaticity, such that the influence of Automatic processing is irrelevant whenever the Controlled process succeeds. However, and importantly, the PDP is agnostic about the temporal order of processing. For example, Automatic processing might be activated very quickly after stimulus onset, but if Controlled processing successfully intervenes at any time before a response is made then it will override the response tendency that would be produced by Automatic processing (Klauer & Voss, 2008; Laukenmann et al., 2023). Thus, the equations of the PDP are specified

[5] We use the terms Controlled and Automatic to refer to the PDP's parameters in correspondence with the terminology of Payne (2001). However, we expressly refrain from assuming that the processes assessed by either parameter possess features of controllability (e.g., conscious, intentional, resource-dependent, slow) or automaticity (e.g., unconscious, unintentional, efficient, fast) as articulated in traditional dual-process perspectives of cognition (e.g., Metcalfe & Mischel, 1999; Shiffrin & Schneider, 1977).

such that the success of the Controlled process (*C*) will always produce the correct response (+). The Automatic process can only produce a response in the absence of influence from the Controlled process (1 – *C*), with the probability (*A*) representing a "gun" response and with the counter-probability (1 – *A*) representing a "tool" response. When the target stimulus is a gun, (*A*) produces a correct response (+), and (1 – *A*) produces an incorrect response (–). In contrast, when the target stimulus is a tool, (*A*) produces an incorrect response (–), and (1 – *A*) produces a correct response (+).

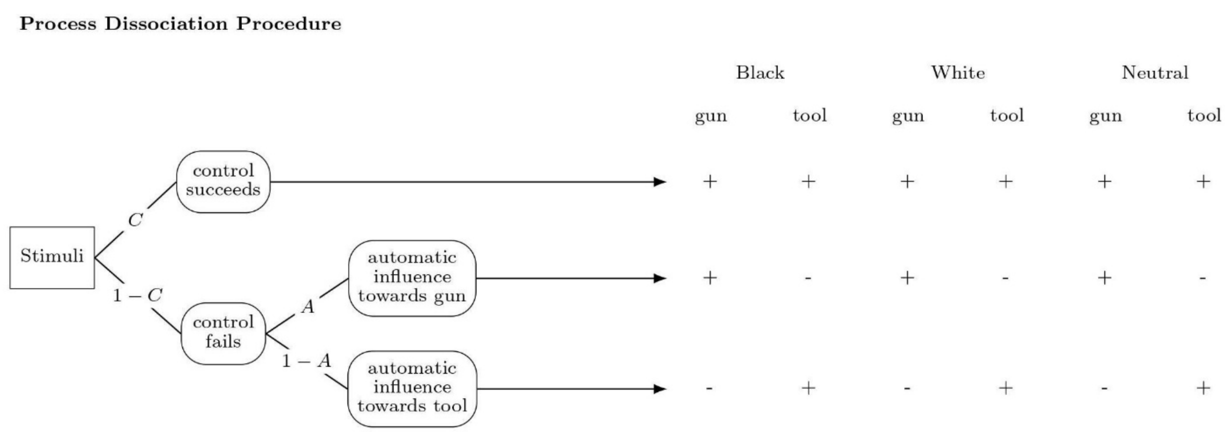

Figure 1. Multinomial processing trees of the process dissociation procedure (PDP). Branches lead to correct (+) and incorrect (–) responses. *C* = Controlled process parameter, *A* = Automatic process parameter.

The present research

The aim of the present research is to investigate correspondence across three indirect measures of racial bias – the WIT, FPST, and IAT – configured to assess racial threat stereotypes. Importantly, we retain each measure's traditional structure (i.e., sequential, concurrent, serial, respectively) but align them on all other procedural dimensions. Specifically, we used the same stimuli, response time window, and number of trials in all three measures. In doing so, we can investigate the extent to which the small correlations among racial bias estimates identified in previous research (Glaser & Knowles, 2008; Ito et al., 2015; Payne, 2005) reflect procedural artifacts versus meaningful differences among measures.

In the present research, we relied on a fully within-participants design such that each participant completed all three measures. We investigated correspondence among responses in two main ways. First, we examined summary statistics of performance in terms of differences in response accuracy and latency between critical trials. Next, to provide more theoretically precise insight than can be revealed by summary statistics, we applied the PDP to participants' responses on each indirect measure. Replicating previous research, we expect parameters that reflect controlled processes to correlate moderately-to-strongly between indirect measures (Ito et al., 2015). Additionally, and in contrast to previous research, we expect parameters that reflect automatic processes and racial bias to correlate moderately-to-strongly between measures because we have aligned stimuli across them.

All data, preregistration, analytic code and supplementary material are available at https://osf.io/2whze/.

Methods

Participants

We recruited 547 undergraduate students at a large, public Southern California university to participate for partial course credit. As required by the Institutional Review Board at the university where the study took place, we excluded participants from analysis who chose to reject data inclusion at the end of the study (44 rejected inclusion). We excluded participants whose error rate was >50% for any task (9 WIT; 2 FPST; 5 IAT; 1 in at least two tasks) to account for random responding. We also excluded participants if they missed trials (i.e., trials with no response, a too fast response <100ms, or a too slow response >1700ms) at a rate 1.5 times the interquartile range above the median (i.e., Tukey's criterion) for any task (43 WIT; 26 FPST; 6 IAT; 46 in at least two tasks) to account for inattentive responding. Because we collected these data online, we adopted relatively conservative inclusion criteria to ensure data quality. In total, we excluded 26.0% of participants who agreed to be included in data analysis.

The final sample comprised 372 participants (age: M_{age} = 19.6, SD_{age} = 2.3; gender: 225 female, 141 male, 5 other, 1 unreported; race: 33 White, 11 Black, 130 Asian, 146 Latino, 48 other, 4 unreported). A post-hoc power analysis for a repeated measurement ANOVA with a sample size of N = 372, and a Type-1 error level of α = .05 (Faul et al., 2007) afforded a test power of at least $1-\beta > .49$ to detect small effects (f = .1 or η_p^2 = .01 in the underlying population, cf. Cohen, 1992). If effects are at least of medium size (f = .25 or η_p^2 = .06 in the underlying population, cf. Cohen, 1992), the power increases to $1-\beta > .99$ under otherwise identical conditions. A post-hoc power analysis for a bivariate correlation and a two-tailed Type-1 error level of α = .05 afforded a test power of at least $1-\beta = .49$ to detect a small correlation (r = .10) and increased to a test power of at least $1-\beta > .97$ to detect a correlation of r =. 20 given identical conditions. Overall, our study is sufficiently powered.

Procedure

We collected data online using the software lab.js (Henninger et al., 2022). After providing consent, participants completed the three indirect measures in random order. Then, participants completed a basic demographic questionnaire (age, gender, ethnicity) and were asked what they thought the study was about. Participants were thanked and debriefed at the end. The response keys D and L were counterbalanced between participants across all measures, and the response keys corresponding to gun and tool targets were held constant within participants.

Weapon Identification Task. Participants completed an adapted version of the WIT (Payne, 2001; Rivers, 2017). Participants were instructed to identify as quickly and accurately as possible a target object (i.e., gun, tool) preceded by a face image as prime. In addition to the Black and White male faces, the adapted WIT also included an outline of a face as a neutral prime (Rivers, 2017).[6] In each trial, participants were presented with a fixation cross (500 ms), a face prime (200 ms), a target object (200 ms), a pattern mask (500 ms),[7] and a feedback screen (1000 ms), each presented in the center of the screen. On practice trials, participants received the following feedback: 'correct!', 'false!', 'too slow!'. Slow responses were operationalized as responses made 700 ms or more after target object onset. On experimental trials, participants only received feedback if their response was too slow. Participants' response latency was recorded, even if it exceeded the 700 ms limit. Participants first completed 20 practice trials containing only neutral face outlines as primes, half of which were paired with a gun target and half of which were paired with a tool target. Next, participants completed 240 experimental trials with 80 trials for each prime race by target object combination in random order. Participants had two self-paced breaks after 80 and 160 trials.

---

[6] We included neutral face primes to provide enough degrees of freedom for the MPT model comparison that we initially planned for this project (see Footnote 4). We still use responses to neutral faces as reference category to investigate differences in response tendency for different face races compared to a neutral face outline.

[7] In a first wave of preliminary data collection (*N* = 42), we imposed a response time limit of 500 ms after target onset (with a pattern mask shown for 300 ms) which resulted in a substantially higher error rate for the WIT (*M* = 33.9%) compared to the FPST (*M* = 13.8%) and the IAT (*M* = 12.6%). To improve comparability between indirect measures, we aim to have similar error rates for participants across measures and thus adjusted the response time limit to 700 ms for the WIT. This preliminary data was not included in further analysis.

First-Person Shooter Task. Participants completed an adapted version of the FPST (Correll et al., 2002, 2014; Unkelbach et al., 2008). Participants were instructed to “shoot” or “don’t shoot” as quickly and accurately as possible an image of a Black male face, a White male face, or the outline of a face as a neutral image, that was paired with either a gun (“armed”) or tool (“unarmed”) target object. The target object was displayed on either the left or right side of the face with its handle pointing towards the face. In each trial, participants were presented with a fixation cross (500 ms), a face image and target object presented simultaneously at one of nine random positions on the screen (700 ms), and a feedback screen (1000 ms). On practice trials, participants received the following feedback: ‘correct!’, ‘false!’, ‘too slow!’. Slow responses were operationalized as response made 700 ms or more after the onset of the target object with the face image. On experimental trials, participants only received feedback if their response was too slow. Participants first completed 20 practice trials containing only neutral face outlines as face images, half of which were paired with a gun target and half of which were paired with a tool target. Next, participants completed 240 experimental trials with 80 trials for each face race by target object combination in random order. Participants had two self-paced breaks after 80 and 160 trials.

Implicit Association Test. Participants completed an adapted version of the IAT (Greenwald et al., 1998). Participants were instructed to categorize as quickly and accurately as possible target objects (a gun or a tool) and face images (a Black male face, a White male face, or a neutral face outline). In each trial, participants were presented with a fixation cross (500 ms), a target object or face image (700 ms), and a feedback screen (1000 ms), each presented in the center of the screen. On practice trials, participants received the following feedback: ‘correct!’, ‘false!’, ‘too slow!’. Slow responses were operationalized as responses made 700 ms or more after the onset of the target object or face image. On experimental trials, participants only received feedback if their response was too slow. In each trial, key assignments for the categories (target objects: ”gun”/“tool”; race categories: “Black”/“White”/“neutral”) were continuously displayed on the left or right lower corner on screen. In total, the IAT consisted of 13 blocks, seven practice blocks with 20 trials and six experimental blocks with 40 trials each. The first practice block consisted of learning the assignment of the target objects *gun* and *tool* to the left or right response key. Next, each of the three combinations of face image category pairings (Black versus White male faces; Black male faces versus neutral outline; White male faces versus neutral outline) were presented in a grouping of four blocks. These four blocks each consisted of a practice block in which participants learned the assignment of the race categories to the left and right response keys, an experimental block combining target objects and face images, a practice block with reversed assignment of the race categories to the response keys, and an experimental block combining target objects and face images with reversed key assignment for the face images. The order of key pairings for target objects with race categories was randomized within each race category pairing. The four blocks of each race category pairing were presented in a random order for each participant.

Materials

Each indirect measure included the same stimulus material taken from Rivers (2017) and Phills et al. (2011) and were displayed with a 300 × 300-pixel resolution. Face primes consisted of 24 Black male faces, 24 White male faces, and one neutral image of the outline of a face. Target objects consisted of 5 drawings of guns and of 5 tools presented horizontally.

Results

Data pre-processing

Prior to all analyses, we excluded trials with latencies <100 ms and >1700 ms, resulting in exclusion of 1.05% of trials. The indirect measures were presented in random order for each participant to rule out order effects. We first analyzed error rates, correct response times, and PDP parameters from the three indirect measures. In a second exploratory set of analyses, we investigated whether the order in which participants completed the measures affected PDP parameters.

Error rates analysis[8]

Figure 2 shows the error rates of all three indirect measures by race and target. A 3 (task) × 3 (race) × 2 (target) mixed analysis of variance (ANOVA) with Greenhouse-Geyser correction on the error rates yielded a main effect of measure $F(1.68, 622.45) = 60.75$, $p < .001$, $\eta_p^2 = .141$, which indicates that error rates varied across measures. To investigate these differences, we used paired *t*-test comparisons. The WIT had a significantly higher error rate ($M = 14.0\%$, $SD = 9.4$) than the FPST ($M = 11.1\%$, $SD = 6.9$), $t(371) = 6.37$, $p < .001$, $d_Z = 0.25$, and the IAT ($M = 9.7\%$, $SD = 5.2$), $t(371) = 9.92$, $p < .001$, $d_Z = 0.38$. The FPST had a significantly higher error rate than the IAT, $t(371) = 5.01$, $p < .001$, $d_Z = 0.17$.

The ANOVA also yielded the expected race × target interaction, $F(1.96, 726.08) = 63.08$, $p < .001$, $\eta_p^2 = .145$, indicating that error rates for each target type varied as a function of race. A three-way interaction also emerged, $F(3.71, 1374.92) = 3.53$, $p = .009$, $\eta_p^2 = .009$, indicating that the racial bias effect varied across measures. Nevertheless, analyzing each measure separately, the race by target interaction remained significant for all three measures: WIT $F(1.86, 690.93) = 26.53$, $p < .001$, $\eta_p^2 = .067$, FPST $F(1.98, 735.68) = 17.56$, $p < .001$, $\eta_p^2 = .045$, and IAT $F(1.98, 735.79) = 38.13$, $p < .001$, $\eta_p^2 = .093$.

[8] We report means, standard deviations and further statistical analysis of error rates and correct response times in the supplement.

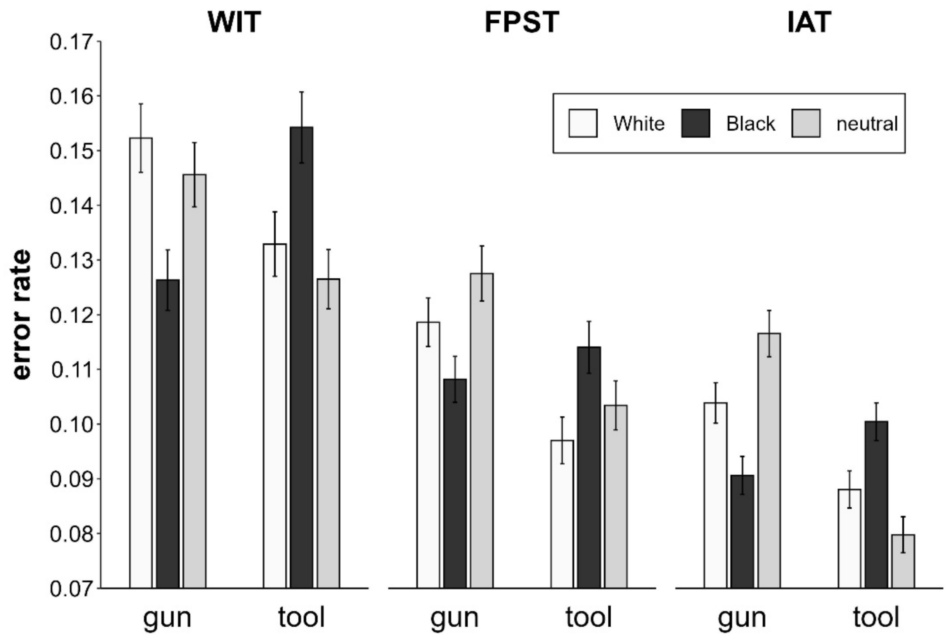

Figure 2. Error rates of the three indirect measures by race and target. WIT = Weapon Identification Task, FPST = First-Person Shooter Task, IAT = Implicit Association Task. Error bars represent one standard error.

To investigate correspondence between racial bias estimates, we analyzed the correlations between summary statistics of accuracy as proxies for racial bias. To do so, we operationalized racial bias as the difference between errors for gun and tool targets for each face type. Specifically, we calculated Black versus White accuracy bias as: ((errors(tool|Black) – errors (gun|Black)) – ((errors (tool|White) – errors (gun|White)). Conceptually replicating previous research, accuracy-based racial bias estimates correlated significantly between the WIT and FPST (r = .18, p < .001), but not between the WIT and IAT (r = .06, p = .27) or between the FPST and IAT (r = .02, p = .77).

Correct response time analysis

Figure 3 shows the correct response times (RT) of all three indirect measures by race and target. A 3 (task) × 3 (race) × 2 (target) mixed analysis of variance (ANOVA) with Greenhouse-Geyser correction on the correct RTs yielded a main effect of measure $F(1.59, 591.06) = 426.02$, p <.001, $\eta_p^2 = .535$, which indicates that correct RTs varied across measures. To investigate these differences, we used paired *t*-test comparisons. The FPST had significantly slower mean correct RTs (M = 567.8ms, SD = 182.7) than the WIT (M = 505.9ms, SD = 168.8), $t(371) = -22.84$, $p < .001$, $d_Z = -1.01$, and the IAT (M = 506.4ms, SD = 105.4), $t(371) = -61.41$, $p < .001$, $d_Z = -1.13$. However, mean correct RTs for the WIT and IAT were similar, $t(371) = 0.26$, $p < .001$, $d_Z = 0.01$.

The ANOVA also yielded the expected race × target interaction, $F(1.99, 736.93) = 34.07$, $p < .001$, $\eta_p^2 = .084$, indicating that correct RTs for each target type varied as a function of race. A three-way interaction also emerged, $F(2.90, 1445.95) = 25.33$, $p < .001$, $\eta_p^2 = .064$, indicating that the racial bias effect varied across measures. Analyzing each measure separately showed that the race by target interaction is only significant for the IAT: WIT $F(1.99, 738.36) = 0.57$, $p = .565$, $\eta_p^2 = .002$, FPST $F(2.00, 740.76) = 1.25$, $p = .286$, $\eta_p^2 = .003$, and IAT $F(1.99, 737.35) = 101.70$, $p < .001$, $\eta_p^2 = .215$.

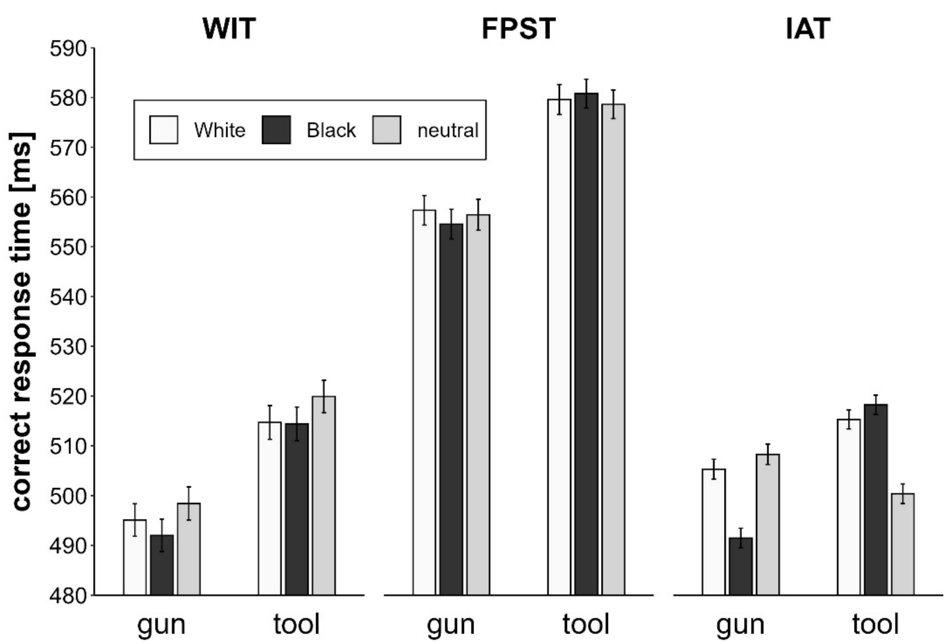

Figure 3. Correct response times of the three indirect measures by race and target. WIT = Weapon Identification Task, FPST = First-Person Shooter Task, IAT = Implicit Association Task. Error bars represent one standard error.

To investigate correspondence between racial bias estimates, we analyzed the correlations between summary statistics of correct RTs as proxies for racial bias. Similarly to above, we operationalized racial bias as the difference between correct RTs for gun and tool targets for each face type, i.e.,: ((RT(tool|Black) – RT (gun|Black)) – ((RT (tool|White) – RT (gun|White)). Correct RT-based racial bias estimates correlated significantly between the WIT and FPST (r = .15, p = .004), but not between the WIT and IAT (r = .02, p = .63) or between the FPST and IAT (r = .01, p = .80).

Multinomial processing tree estimation

Modeling procedure. We used the hierarchical, latent-trait MPT modeling approach of Klauer (2010) as a framework for all analyses. To assess goodness-of-fit, we used Bayesian posterior predictive p-values corresponding to the test statistics T_1 and T_2 (Klauer, 2010; Klauer et al., 2015). A p >.05 reflects no credible difference between observed and predicted frequencies and can be interpreted as evidence that model assumptions are in line with the observed data (Klauer, 2010; Klauer et al., 2015).

We used the R package TreeBUGS (Heck et al., 2018) to fit the hierarchical latent-trait MPT with default priors. We used the Markov Chain Monte Carlo algorithm for three independent estimation chains with 1,000,000 iterations each, of which 250,000 were removed as a burn-in period. Every 500th iteration was retained to compute summary statistics.

Joint modeling across indirect measures. We conducted a joint MPT-model estimation which applied the PDP to each indirect measure (PDP-joint model). This approach allows us to estimate PDP parameters and their correlations within and between measures in one modeling step. We specified one Controlled parameter for each measure and three Automatic parameters for each race condition within each measure.

The Rubin-Gelman statistic was smaller than $\hat{R}$ < 1.05 for all parameter estimates, indicating acceptable convergence of estimation chains. The PDP-joint model fit the data well, $p(T_1)$ = .08, $p(T_2)$ = .45. We report Controlled and Automatic process parameter estimates, racial bias estimates quantified

Table 2

Mean parameter estimates and intercorrelations for PDP-joint model.

Parameter	*Mean*	*95%-BCI*	*1*	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	*11*	*12*	*13*	*14*
WIT																
1. C	.75	[.73 – .77]														
2. A_B	.54	[.53 – .56]	-.13													
3. A_W	.46	[.44 – .48]	-.06	.15												
4. A_N	.46	[.44 – .48]	-.04	**.43**	**.51**											
5. ΔA_{BW}	.08	[.05 – .12]	-.03	**.61**	**-.69**	-.16										
FPST																
6. C	.80	[.78 – .81]	**.56**	-.18	-.12	-.05	-.03									
7. A_B	.51	[.48 – .53]	-.13	**.65**	**.31**	**.56**	**.22**	-.09								
8. A_W	.43	[.41 – .46]	-.10	**.54**	**.62**	**.71**	-.10	**-.21**	**.64**							
9. A_N	.44	[.41 – .46]	**-.18**	**.46**	**.51**	**.70**	-.07	-.08	**.65**	**.76**						
10. ΔA_{BW}	.07	[.05 – .10]	-.03	.10	**-.37**	-.19	**.38**	.14	**.39**	**-.44**	-.15					
IAT																
11. C	.82	[.81 – .83]	**.52**	-.08	.08	.02	-.12	**.65**	-.10	-.08	.02	-.02				
12. A_B	.53	[.51 – .55]	.18	**.38**	**.40**	**.46**	-.04	.24	**.47**	**.43**	**.39**	.03	.20			
13. A_W	.46	[.43 – .48]	-.13	**.34**	**.56**	**.47**	-.20	-.16	**.44**	**.50**	**.43**	-.09	-.02	.36		
14. A_N	.41	[.38 – .43]	-.10	**.46**	**.51**	**.56**	-.07	**-.23**	**.56**	**.64**	**.63**	-.10	.05	.32	**.55**	
15. ΔA_{BW}	.08	[.05 – .11]	**.26**	-.02	-.22	-.08	.16	**.34**	-.05	-.14	-.10	.11	.18	.44	**-.65**	-.26

Note. Correlations are calculated on the probit scale. Correlations in bold indicate that the 95% BCI of the correlation coefficient does not contain zero. BCI = Bayesian Credibility Interval. C = Controlled process parameter, A = Automatic process parameter, B = Black, W = White, N = neutral.

Racial bias estimate: $\Delta A_{BW} = A_B – A_W$.

as the difference between Automatic process parameters for Black versus White targets, and their correlations in Table 2.

Bayesian hierarchical latent-trait MPT models provide posterior distributions of their parameter estimates. We can use these posterior estimates to calculate differences between parameters and correlations which have not already been estimated by the MPT modeling procedure. In addition, we can estimate Bayes factors for parameter differences and their correlations using the Savage-Dickey density ratio test (Wagenmakers et al., 2010), which compares the posterior and prior density at null value of the null hypothesis (i.e., 0 to indicate no differences and zero correlations). Bayes factors quantify the likelihood of one hypothesis versus another, so in this context BF_{10} quantifies the likelihood of the alternative hypothesis compared to the null hypothesis. Following Jeffreys' (1961) Bayes factor interpretation scale, BF_{10} = 3 – 10 reflects moderate, BF_{10} = 10 – 30 reflects strong, BF_{10} = 30 – 100 reflects very strong, and $BF_{10} > 100$ reflects extreme evidence for the alternative hypothesis. Conversely, BF_{10} = 0.10 – 0.33 reflects moderate, BF_{10} = 0.03 – 0.10 reflects strong, BF_{10} = 0.01 – 0.03 reflects very strong, and $BF_{10} < 0.01$ reflects extreme evidence for the null hypothesis. In the following, we provide Bayes factors when we investigate parameter differences and their correlations.

Racial bias estimates across indirect measures. The racial bias estimate for Black versus White male faces in the PDP is reflected in the difference between Automatic process parameters for Black versus White faces: $\Delta A_{BW} = A_B - A_W$. We operationalize racial bias in this way analogously to how racial bias has been quantified before (e.g., Payne, 2005; Ito et al., 2015). We report mean values, their 95% Bayesian Confidence Intervals (BCIs), and Bayes factors comparing the alternative hypothesis of a credible difference ($\Delta A_{BW} \neq 0$) to the null hypothesis of no difference ($\Delta A_{BW} = 0$) using a weakly informative prior for the difference between normally distributed variables $N(0,2)$.[9]

Positive ΔA_{BW} values (Table 2) with 95% BCIs that do not contain zero emerged for all three indirect measures, with Bayes factors suggesting extreme evidence for the alternative hypothesis of a credible difference: WIT (ΔA_{BW} = .08 [.05 – .11], BF_{10} > 100), FPST (ΔA_{BW} = .07 [.05 – .10], BF_{10} > 100), and IAT (ΔA_{BW} = .08 [.05 – .11], BF_{10} > 100). Thus, all three indirect measures indicate that threatening objects (i.e., guns) are credibly associated more strongly with Black than White male faces.

To get a better understanding of the nature of racial bias in these indirect measures, we used Automatic process parameters estimated from the face outline stimuli (A_N) as a neutral reference against which to interpret participants' response tendency towards guns versus tools. For the WIT, the response tendency to respond "gun" was credibly higher for Black than neutral faces ($\Delta A_{BN} = A_B - A_N$ = .08 [.05 – .11], BF_{10} > 100), but not different from neutral for White faces ($\Delta A_{WN} = A_W - A_N$ = -.001 [-.03 – .03], BF_{10} = 0.02). The same pattern emerged for the FPST, such that the response tendency to respond "gun" was credibly higher for Black than neutral faces (ΔA_{BN} = .07 [.04 – .10], BF_{10} > 100), but not different from neutral for White faces (ΔA_{WN} = -.004 [-.03 – .02], BF_{10} = 0.02). The IAT demonstrated a different pattern

[9] We conducted the Savage-Dickey ratio test for MPT parameters on the probit scale to ensure normal distribution of parameters. The default prior distribution of parameters in latent-trait MPT models in TreeBUGS is a multivariate normal distribution (inverse-Wishart) reflecting a weakly informative prior which also allows for correlations between parameters. However, to approximate the role of this weakly informative prior distribution we choose N(0,2) as a weakly informative prior for parameter differences based on the assumption that the difference between two weakly informative normally distributed parameter priors with N(0,1) results in N(0,2) for their difference: N(0,1) – N(0,1).

of results, such that the response tendency to respond “gun” was credibly higher for both Black (ΔA_{BN} = .13 [.10 – .16], BF_{10} > 100) and White faces (ΔA_{WN} = .05 [.02 – .08], BF_{10} = 4.42) than for neutral. Taken together, in all three measures Black faces are more strongly associated with threat than is the neutral face outline. However, only in the IAT White faces are more strongly associated with threat than is the neutral face outline.

To further investigate racial bias in the IAT against the baseline of the neutral face outline, we exploratorily estimated a PDP model specified with Automatic process parameters for each IAT block combination of face-image category pairings (i.e., Black versus White male faces; Black male faces versus neutral outline; White male faces versus neutral outline) and target type (object versus face).[10] This model specification positions us to compare two different operationalizations of racial bias. The first operationalization reflects how IAT bias is traditionally conceptualized and is based on the difference between Automatic process parameters for the IAT block directly contrasting Black and White faces. We found a credible racial bias estimate larger than zero for both target types (objects: ΔA_{BW} = .14 [.07 – .20], BF_{10} > 100; faces: ΔA_{BW} = .12 [.06 – .19], BF_{10} = 18.92). The second operationalization conceptualizes bias against a neutral standard and is based on the difference between Automatic process parameters for the two IAT blocks directly contrasting the Black or White face with the neutral face outline (i.e., the difference between the Automatic process parameters of the Black faces from the Black versus neutral IAT block and of the White faces from the White versus neutral IAT block). We found a racial bias estimate that was not credibly different from zero for either target type (objects: ΔA_{BW} = .01 [-.04 – .06], BF_{10} = 0.03; faces: ΔA_{BW} = .08 [-.01 – .16], BF_{10} = 0.23). This exploratory analysis indicates that racial bias

[10] We report the full model estimation with model fit, parameter estimates, racial bias estimates, and correlations in the supplement (Table S9).

in the IAT is driven by directly contrasting Black and White race categories (i.e., Black versus White IAT blocks). However, when contrasting either Black or White race categories against the neutral face outline, Black and White faces were similarly strongly associated with threat. Taken together, racial bias only emerges in IAT blocks in which Black and White race categories are directly contrasted, and not in blocks in which either race category is contrasted against the neutral face outline.

Parameter and racial bias correlations across tasks. In the PDP-joint model, Controlled process parameters correlated strongly (r = .52 – .65) among indirect measures, indicating that processes that contribute to overall accuracy across measures correspond highly within participants. Automatic process parameters correlated moderately to strongly within measures (WIT: r = .15 – .51; FPST: r = .64 – .76; IAT: r = .32 – .55) and moderately to strongly between WIT and FPST (r = .31 – .71), WIT and IAT (r = .34 – .56), and FPST and IAT (r = .39 – .64). This pattern of results indicates that participants' response tendency towards "gun" over "tool" corresponded highly within and across indirect measures.

Racial bias estimates (i.e., $\Delta A_{BW} = A_B - A_W$) correlated moderately between the WIT and FPST (r = .38 [.09. – .66]), but not between the IAT and WIT (r = .16 [-.19. – .49]) or between the IAT and FPST (r = .11 [-.29 – .49]). To quantify the evidence for the correlations, we also computed Bayes factors (BF_{10}) comparing the alternative hypothesis ($|r| > 0$) against the null hypothesis ($r = 0$) using an uninformative prior (i.e., a uniform distribution ranging from -1 to 1) for correlations (Guan et al., 2020). Bayes factors indicate moderate evidence for a correlation of racial bias estimates between the WIT and FPST (BF_{10} = 4.00), but moderate evidence for the null hypothesis for the correlation of racial bias estimates between the IAT and WIT (BF_{10} = 0.33) and the IAT and FPST (BF_{10} = 0.32).This pattern of results indicates good correspondence between the WIT and FPST in racial bias estimates, but no correspondence in racial bias estimates between the IAT and the other two measures.

Exploring the influence of indirect measure order on PDP parameters

Modeling procedure. We counterbalanced the order in which participants completed indirect measures to minimize the influence of order effects on our primary analyses. Nevertheless, task order may have influenced participants' responses in important ways. With a consent form, three indirect measures, a demographics survey, and debriefing, our study took up to one hour for participants to complete. Thus, one possibility is that participants experienced fatigue on measures they completed later relative to measures they completed earlier in the testing session. Another possibility is that completing the IAT changed participants' responses on measures that came after it. On the WIT and FPST, participants can ignore the face images altogether and make correct responses solely by attending to weapons and harmless objects. In contrast, on the IAT participants must categorize the face images of Black and White men on 50% of trials. Consequently, the IAT uniquely relies on racial categories, and such category salience may carry over to influence participants' responses on subsequent tasks.

To investigate the influence of the order in which participants completed the indirect measures on process parameters, we exploratorily estimated two additional PDP-joint models that included measure order as a factorial covariate. The first model (indirect measure order model) estimated a separate PDP parameter for each indirect measure depending on whether it was completed first, second, or third. If participants experienced fatigue – in the form of depleted executive function and/or reduced motivation – as they completed the study, we might expect for Controlled process parameters to be lower on later tasks compared to earlier tasks; however we might expect for Automatic process parameters and/or racial bias estimates (quantified as the difference between Automatic process parameters for Black versus White targets) to be unaffected by potential fatigue effects (Govorun & Payne, 2006).

The second model (IAT position model) estimated a separate PDP parameter for each indirect measure depending on whether the IAT was completed first, second, or third. When participants complete the IAT, they are directly instructed to categorize Black and White male faces. In contrast, racial categories are incidental to participants' responses on the WIT and FPST, in which they are instructed to discern whether objects are weapons or tools. Consequently, the IAT may function as a race category salience manipulation, which could boost correspondence of racial bias estimates between the IAT and other measures that participants complete subsequently (Olson & Fazio, 2003). If completing the IAT increases racial category salience, we might expect for racial bias estimates (quantified as the difference between Automatic process parameters for Black versus White targets) to be higher on WITs and FPSTs completed after versus before the IAT. Furthermore – and extending beyond the predictions of the indirect measure order model – we might expect for racial bias estimate correlations among tasks to be stronger when participants completed the IAT first compared to later in the study.

We used the same Bayesian hierarchical latent-trait estimation approach here as we did with the PDP-joint model reported above. For both models, the Rubin-Gelman statistic was smaller than $\hat{R} < 1.05$ for all parameter estimates, indicating acceptable convergence of estimation chains. Both models fit the data well: indirect measure position model $p(T_1)$ =. 085, $p(T_2)$ = .452; IAT position model $p(T_1)$ = .495, $p(T_2)$ = .323.

Influence of indirect measure position on PDP parameters. In the indirect measure order model (Table 3), Controlled process parameters were lower for measures that participants completed later in the study. For all three measures, Bayes factors indicate at least strong evidence (all BF_{10} > 24.40) that Controlled process parameters are higher for measures completed first versus second or third. For the WIT and FPST the drop happened between the first and second measure, such that Controlled process parameters did not differ between measures completed second versus third. In contrast, for the IAT the drop happened between the second and third measure, such that Controlled process parameters did not differ between measures completed first versus second.

Our model produced a total of 27 parameters representing Automatic processes and 3 racial bias estimates. Of those 27 Automatic process parameters, Bayes factors associated with 22 of them (all BF_{10} < 0.22) indicate moderate to strong evidence for the null hypothesis (i.e., that estimates do not vary as a function of measure position). The Bayes factors associated with the remaining 5 tests (BF_{10} between 0.34 – 2.94) were equivocal, providing no clear evidence for either the alternative or null hypothesis. Bayes factors for all racial bias estimates (all BF_{10} < 0.09) indicate strong evidence for the null hypothesis.

Table 3

Mean parameter estimates by indirect measure position and Bayes factor for influence of indirect measure position.

	position of indirect measure						*Bayes factor (BF_{10}) of measure position influence*		
	1^{st} position		*2^{nd} position*		*3^{rd} position*		*1^{st} – 2^{nd}*	*1^{st} – 3^{rd}*	*2^{nd} – 3^{rd}*
Parameter	*Mean*	*95%-BCI*	*Mean*	*95%-BCI*	*Mean*	*95%-BCI*			
WIT	*N* = 123		*N* = 135		*N* = 114				
1. C	.89	[.82 – .97]	.68	[.61 – .75]	.67	[.60 – .75]	**32.13**	**24.40**	*0.09*
2. A_B	.55	[.47 – .63]	.53	[.45 – .60]	.55	[.48 – .63]	*0.08*	*0.08*	*0.09*
3. A_W	.43	[.35 – .51]	.49	[.42 – .57]	.46	[.38 – .53]	0.13	*0.09*	*0.09*
4. A_N	.53	[.46 – .60]	.43	[.37 – .50]	.41	[.35 – .48]	0.34	0.49	*0.07*
5. ΔA_{BW}	0.12	[0.01 – 0.23]	0.03	[-0.07 – 0.14]	0.10	[-0.01 – 0.20]	*0.09*	*0.06*	*0.07*
FPST	*N* = 115		*N* = 126		*N* = 111				
6. C	.94	[.89 – >.99]	.77	[.72 – .82]	.69	[.64 – .74]	**>100.00**	**>100.00**	0.35
7. A_B	.54	[.45 – .62]	.50	[.43 – .57]	.49	[.42 – .56]	*0.10*	*0.10*	*0.08*
8. A_W	.48	[.40 – .56]	.44	[.37 – .51]	.39	[.32 – .46]	*0.10*	0.19	0.11
9. A_N	.56	[.48 – .64]	.37	[.30 – .45]	.39	[.31 – .46]	2.94	2.44	*0.08*
10. ΔA_{BW}	0.06	[-0.05 – 0.17]	0.06	[-0.04 – 0.15]	0.10	[0.02 – 0.19]	*0.06*	*0.06*	*0.06*
IAT	*N* = 134		*N* = 111		*N* = 127				
11. C	.89	[.85 – .93]	.87	[.82 – .92]	.70	[.66 – .75]	0.11	**>100.00**	**>100.00**
12. A_B	.58	[51. – .66]	.50	[.42 – .58]	.52	[.45 – .58]	0.19	0.14	*0.08*
13. A_W	.53	[.46 – .61]	.44	[.37 – .52]	.39	[.32 – .46]	0.22	1.26	0.12
14. A_N	.38	[.31 – .46]	.45	[.37 – .52]	.39	[.32 – .46]	0.14	*0.08*	0.13
15. ΔA_{BW}	0.05	[-0.05 – 0.15]	0.06	[-0.05 – 0.16]	0.12	[0.03 – 0.22]	*0.06*	*0.09*	*0.06*

Note. BCI = Bayesian Credibility Interval. *C* = Controlled process parameter, *A* = Automatic process parameter, *B* = Black, *W* = White, *N* = neutral. Racial bias estimate: $\Delta A_{BW} = A_B – A_W$. Bayes factors are quantifying the evidence of the alternative hypothesis, that is a difference (Δ ≠ 0) to no difference (Δ = 0) by indirect measure position. Bayes factors (BF_{10}) >10, indicating at least strong evidence for the alternative hypothesis, are in bold. Bayes factors (BF_{10}) <.10, indicating at least strong evidence for the null hypothesis, are in italics.

Taken together, this model suggests that Automatic process parameters and racial bias estimates did not vary as a function of measure position.

Overall, this pattern of results shows that participants' Controlled process parameters decreased over the course of the experiment, which suggests fatigue or decreasing motivation effects. In contrast, Automatic process parameters and racial bias estimates did not vary as a function of measure position, which indicates that response bias and racial bias is not influenced by fatigue or motivation effects.

Variation in racial bias estimates and their correlations as a function of IAT position. In the IAT position model (Table 4) racial bias estimates were similar for the WIT and FPST independent of when participants completed the IAT (all $BF_{10} < 0.17$). Similarly, racial bias estimates for the IAT did not vary as a function of when participants completed it (all $BF_{10} < 0.08$).

To investigate the influence of IAT position on racial bias correlations among indirect measures, we transformed correlation estimates using Fisher's z-transformation and tested their difference against zero with a weakly informative prior $N(0,2)$. The size of racial bias correlations between WIT and FPST (IAT-first: r = .41 [.06 - .68]; IAT-second: r = .34 [-.01 - .65]; IAT-third: r = .36 [.03 - .65]) did not vary as a function of when participants completed the IAT (all BF_{10} <.09). Similarly, the position in which participants completed the IAT had no influence (all $BF_{10} < 0.07$) on the correlation of racial bias estimates between the IAT and either the WIT (IAT-first: r = .13 [-.25 - .49]; IAT-second: r = .15 [-.24 - .51]; IAT-third: r = .17 [-.18 - .52]) or the FPST (IAT-first: r = .14 [-.28 - .52]; IAT-second: r = .13 [-.30 - .53]; IAT-third: r = .17 [-.25 - .55]).

Taken together, these results indicate that IAT position does not influence the size of racial bias estimates or their correlations across indirect measures.

Table 4.

Mean parameter estimates by IAT position and Bayes factor for influence of IAT position.

	position of IAT						*Bayes factor (BF_{10}) of IAT position influence*		
	1st position (*N* = 134)		*2nd position* (*N* = 111)		*3rd position* (*N* = 127)		*1st – 2nd*	*1st – 3rd*	*2nd – 3rd*
Parameter	*Mean*	*95%-BCI*	*Mean*	*95%-BCI*	*Mean*	*95%-BCI*			
WIT									
1. C	.70	[.62 – .79]	.73	[.64 – .83]	.81	[.72 – .90]	0.13	0.38	0.22
2. A_B	.64	[.56 – .71]	.52	[.44 – .60]	.47	[.39 – .55]	0.34	1.79	0.11
3. A_W	.51	[.43 – .59]	.43	[.34 – .51]	.43	[.35 – .52]	0.19	0.17	*0.09*
4. A_N	.53	[.46 – .60]	.43	[.37 – .50]	.41	[.35 – .48]	*0.08*	0.38	0.61
5. ΔA_{BW}	0.12	[0.01 – 0.23]	0.09	[-0.02 – 0.21]	0.04	[-0.08 – 0.15]	*0.06*	*0.09*	*0.07*
FPST									
6. C	.74	[.68 – .81]	.79	[.72 – .86]	.86	[.80 – .93]	0.13	0.92	0.24
7. A_B	.53	[.46 – .60]	.51	[.43 – .59]	.48	[.41 – .56]	*0.08*	*0.10*	*0.09*
8. A_W	.44	[.37 – .52]	.38	[.30 – .45]	.48	[.40 – .56]	0.15	*0.09*	0.26
9. A_N	.41	[.33 – .49]	.42	[.34 – .51]	.47	[.39 – .55]	*0.01*	0.12	0.12
10. ΔA_{BW}	0.08	[-0.01 – 0.17]	0.14	[0.04 – 0.23]	0.06	[-0.09 – 0.11]	*0.07*	*0.08*	0.17
IAT									
11. C	.90	[.85 – .95]	.85	[.79 – .91]	.71	[.66 – .76]	0.20	**>100.00**	4.55
12. A_B	.60	[.52. – .67]	.49	[.41 – .57]	.51	[.44 – .58]	0.33	0.21	*0.08*
13. A_W	.55	[.48 – .62]	.44	[.36 – .51]	.39	[.32 – .46]	0.34	2.10	0.11
14. A_N	.40	[.32 – .48]	.43	[.35 – .51]	.39	[.31 – .46]	*0.10*	*0.09*	0.11
15. ΔA_{BW}	0.05	[-0.05 – 0.15]	0.05	[-0.05 – 0.16]	0.12	[0.03 – 0.22]	*0.06*	*0.08*	*0.08*

Note. BCI = Bayesian Credibility Interval. *C* = Controlled process parameter, *A* = Automatic process parameter, *B* = Black, *W* = White, *N* = neutral. Racial bias estimate: $\Delta A_{BW} = A_B - A_W$. Bayes factors are quantifying the evidence of the alternative hypothesis, that is a difference ($\Delta \neq 0$) to no difference ($\Delta = 0$) by indirect measure position. Bayes factors (BF_{10}) >10, indicating at least strong evidence for the alternative hypothesis, are in bold. Bayes factors (BF_{10}) <.10, indicating at least strong evidence for the null hypothesis, are in italics.

General Discussion

In the present research, we investigated correspondence among three indirect measures of racial bias: the Weapon Identification Task (WIT), the First-Person Shooter Task (FPST) and the Implicit Association Test (IAT). Previous research showing small or null correlations among these measures is confounded by differences in procedures or stimuli. In contrast, we aligned procedures and stimuli across all three measures, allowing them to vary only in structure. With the measures aligned in this way, we found that participants' responses corresponded across measures both in terms of accuracy-oriented controlled responding – replicating previous research – and also in terms of the tendency to respond "gun" versus "tool". That said, racial bias operationalized in terms of differences between Black and White Automatic process parameters did not correspond as consistently across measures: whereas racial bias estimates correlated moderately between the WIT and FPST, racial bias estimates did not correlate between the IAT and the other measures.

Moreover, the present research extends our understanding of racial bias as an individual difference. Indirect measures such as the WIT, FPST, and IAT routinely reveal racial bias effects in the aggregate (Payne et al., 2017; Payne & Correll, 2020), but the extent to which racial bias is a property of individuals has been a matter of considerable debate (e.g., Granados Samayoa & Fazio, 2017; Kurdi & Banaji, 2017). Low correlations among individuals' summary statistics of racial bias across measures along with consistent evidence of group-level racial bias in the aggregate has been interpreted to indicate that the construct assessed by indirect measures is a property of contexts rather than individuals (Payne et al., 2017). However, the high correlations among parameters across measures that we observed in the present research suggests that individuals may be able to rely on the same abilities across tasks - which would indicate that these measures tap into a common underlying construct (e.g., Ito et al., 2015; Olson & Fazio, 2003; Payne et al., 2005). Below we discuss how the patterns of correspondence across indirect measures we identified here provide insight into properties of individuals versus reveal important differences among the tasks themselves.

Correspondence across indirect measures of racial bias

Replicating previous research, Controlled process estimates correlated strongly across all indirect measures. This pattern of results indicates that participants' cognitive control abilities generalize across measures. These findings dovetail with other work showing that the WIT and FPST are related to executive functions like inhibition (Ito et al., 2015; Payne, 2005), and the IAT is related to updating and task switching (Ito et al., 2015; Klauer et al., 2010). Hence, the present research joins a body of literature connecting indirect measures of racial bias with a broad constellation of executive functions and other higher cognitive abilities.

An important finding to emerge from the present research is that Automatic process estimates correlated moderately to strongly across all indirect measures. By aligning measures across procedural features and stimuli and using MPT modeling to disentangle the joint contributions of multiple processes, we showed that the constructs reflected in each of the Automatic process estimates – one for each target group – correspond well across measures. The most straightforward interpretation of the Automatic process parameter is that it reflects participants' preference to respond with gun in comparison to tool. This preference can reflect the threat stereotypes that indirect measures are configured to assess, but it can also reflect recoding (Meissner & Rothermund, 2013), a simple hand-side preference, other dispositional characteristics like a threat-related attention bias driven by anxiety (Bar-Haim et al., 2007), a general threat superiority effect for gun targets (Rivera-Rodriguez et al., 2021; Subra et al., 2018), or guessing. Importantly, our findings show that this response tendency generalizes across indirect measures, which suggests that the preference to respond with gun versus tool may be a property of participants.

Additionally, Automatic process parameters provide insight into racial biases when we compare parameters estimated from trials that include Black versus White faces. Across all measures, participants demonstrated higher Automatic process parameters for “Black” than “White” targets, corresponding with the cultural stereotype of Black men as dangerous – which, in turn, provides a degree of validity evidence that these measures can assess their intended construct. Nevertheless, future research is necessary to better understand factors that moderate correspondence among processes that contribute to racially biased responses on these measures.

Differences between indirect measures of racial bias

Though Controlled and Automatic process parameters correlated moderately-to-strongly across all indirect measures, correlations between parameter estimates in isolation provide little insight into the extent to which Black versus White people are differentially stereotyped as threatening. When we operationalized racial bias as the difference between Black and White Automatic process parameters, we found a more nuanced pattern of results. Specifically, racial bias operationalized in terms of the difference between Automatic process parameters correlated moderately between the WIT and FPST, but not between the IAT and the other two measures.

These findings dovetail with previous work demonstrating correspondence between the WIT and FPST (Payne & Correll, 2020; Ito et al., 2015). The WIT and FPST are both configured to assess the influence of race on fast behavioral responses to guns and non-threatening target objects. Furthermore, both seminal articles introducing these paradigms explicitly refer to the behavior of police officers in shooting unarmed Black men (Correll et al., 2002; Payne, 2001; Payne & Correll, 2020). Although the traditional configurations of the WIT and FPST differ in stimulus presentation order (sequential versus concurrent) and style (e.g., isolated stimulus in the center of the screen versus full-body images with different backgrounds and different screen positions), they appear to assess similar constructs. In both measures, participants view pictures of Black and White male faces, but do not need to attend to race – or to the faces at all – to make a correct response on the WIT (i.e., categorizing “gun”, “tool”) or the FPST (i.e., decide to “shoot”, “don’t shoot”). Interestingly, in the present research, response tendencies towards the response “gun” were the same between White faces and the neutral face outline but enhanced for Black faces for both the WIT and FPST. This finding may indicate that only Black faces elicit an additional response tendency towards “gun”, which aligns with claims that the WIT and FPST assess threat stereotypes based on racial exemplars (Olson & Fazio, 2003; Livingston & Brewer, 2002).

In contrast, lack of correspondence between the IAT and the other two indirect measures may reflect heightened category salience and direct contrasting of race categories as a function of the structure of the IAT. Because of its dual-categorization structure, only the IAT requires participants to attend to race to make a correct response to half of trials (i.e., to face stimuli). Because IAT task instructions differ from the other two measures in this regard, they may lead participants to rely on response strategies that are otherwise unnecessary or irrelevant on the other two measures. For example, recoding is a task-simplification strategy by which participants can conceptually reduce the IAT from four categories to two. Recoding guns and Black male faces into a single “threatening” category, and recoding tools and White male faces into a single “non-threatening” category, simplifies the IAT to a binary decision: threatening or not? Such simplification might be easier for one IAT block type (e.g., Black-guns) compared to the other block type (e.g., Black-tools), resulting in less errors in one than the other, which itself represents a racial bias effect (Meissner & Rothermund, 2013). In the PDP, these recoding effects would map on the Automatic process parameter together with other stereotype influences.[11] Response strategies based on race categories may explain the observed Automatic process parameter pattern for the three different IAT block pairings (i.e., Black versus White male faces; Black versus neutral outline; White male faces versus neutral outline). Whereas the IAT block directly

[11] Meissner and Rothermund (2013) implemented the ReAL model as MPT model which estimates the contribution of recoding along with controlled and automatic processes to IAT responses. However, the ReAL model is tailored to the IAT and requires specific procedures to produce valid parameter estimates (Calanchini et al., 2021).

contrasting Black and White faces produced the expected racial bias effect (i.e., stronger threat stereotypes for Black than White male faces), Black and White faces both elicited a threat response tendency in blocks contrasted with the neutral face outline. This pattern of results suggests that male faces in general (i.e., the superordinate category shared by both Black and White face stimuli) are perceived as more threatening than the neutral face outline. Thus, our findings would seem to align with other researchers who conclude that the IAT assesses threat stereotypes based on racial categories (De Houwer, 2001; Olson & Fazio, 2003).

Overall, all three of the indirect measures we investigated in the present research appear to assess racial stereotypes linking Black versus White male faces with threat. However, task specific procedural differences may play a large role in racial bias expression even when measures are aligned on several procedural dimensions (i.e., stimuli material, trial number, response time). Consequently, each indirect measure may reflect a specific facet of racial bias. As a sequential-priming task, the WIT taps into the early interference of threat-stereotypes in object identification (Klauer & Voss, 2008; Laukenmann et al., 2023; Payne et al., 2005). In contrast, the FPST assesses the behavioral decision to “shoot” or “don’t shoot” at an armed or unarmed person, which ties the behavioral decision of target object identification to the person holding it (Correll et al., 2015). The IAT, in its dual-categorization structure, explicitly contrasts the two reference group categories, which may highlight racial categories in participants’ responding (De Houwer, 2001; Meissner & Rothermund, 2013). Hence, these measures would seem to tap into different aspects of racial bias.

Summary statistics of racial bias versus MPT model parameters

Racial bias effects emerged across all three indirect measures in terms of accuracy-based summary statistics, such that all three indicate that guns are associated more strongly with Black than White male faces in our sample of participants. However, these accuracy-based summary statistics only corresponded weakly between the WIT and FPST, and not between the IAT and the other two measures. Latency-based summary statistics revealed a racial bias effect only for the IAT, but the pattern of correlations among measures was similar to what we observed with accuracy-based summary statistics, such that the WIT and FPST corresponded weakly with each other but not at all with the IAT.

In contrast, MPT modeling revealed correspondence among processes across measures, such that both Controlled and Automatic process parameters correspond well with conceptually analogous parameters across measures. In addition, racial bias estimates operationalized in terms of differences between model parameters correlated moderately between the WIT and FPST, demonstrating that MPT modeling can reveal a stronger correspondence between measures that summary statistics may otherwise obscure. Additionally, model-based analyses localized fatigue effects on Control process parameters, with Automatic parameters remaining consistent across measures. This finding aligns with previous research demonstrating that Control, but not Automatic process parameters and their derived racial bias estimates, may be tempered by prior mental exertion (Govorun & Payne, 2006). Taken together, the present research highlights one way in which summary statistics of racial bias can belie process-level relationships, and at the same time demonstrates the value of the theoretical precision provided by MPT models.

Limitations

Despite the strengths of the present research, it is also limited in some ways. For example, to test our primary research question about correspondence among indirect measures of racial bias, we aligned all three in terms of stimuli and procedural features. Such alignment positioned us to make apples-to-apples comparisons across measures. However, this comparability came at the cost of modifying these measures from their traditional form, which may affect the generalizability of our findings to how most researchers use these measures.

With that said, readers interested in the extent to which the traditional versions of the WIT, FPST, and IAT correspond with each other can consult Ito et al. (2015), who examined that very question. Moreover, though we did not rely on traditional versions of each of these measures in the present research, we varied them all on procedural dimensions that have varied in previous research. Indeed, aligning measures on these dimensions allowed us to control for potential confounding stimuli effects (Correll et al., 2002; 2011) and to obtain comparable levels of error rates (Ito et al., 2015).

Even what is arguably the biggest procedural change we implemented in the present research – the inclusion of the neutral face outline along with Black and White faces – has precedent in this literature (Rivers, 2017). The inclusion of a neutral face outline did not change the traditional structure of the WIT or FPST, but it did change the structure of the IAT because it required additional sets of blocks (i.e., Black/White, Black/neutral, White/neutral). However, this expanded IAT format closely aligns with an existing version of the IAT: the multi-category IAT (Axt et al., 2014). Moreover, the expected bias effects – indicating stronger threat stereotypes for Black versus White men – emerged across all three measures, which suggest that these changes did not significantly alter the measures. Nevertheless, previous research has demonstrated that task procedures (such as number of trials) affect the extent to which different cognitive processes influence responses on indirect measures of racial bias (Calanchini et al., 2021), so future research should continue to investigate the role of stimulus and structural effects in indirect measures of racial bias.

Conclusion

The present research makes two key contributions to social cognition literature. The first contribution is in demonstrating that three commonly used indirect measures of racial bias correspond well on controlled responding and automatic response tendencies when procedurally aligned. The second contribution is that racial bias may only correspond moderately across measures, suggesting that each of these measures may assess specific aspects of racial threat stereotypes. Taken together, this work offers a useful template for future researchers who seek to incorporate multiple operationalizations of racial bias in their work and gain theoretically precise insight into the processes that contribute to responses on indirect measures of racial bias.

References

Axt, J. R., Ebersole, C. R., & Nosek, B. A. (2014). The Rules of Implicit Evaluation by Race, Religion, and Age. *Psychological Science, 25*(9), 1804–1815. https://doi.org/10.1177/0956797614543801

Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, M. J., & van IJzendoorn, M. H. (2007). Threat-related attentional bias in anxious and nonanxious individuals: A meta-analytic study. *Psychological Bulletin, 133*(1), 1–24. https://doi.org/10.1037/0033-2909.133.1.1

Bar-Anan, Y., & Nosek, B. A. (2014). A comparative investigation of seven indirect attitude measures. *Behavior research methods*, *46*(3), 668–688. https://doi.org/10.3758/s13428-013-0410-6

Bishara, A. J., & Payne, B. K. (2009). Multinomial process tree models of control and automaticity in weapon misidentification. *Journal of Experimental Social Psychology*, *45*(3), 524-534. https://doi.org/10.1016/j.jesp.2008.11.002

Calanchini, J. (2020). How multinomial processing trees have advanced, and can continue to advance, research using implicit measures. *Social Cognition*, *38*(Supplement), s165-s186. https://doi.org/10.1521/soco.2020.38.supp.s165

Calanchini, J., Meissner. F., & Klauer, K. C. (2021). The role of recoding in implicit social cognition: Investigating the scope and interpretation of the ReAL model for the implicit association test. *PLOS ONE, 16(4)*, e0250068. https://doi.org/10.1371/journal.pone.0250068

Cohen, J. (1992). A power primer. *Psychological Bulletin, 112*(1), 155–159.

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*(4), 469–487. https://doi.org/10.1037/0022-3514.89.4.469

Corneille, O., & Hütter, M. (2020). Implicit? What do you mean? A comprehensive review of the delusive implicitness construct in attitude research. *Personality and Social Psychology Review*, *24*(3), 212-232. https://doi.org/10.1177/1088868320911325

Correll, J., Hudson, S. M., Guillermo, S., & Ma, D. S. (2014). The Police Officer's Dilemma: A Decade of Research on Racial Bias in the Decision to Shoot: The Police Officer's Dilemma. *Social and Personality Psychology Compass*, *8*(5), 201–213. https://doi.org/10.1111/spc3.12099

Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology*, *83*(6), 1314–1329. https://doi.org/10.1037/0022-3514.83.6.1314

Correll, J., Park, B., Judd, C. M., Wittenbrink, B., Sadler, M. S., & Keesee, T. (2007). Across the thin blue line: Police officers and racial bias in the decision to shoot. *Journal of Personality and Social Psychology*, *92*(6), 1006–1023. https://doi.org/10.1037/0022-3514.92.6.1006

Correll, J., Wittenbrink, B., Crawford, M. T., & Sadler, M. S. (2015). Stereotypic vision: How stereotypes disambiguate visual stimuli. *Journal of Personality and Social Psychology*, *108*(2), 219–233. https://doi.org/10.1037/pspa0000015

Correll, J., Wittenbrink, B., Park, B., Judd, C. M., & Goyle, A. (2011). Dangerous enough: Moderating racial bias with contextual threat cues. *Journal of experimental social psychology*, *47*(1), 184-189. https://doi.org/10.1016/j.jesp.2010.08.017

Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit Attitude Measures: Consistency, Stability, and Convergent Validity. *Psychological Science, 12*(2), 163–170. https://doi.org/10.1111/1467-9280.00328

De Houwer, J. (2001). A structural and process analysis of the Implicit Association Test. *Journal of Experimental Social Psychology*, *37*(6), 443–451. https://doi.org/10.1006/jesp.2000.1464

Erdfelder, E., Auer, T.-S., Hilbig, B. E., Aßfalg, A., Moshagen, M., & Nadarevic, L. (2009). Multinomial Processing Tree Models: A Review of the Literature. *Zeitschrift Für Psychologie / Journal of Psychology*, *217*(3), 108–124. https://doi.org/10.1027/0044-3409.217.3.108

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. https://doi.org/10.3758/BF03193146

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*(6), 1013–1027. https://doi.org/10.1037/0022-3514.69.6.1013

Foroni, F., & Bel-Bahar, T. (2010). Picture-IAT versus Word-IAT: level of stimulus representation influences on the IAT. *European Journal of Social Psychology*, *40*(2), 321–337. https://doi.org/10.1002/ejsp.626

Frenken, M., Hemmerich, W., Izydorczyk, D., Scharf, S., & Imhoff, R. (2022). Cognitive processes behind the shooter bias: Dissecting response bias, motor preparation and information accumulation. *Journal of Experimental Social Psychology*, *98*, 104230. https://doi.org/10.1016/j.jesp.2021.104230

Gawronski, B., De Houwer, J., & Sherman, J. W. (2020). Twenty-five years of research using implicit measures. *Social Cognition*, *38*(Supplement), s1-s25. https://doi.org/10.1521/soco.2020.38.supp.s1

Gawronski, B., Cunningham, W. A., LeBel, E. P., & Deutsch, R. (2010). Attentional influences on affective priming: Does categorisation influence spontaneous evaluations of multiply categorisable objects? *Cognition and Emotion*, *24*(6), 1008–1025. https://doi.org/10.1080/02699930903112712

Glaser, J., & Knowles, E. D. (2008). Implicit motivation to control prejudice. *Journal of Experimental Social Psychology*, *44*(1), 164–172. https://doi.org/10.1016/j.jesp.2007.01.002

Govorun, O., & Payne, B. K. (2006). Ego—depletion and prejudice: Separating automatic and controlled components. *Social Cognition*, *24*(2), 111-136. https://doi.org/10.1521/soco.2006.24.2.111

Granados Samayoa, J. A., & Fazio, R. H. (2017). Who Starts the Wave? Let's Not Forget the Role of the Individual. *Psychological Inquiry*, *28*(4), 273–277. https://doi.org/10.1080/1047840X.2017.1373554

Greenwald, A. G., & Lai, C. K. (2020). Implicit social cognition. *Annual review of psychology*, *71*, 419–445. https://doi.org/10.1146/annurev-psych-010419-050837

Greenwald, 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*(6), 1464–1480. https://doi.org/10.1037/0022-3514.74.6.1464

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*(2), 197–216. https://doi.org/10.1037/0022-3514.85.2.197

Guan, M., Stokes, R., Vandekerckhove, J., & Lee, M. D. (2020). A cognitive modeling analysis of risk in sequential choice tasks. *Judgment and Decision making*, *15*(5), 823-850. https://doi.org/10.1017/S1930297500007956

Heck, D. W., Arnold, N. R., & Arnold, D. (2018). TreeBUGS: An R package for hierarchical multinomial-processing-tree modeling. *Behavior Research Methods*, *50*(1), 264–284. https://doi.org/10.3758/s13428-017-0869-7

Henninger, F., Shevchenko, Y., Mertens, U. K., Kieslich, P. J., & Hilbig, B. E. (2022). lab. js: A free, open, online study builder. *Behavior Research Methods*, *54*(2), 556-573. https://doi.org/10.3758/s13428-019-01283-5

Hütter, M.**,** & Klauer, K. C. (2016). Applying processing trees in social psychology. *European Review of Social Psychology, 27*, 116–159. https://doi.org/10.1080/10463283.2016.1212966

Ito, T. A., Friedman, N. P., Bartholow, B. D., Correll, J., Loersch, C., Altamirano, L. J., & Miyake, A. (2015). Toward a comprehensive understanding of executive cognitive function in implicit racial bias. *Journal of Personality and Social Psychology*, *108*(2), 187–218. https://doi.org/10.1037/a0038557

Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*(5), 513–541. https://doi.org/10.1016/0749-596X(91)90025-F

Jeffreys, H. (1961). *Theory of probability* (3rd ed.). Oxford University Press.

Klauer, K. C. (2010). Hierarchical Multinomial Processing Tree Models: A Latent-Trait Approach. *Psychometrika*, *75*(1), 70–98. http://dx.doi.org/10.1007/s11336-009-9141-0

Klauer, K. C., Dittrich, K., Scholtes, C., & Voss, A. (2015). The invariance assumption in process-dissociation models: An evaluation across three domains. *Journal of Experimental Psychology: General*, *144*(1), 198–221. https://doi.org/10.1037/xge0000044

Klauer, K. C., Schmitz, F., Teige-Mocigemba, S., & Voss, A. (2010). Understanding the role of executive control in the Implicit Association Test: Why flexible people have small IAT effects. *Quarterly Journal of Experimental Psychology*, *63*(3), 595–619. https://doi.org/10.1080/17470210903076826

Klauer, K. C., & Voss, A. (2008). Effects of Race on Responses and Response Latencies in the Weapon Identification Task: A Test of Six Models. *Personality and Social Psychology Bulletin*, *34*(8), 1124–1140. https://doi.org/10.1177/0146167208318603

Kurdi, B., & Banaji, M. R. (2017). Reports of the Death of the Individual Difference Approach to Implicit Social Cognition May Be Greatly Exaggerated: A Commentary on Payne, Vuletich, and Lundberg. *Psychological Inquiry*, *28*(4), 281–287. https://doi.org/10.1080/1047840X.2017.1373555

Laukenmann, R., Erdfelder, E., Heck, D. W., & Moshagen, M. (2023). Cognitive processes underlying the weapon identification task: A comparison of models accounting for both response frequencies and response times. *Social Cognition*, *41*(2), 137–164. https://doi.org/10.1521/soco.2023.41.2.137

Lindsay, D. S., & Jacoby, L. L. (1994). Stroop process dissociations: The relationship between facilitation and interference. *Journal of Experimental Psychology: Human Perception and Performance, 20*(2), 219–234. https://doi.org/10.1037/0096-1523.20.2.219

Livingston, R. W., & Brewer, M. B. (2002). What are we really priming? Cue-based versus category-based processing of facial stimuli. *Journal of Personality and Social Psychology*, *82*(1), 5–18. https://doi.org/10.1037/0022-3514.82.1.5

McConahay, J. B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In J. F. Dovidio & S. L. Gaertner (Eds.), *Prejudice, discrimination, and racism* (pp. 91–125). Academic Press.

Meissner, F., & Rothermund, K. (2013). Estimating the contributions of associations and recoding in the Implicit Association Test: The ReAL model for the IAT. *Journal of Personality and Social Psychology, 104*(1), 45–69. https://doi.org/10.1037/a0030734

Mekawi, Y., & Bresin, K. (2015). Is the evidence from racial bias shooting task studies a smoking gun? Results from a meta-analysis. *Journal of Experimental Social Psychology*, *61*, 120–130. https://doi.org/10.1016/j.jesp.2015.08.002

Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review, 106*(1), 3–19. https://doi.org/10.1037/0033-295X.106.1.3

Olson, M. A., & Fazio, R. H. (2003). Relations between implicit measures of prejudice: What are we measuring? *Psychological Science*, *14*(6), 636–639. https://doi.org/10.1046/j.0956-7976.2003.psci_1477.x

Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, *81*(2), 181–192. https://doi.org/10.1037/0022-3514.81.2.181

Payne, B. K. (2005). Conceptualizing Control in Social Cognition: How Executive Functioning Modulates the Expression of Automatic Stereotyping. *Journal of Personality and Social Psychology*, *89*(4), 488–503. https://doi.org/10.1037/0022-3514.89.4.488

Payne, B. K., & Correll, J. (2020). Race, weapons, and the perception of threat. In *Advances in experimental social psychology* (Vol. 62, pp. 1-50). Academic Press. https://doi.org/10.1016/bs.aesp.2020.04.001

Payne, B. K., Shimizu, Y., & Jacoby, L. L. (2005). Mental control and visual illusions: Toward explaining race-biased weapon misidentifications. *Journal of Experimental Social Psychology*, *41*(1), 36–47. https://doi.org/10.1016/j.jesp.2004.05.001

Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017). The Bias of Crowds: How Implicit Bias Bridges Personal and Systemic Prejudice. *Psychological Inquiry*, *28*(4), 233–248. https://doi.org/10.1080/1047840X.2017.1335568

Phills, C. E., Kawakami, K., Tabi, E., Nadolny, D., & Inzlicht, M. (2011). Mind the gap: Increasing associations between the self and blacks with approach behaviors. *Journal of Personality and Social Psychology, 100*(2), 197–210. https://doi.org/10.1037/a0022159

Plant, E. A., Peruche, B. M., & Butz, D. A. (2005). Eliminating automatic racial bias: Making race non-diagnostic for responses to criminal suspects. *Journal of Experimental Social Psychology*, *41*(2), 141–156. https://doi.org/10.1016/j.jesp.2004.07.004

Riefer, D. M., & Batchelder, W. H. (1988). Multinomial modeling and the measurement of cognitive processes. *Psychological Review, 95*(3), 318–339. https://doi.org/10.1037/0033-295X.95.3.318

Rivera-Rodriguez, A., Sherwood, M., Fitzroy, A. B., Sanders, L. D., & Dasgupta, N. (2021). Anger, race, and the neurocognition of threat: attention, inhibition, and error processing during a weapon identification task. *Cognitive research: principles and implications*, *6*, 1–27. https://doi.org/10.1186/s41235-021-00342-w

Rivers, A. M. (2017). The Weapons Identification Task: Recommendations for adequately powered research. *PLOS ONE*, *12*(6), e0177857. https://doi.org/10.1371/journal.pone.0177857

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General, 104*(3), 192–233. https://doi.org/10.1037/0096-3445.104.3.192

Schimmack, U. (2021). The Implicit Association Test: A Method in Search of a Construct. *Perspectives on Psychological Science, 16*(2), 396–414. https://doi.org/10.1177/1745691619863798

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84*(2), 127–190. https://doi.org/10.1037/0033-295X.84.2.127

Sriram, N., & Greenwald, A. G. (2009). The brief implicit association test. *Experimental Psychology*, *56*(4), 283–294. https://doi.org/10.1027/1618-3169.56.4.283

Subra, B., Muller, D., Fourgassie, L., Chauvin, A., & Alexopoulos, T. (2018). Of guns and snakes: testing a modern threat superiority effect. *Cognition and emotion*, *32*(1), 81–91. https://doi.org/10.1080/02699931.2017.1284044

Unkelbach, C., Forgas, J. P., & Denson, T. F. (2008). The turban effect: The influence of Muslim headgear and induced affect on aggressive responses in the shooter bias paradigm. *Journal of Experimental Social Psychology*, *44*(5), 1409–1413. https://doi.org/10.1016/j.jesp.2008.04.003

Unkelbach, C., Goldenberg, L., Müller, N., Sobbe, G. & Spannaus, N. (2009). A shooter bias in Germany against people wearing Muslims headgear. *Revue internationale de psychologie sociale*, 22, 181–201. https://www.cairn.info/revue--2009-3-page-181.htm.

Volpert-Esmond, H. I., Scherer, L. D., & Bartholow, B. D. (2020). Dissociating automatic associations: Comparing two implicit measurements of race bias. *European Journal of Social Psychology*, *50*(4), 876–888. https://doi.org/10.1002/ejsp.2655

Wagenmakers, E. J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive psychology*, *60*(3), 158-189. https://doi.org/10.1016/j.cogpsych.2009.12.001