**Cognitive processes behind the shooter bias: Dissecting response bias, motor
preparation and information accumulation**

Frenken, Marius¹, Hemmerich, Wanja¹, Izydorczyk, David², Scharf, Sophie², & Imhoff,
Roland¹

¹Johannes Gutenberg-University Mainz, Germany

²University of Mannheim, Germany

Accepted Manuscript:

Frenken, M., Hemmerich, W., Izydorczyk, D., Scharf, S., & Imhoff, R. (in press). Cognitive processes behind the shooter bias: Dissecting response bias, motor preparation and information accumulation. *Journal of Experimental Social Psychology*.

September 2021

Abstract

A rich body of research points to racial biases in so-called police officer dilemma tasks: participants are generally faster and less error-prone to “shoot” (vs. not “shoot”) Black (vs. White) targets. In three experimental (and two supplemental) studies (total $N = 914$), we aimed at examining the cognitive processes underlying these findings under fully standardized conditions. To be able to dissect a-priori decision bias, biased information processing and motor preparation, we rendered video sequences of virtual avatars that differed in nothing but the tone of their skin. Modeling the data via drift diffusion models revealed that the threat of a social group can be explicitly learned and mapped accordingly on an a-priori response bias within the model (Study 1). Studies 2 and 3 replicated the racial shooter bias as apparent in faster reaction times in stereotype-consistent trials. This, however, appears to result from stereotype-consistent motoric preparations and execution readiness, but not from pre-judicial threat biases. The results have implications especially for automatic stereotypes in the public.

Keywords: implicit stereotyping, shooter task, shooter bias, racial bias, diffusion model

Cognitive processes behind the shooter bias: Dissecting response bias, motor preparation and information accumulation

Over the past years—and especially after the recent murder of George Floyd in the Summer of 2020—social movements acting against White supremacy and violence against Black people like *BlackLivesMatter* have gained momentum in global dimensions: in Germany, where the BlackLivesMatter demonstrations were among the largest outside the US in 2020, claims for investigations of *institutional racism* against different minorities—as in the alleged racial profiling conducted by the German police against people of African descent (UN Human Rights, 2017)—became part of the public debate (Knight, 2020; Safronova, 2020). Also other forms of discrimination and racism against Black people have been registered in Germany for a long time (for a summary about Black life in Germany, see Wandert et al., 2009). One approach to tackle stereotypical biases empirically is the First Person Shooter Task paradigm (FPST, Correll et al., 2002) as an implicit measure of stereotypical social cognition.

Indeed, numerous laboratory studies have shown that the race of a displayed person can cause differences in measurable (shooting) behavior (for a meta-analysis, see Mekawi & Bresin, 2015). The robustness of this effect was corroborated in different populations (e.g. Germany; Essien et al., 2017), with different target features (e.g. Muslim headgear instead of race; Unkelbach et al., 2008), and regardless of the social context (Correll et al., 2002). Its independence of explicit prejudice (Correll et al., 2002) suggests that the shooter bias is a spontaneous decision-making behavior in an indirect performance task as fast, less volatile decisions are less susceptible to post-perceptual effects, such as faking and social desirability. The present research seeks to explore the underlying cognitive mechanisms of the shooter bias in a predominantly German sample while also maximizing experimental standardization to rule out alternative methodological explanations based on confounds in materials.

The Shooter paradigm and its' standardization

In classic FPSTs, participants are tasked *to shoot* armed targets and *to not shoot* unarmed targets of either Black or White ethnicity as quickly as possible. Previous meta-analytic results indicate a racial bias (Mekawi & Bresin, 2015): reaction times are shorter for Black males holding a gun and prolonged for Black males holding a harmless object compared to White target males while at the same time the shooting threshold is generally more liberal for Black targets (i.e., more incorrect classifications as armed targets). Overall, studies confirm the interaction effect of race and weapon on reaction times as well as an increase in shooting errors for unarmed Black compared to White targets in lay populations but demonstrated mixed findings in samples of police officers (Cox & Devine, 2016).

Although most published studies either re-used the original or adopted similar stimuli to the ones originally used by Correll et al. (2002), these images lack full standardization and contain possible confounds and uncontrolled noise. Primarily, the depicted target persons differ in appearance, clothing, posture, and are positioned in front of multiple backgrounds and landscapes. These varying factors are not fully crossed over the experimental conditions and could thus be confounded with target race. Furthermore, the physical contrast between the object in hand and the skin color has not been addressed. Dark objects may be harder to discern against darker skin color, introducing potential alternative explanations based solely on color contrast alone. Another potential source of noise is the different style of clothing. To the extent that some Black targets wear stereotypically Black attires (e.g., wide t-shirts), effects are ambiguous as to whether they rest on race (as most prominently indicated by skin tone) or on a latent cascade of tertiary markers of stereotypical threat. Paying attention to such confounds is necessary as the sensitivity of the pattern of results to moderating variables and methodological variations appears to be of importance for the shooter bias (Cox et al., 2014).

Stereotypes and intergroup processes behind the shooter bias

Given the uniform pattern of findings regarding faster shoot decisions (respectively slower non-shoot decisions) for Black compared to White targets in lay populations, it is important to gain further insights into the underlying procedural architecture of this effect. Some authors have pointed to the relevance that Black people are stereotypically associated with threat, danger, and crime (Eberhardt et al., 2004; Lundberg et al., 2018; Pickel & Sneyd, 2018) which is considered as an explanation for the racial bias in the shooter paradigm (Correll et al., 2011; Plant et al., 2011). From the perspective of cultural stereotypes about outgroups, the activation of a group category (which is stereotypically perceived as threatening) facilitates aggressive and stereotype-consistent responses.

This interaction is supported by several earlier studies: racial biases are more pronounced when clothing or environmental cues are perceived as threatening (Correll et al., 2011; Kahn & Davies, 2017); the mere activation of threat-related social categories results in more aggressive responses (Mange et al., 2012); the bias is absent for stereotypically non-threatening social groups such as Asian or female targets (Miller et al., 2012; Plant et al., 2011; Sadler et al., 2012). Mekawi et al. (2019) found under experimental conditions that the dehumanization of African-Americans led to accelerated shooting decisions, relative to the level of perceived threat. Similarly, the neural foundation of this racial bias is believed to be an early automatic attention process related to an induced perceived threat (Correll et al., 2006; Fleming et al., 2010), preparing for a more aggressive response compared to neutral targets. Vice versa, Kahn et al. (2016) provided evidence from a reversed ingroup perspective that stereotypes about Whites can have protective effects.

There is no clear answer to the question of whether an established stereotype of threat is necessary or if the mere perception of outgroup membership suffices to elicit the outgroup bias. On the one hand, the shooter bias emerges in a minimal group paradigm without

preexisting stereotypes (Miller et al., 2012), yet, priming participants with a task-irrelevant outgroup only then leads to racial shooter biases if the outgroup is stereotypically perceived as threatening (Mange et al., 2016). Individual factors such as racial ingroup identity and chronic beliefs about interpersonal danger are associated with shooting behavior (Kenworthy et al., 2011; Miller et al., 2012) and might play a moderating role here. In a nutshell, the mere outgroup categorization can sometimes lead to biased shooting behavior, but the effect is more reliable when the outgroup is associated with threat.

Differentiating the underlying processes

While these are important first steps in understanding that intergroup processes and their related threat stereotypes contribute to the racial shooter bias, it remains unclear what kind of stereotypes and what kind of process mediates this effect. In the following, we introduce three likely stereotypes that may plausibly align with the underlying, institutionalized notion (the association between *Black* and *danger*), but are likely not identical and have different implications at the process level.

It is conceivable that people hold *frequency stereotypes* (i.e., a higher proportion of Black people will carry a gun) and hence favor a prejudiced decision before processing information relevant for the task. The frequency stereotype would suggest that people expect an association between gun ownership and skin color and might see it as instrumental to prepare for a shoot decision if they see a Black person (e.g., because of experiences from news and media which have been found to be biased towards Black people, see Sommers et al., 2006). To that end, Johnson and Chopik (2019) found a correlational relationship suggesting that the frequency bias can account for (mainly explicit) stereotypes. They found that state-level violent crime rates among White and Black Americans were associated with the strength of the Black-violence stereotype. Altogether, frequency stereotypes point to a

prior decision bias in favor of a stereotype-concordant option before having considered any additional evidence relevant for the decision.

The *typicality hypothesis* (i.e., the typical Black person is threatening) would potentially translate more into a situation whereby gun and Black skin are independently both seen as general cues of threat, whereas harmless objects and White skin are believed to be innocuous. If both cue types are in agreement, concordant responses will be accelerated as they activate overlapping conceptual clusters, whereas incompatible cues will impede responses as they activate competing clusters. This would not have implications for a prior decision bias but for task-relevant information processing (i.e., interpretation and detection of information).

Another possible path of how the stereotypical race-threat association impacts shooting behavior and accelerates stereotype-consistent responses is a *response execution bias*. Here, stereotypes influence the peripheral, non-decisional routes, i.e. the action control for the motoric response (Payne et al., 2005), and possibly initiate preparations on a motoric—not on a cognitive-decisional level as in both other possible hypotheses. For example, seeing a Black person might make observers tighten their grip on the trigger to facilitate “shoot” but delay “no-shoot” decisions. All three given kinds of processes can be separated by process diffusion models.

Applying the drift diffusion model

Which cognitive mechanisms drive stereotypical-consistent behavior and what is their psychological meaning? Drift diffusion models (see, Ratcliff & McKoon, 2008) offer insights into hypothesized cognitive processes for simple two-choice decision tasks by evaluating responses and reaction times conjointly, with frequent applications in cognitive sciences (Ratcliff et al., 2016). Further advances extended the mathematical framework of drift

diffusion models to include Bayesian and hierarchical estimators (see Methods section in Study 1 and Supplements).

Drift diffusion models assume that participants process information over time until they reach a decision threshold for one of two options (see Figure 1) and make a decision. The models consist of different parameters to describe possible influences on the decision-making process, separately estimated for all experimental conditions. The resulting parameters can be assigned to cognitive processes and compared between experimental conditions. We will briefly review the most important parameters in the following paragraphs.

Figure 1. The decision process and parameters of the drift diffusion model. To reach one of the two decision thresholds, participants accumulate information relevant to the decision over time. Parameters are the speed of accumulation (δ), the distance between the thresholds (α), the starting point of the accumulation process (β , relative to α) and the non-decision time (τ). The non-decision time represents all processes that are not part of the decision and—in contrast to the figure—can occur simultaneously at any point in the decision process.

The distance between both decision thresholds (α) represents the amount of information required by an individual to reach a decision. The greater the distance between the decision options, the more evidence is needed for a decision and the more cautious the participants. For example, participants may try to avoid appearing biased by responding particularly cautiously when the avatar is Black and, accordingly, possess a higher boundary separation for Black than for White avatars.

The rate of information accumulation (δ), the *drift rate*, describes the speed of an individual's information uptake. Steeper drift rates indicate more rapid information processing, for instance because two cues converge in the implied decision. For the current

study, the identification of an object may be enhanced when the stimulus object matches the stereotype of the avatar (e.g., Black – gun). Vice versa, identification of objects is delayed when object and stereotype are incompatible.

The relative *starting point* ($\beta \cdot \alpha$) is an a-priori bias in favor of one of the options before any information relevant to the decision is processed. Biased participants likely start their information uptake closer to one of the decision thresholds depending on the targets' race, for instance, for Black avatars closer to the shoot threshold (independent of the object). Therefore, participants can reach the a-priori preferred decision threshold faster, requiring less accumulated information compared to the other option.

Lastly, the *non-decision time* (τ) refers to the length of all response processes that are not part of the actual decision-making. Most typically, τ is interpreted as a motoric component of decision making (e.g. van Ravenzwaaij et al., 2012). Likewise—and somewhat reminiscent of the shooter bias—response execution may be faster for Black than White avatars holding a gun (as suggested empirically in a Weapon Identification Task by Payne et al., 2005). Although the parameter is theoretically not limited to motor execution, but also encompasses processes of encoding, the latter is typically related to basal stimulus perception (e.g., effect of interletter space on word recognition, Perea & Gomez, 2012). In contrast, broader top-down influenced encoding, such as perceptual word detection speed after a semantically congruent prime, is presumed to map on information accumulation processes (i.e., the drift rate, see Voss et al., 2013).

Connecting stereotypes and model parameters

Diffusion modeling allows for the testing of the previously described hypotheses about how stereotypes impact information processing which results in racially biased shooting behavior (see Table 1):

The *frequency hypothesis* assumes a skewed representation of gun occurrence, specifically that participants a-priori expect a gun more often in the hands of a Black than a White person. In the framework of diffusion modeling, this bias would be mirrored in the adjustment of the starting point ($\beta \cdot \alpha$) as preparation of the decision. If one expects an increased likelihood of Black avatars holding a gun, it makes sense to shift the starting point into the direction of the shooting decision. In that case, less information is needed to reach the shooting threshold and more information is required for an opposing no-shoot decision.

The *typicality hypothesis* states that the associations of a social group (Black/White) and the object (gun/phone) activate similar concepts and boost stereotype-consistent decisions. Diffusion modeling would map this kind of stereotype on the drift rate parameter (δ), enabling an accelerated information accumulation of objects when they are stereotypically consistent.

Table 1

Model parameters in the context of the current study

Model parameter	Stereotype hypothesis	Interpretation
Drift rate (δ)	Typicality	Accelerated processing of task-relevant objects when they are consistent with the associated stereotype of the avatar holding the object
Starting bias ($\beta \cdot \alpha$)	Frequency	A-priori expectation of stereotypical objects in the hands of the according avatar without considering task-relevant information
Non-decision time (τ)	Execution	Enhanced motoric execution (presumably due to stereotypical pre-activation), independent of the decision process

The *execution hypothesis* proposes that racial information acts as a precursor to a stereotype-driven response—independent of the decision process—that maps on the non-decision parameter (τ) of the diffusion model as an automatic and peripheral motoric component. It may be the result of a task-irrelevant activation specific for the threat association connected to a racial stimulus.

Previous studies trying to apply the drift diffusion model to the shooter task consistently found that the shooter bias is mirrored in the drift rate (Correll et al., 2015; Johnson et al., 2018; Pleskac et al., 2018). Speaking in favor of the typicality hypothesis, participants identified objects faster when the race-object combination was stereotype-compatible. One major caveat in interpreting these findings, however, is that all of them used the FPST stimuli by Correll et al. (2002). In this experimental setup, target race and object appear on screen simultaneously, leaving no time for effects of target race to unfold before the object appears. Specifically, diffusion modeling cannot independently represent the effect of race on the starting bias—which represents an a-priori tendency towards one of the options depending on the race of the target—without also processing the object. Arguably, this not only limits the potential of the diffusion model but also its ecological validity: in most cases, race will be discernible faster and from a greater distance than an object (that might be pulled out of a pocket over the course of an encounter). If race were to modulate the starting bias or enhances motoric preparation, race information must be presented before participants process any object information to gain meaningful results.

This is particularly relevant, as Johnson et al. (2018) showed that offering explicit prior information regarding race via a text overlay eliminated the shooter bias. This study used the standard FPST design but offered prior information about the race of the upcoming man in the picture on a written dispatch for 2 seconds. This textual prime may have unnecessarily highlighted the relevance of race and thus undermined the spontaneous and intuitive nature of the reaction. In another highly popular experimental paradigm to tap into racial biases, the weapon identification task, Black primes led participants to more readily detect a weapon than a tool by shifting the starting bias of the drift diffusion model (Todd et al., 2020).

Here, we introduce a modified shooter task in which virtual avatars appear on scene but pull an object out of their back pocket with a delay. This aligns the procedure to tasks in which race precedes the target in the form of a prime stimulus and allows the starting bias to vary by race, without making race explicitly salient in advance.

The present research

In three experimental studies (pre-registered; <https://osf.io/bg9kw/>), we replicated and extended the existing literature on the shooter bias using a video sequence of a virtual avatar instead of inanimate imagery (screenshot in Figure 2). The virtual setting allowed for full standardization of the shooter task, while mitigating any confounding effects from contrast ratio between skin color and the target or attire. The dynamic presentation of a video sequence showing the avatar before the object establishes an unequivocal chronology of inputs, allowing us to analyze the starting bias as a potential mechanism of how racial biases affect information processing based on a frequency bias about violence among Black males.

Figure 2. Screenshot of a video sequence used in Study 3. In this example, the experimental conditions were Black (skin color) and gun (object). The correct response as a reaction to the appearing gun is to press the shoot button.

Before exploring the underlying mechanisms of race (or skin color) on the decision process, the first study explored one of our starting assumptions, namely, that frequency stereotypes would map on the starting bias. Accordingly, we manipulated frequency stereotypes experimentally (without manipulating race) and analyzed their effect on the resulting diffusion model parameters. The second study is a full FPST, manipulating target race in a virtual environment while testing if the findings of the shooter bias can be replicated and dissecting the underlying cognitive processes using diffusion modeling. To validate the

results of the second study, we conducted a third study including a more diverse pool of avatars and twice as many trials (Study 3). In a file drawer statement, we also briefly report the results of two additional studies which were conducted in non-laboratory settings due to the closure of all institutional testing facilities during the COVID-19 pandemic; detailed analyses can be found in the online supplement.

The hypotheses of all studies are focused on the response times rather than error rates. The reason is that we chose comparably long response time windows where we expect response times to make for a better benchmark of the racial bias than error rates (Mekawi & Bresin, 2015).

The diffusion modeling results contain two different estimation procedures for all studies: a Bayesian estimation and an additional frequentist calculation based on the Maximum-Likelihood (ML)¹ estimation criterion. We report all measures, manipulations, and exclusions for all studies. Furthermore, we confirm that sample sizes for all studies were determined before conducting any data analysis.

Study 1: Pretesting the frequency hypothesis

In the first study, we aimed at manipulating the perceived frequency of dangerous target avatars to test if this would map on an adjusted starting point in the diffusion model accordingly. Therefore, we created associations of the avatar's shirt color congruent with the frequency of dangerous or harmless exemplars in advance of the shooter trials.

Concerning the reaction times, we expected an interaction between the threatening stereotype (indicated by shirt color) and the object type (Hypothesis 1), specifically, that the reaction times are lower for the gun object in combination with the color associated with threat compared to the color considered harmless and vice versa for the phone object. We

¹ The initially preregistered Kolmogorov-Smirnov criterion has been replaced by ML in the Review process.

further predict (Hypothesis 2) that (a) the starting point for the “dangerous” shirt color will experience a shift in the parameter $\beta \cdot \alpha$ to shoot more frequently because of the induced frequency bias; (b) the drift rate will be comparatively similar between the two shirt colors as the perceived threat is expected to map on the starting bias; (c) the threshold separation will be larger for “harmless” than for “dangerous” colors because participants will try to be more cautious.

Methods

Participants and design

The first study included 99 participants (mainly German psychology students; 89 females, 10 males; $M_{\text{age}} = 23.96$, $SD_{\text{age}} = 5.18$), receiving partial course credit for their participation. As preregistered, one participant was excluded from further analysis due to self-reported insufficient participation (final control question). The study builds on a 2×2 within-subjects design with object (gun or mobile phone) and threat stereotype of the target avatar (indicated by dark blue or white shirt color; the association of shirt color with dangerousness was counterbalanced between participants) as within-subject factors, having the presentation order fully randomized during experimental trials.

Procedure

As there were only within-subject factors (except for the theoretically uninteresting counterbalancing factors), the procedure was the same for all participants. After completing the declaration of consent, participants were seated 75 cm from the computer screen (22") in a laboratory room while completing a computer-based simulation of the shooter task constituting 15 training trials and 120 experimental trials in randomized order using Inquisit (version 4).

The instructions were presented on screen and asked participants to imagine being a police officer who is walking through a street. They were tasked to decide as quickly as

possible whether the virtual person appearing in the middle of the street is reaching for a gun or a mobile phone out of his back pocket. The instruction laid out to shoot if the person was reaching for a gun and to not shoot if the object was a phone. Participants logged their decision by pressing the left (“S”) or right (“L”) button on the keyboard. This key assignment was also randomized per participant at the beginning of the experiment and repeatedly displayed between trials as a reminder. To reinforce the stereotype of color, we inserted pictures into the instructions signifying the association between shirt color and threat. The assignment of shirt color (white or dark blue) to the according connotation (either dangerous or harmless) was also made at random. The threat association of shirt color was explained both visually and in written form as a higher base-rate of gun appearance ($\frac{2}{3}$) associated with this particular shirt color. In the harmless condition, we suggested that only $\frac{1}{3}$ of the population of avatars wearing this shirt color were gun-holders. However, the actual appearance of objects within the experiment was equally distributed as only this allowed us to dissect a-priori expectations (as the ones by stereotypes) from actual instrumental behavior.

Participants received direct written feedback after each trial if their response was correct or—in case of prolonged reaction times outside the restricted time interval of 850 ms—encouraging faster responses in future trials. We included no reward or payoff system. The inter-trial interval was as long as participants chose (with a minimum of 1000 ms) as the next trial only started after pressing the space bar. Participants were invited to have a short break after 60 trials. Finally, participants completed a series of demographic questions, were debriefed, and thanked.

Materials

The videos were created in Unity3D using an adaptable White standard avatar (MCS Male) and set in a foggy industrial street environment. The objects in hand were colored in neutral grey. The video sequence had a framerate of 30 FPS and a resolution of 1920×1080.

The scene consisted of a moving avatar grabbing behind his back and taking the relevant object out of his pocket. The movement lasted 1000 ms until the first part of the object was visible. The response time window started with the appearance of the object. Video editing was performed in Shotcut (Melttech LLC, 2018). There was one video clip per experimental condition (4 in total).

Behavioral analysis

We analyzed the behavioral data using *RStudio* and *Jamovi* according to the pre-registered analysis procedure. We excluded 1088 trials in which participants did not respond within the relevant time window between 200 and 850 ms. 11 Participants were excluded due to missing more than ¼ of all trials after exclusions. No remaining participant had an error rate above 40%. The final dataset was comprised of 87 participants. The analysis had 80% power to detect an effect size of $\eta_p^2 = .085$ at $\alpha = .05$ for the interaction effect of interest according to the power analysis in *MorePower* (Campbell & Thompson, 2012).

Diffusion modeling

The data basis for the hierarchical Bayesian diffusion model to test Hypotheses 2a/b/c was nearly identical to the one used in the behavioral analysis, but, additionally, included incorrect responses. The modeled decision process is noisy, resulting in trial-by-trial differences in the parameters. Although this parameter variability distribution is also part of the diffusion model, fixing these variability parameters typically yields good or even better results in many studies (Boehm, Annis, et al., 2018). We report the results of two different model estimators: a model based on hierarchical Bayesian estimation and a model based on the Maximum-Likelihood (ML) criterion.

Bayesian estimations. We modeled the data in R (version 4.1.0; R Core Team, 2021) using Markov-Chain-Monte-Carlo methods (MCMC) as implemented in JAGS (Version 4.3.0, Plummer, 2003) and the JAGS Wiener diffusion model extension (Wabersich &

Vandekerckhove, 2014). The model (see *Supplements*) was embedded in a hierarchical framework drawing parameters from distributions on individual and condition levels (Rouder et al., 2017). The Bayesian procedure estimates these posterior distributions of interest by updating prior distributions with the collected data. The priors express the beliefs about the parameter distributions without taking data into account (Kruschke, 2014). We used weakly informative priors on a sensible range of values for the hyper-prior distributions of the intercept parameters and uninformative priors for all other parameters. We allowed the drift rate and non-decision time to vary as a function of target threat (dangerous/harmless) and object (gun/phone), and the starting point and boundary separation to vary as a function of target threat only.

The hierarchical Bayesian model was embedded in a regression framework (Boehm, Steingroever, & Wagenmakers, 2018), where each of the cognitive parameters (boundary separation, drift rate, non-decision time in seconds, and starting point) was modeled by a regression model with effect coded predictors for the object (gun = 1, phone = -1) and threat stereotype conditions (dangerous = 1, harmless = -1). To test for the difference between conditions, we generated posterior distributions for all parameters per condition, based on the estimated posterior distributions of the parameters of an overarching regression model of each parameter. The reason for this approach was that—in the JAGS Wiener diffusion model extension—the two responses were coded by multiplying the reaction times for one response with -1 (here the shooting response). However, due to this artificial recoding needed to fit the model, the drift rates were negative in the gun (with predominant shooting responses) and positive in the phone condition (with predominant no-shooting responses). To adequately calculate the drift rate interaction effect between object and target threat, the parameter estimates were recoded into absolute values to avoid misinterpretations. This was not possible within the regression framework to initially fit the models and thus required the before

mentioned approach. In favor of a consistent report, we decided to use the same approach for the other parameters, exhibiting an equal pattern of results.

The results of the model evaluation all suggested a good model fit (see Estimations of the posterior distributions, Convergence of chains, Effective Sample Size, and Posterior Predictive Checks in *Supplements*). To quantify the credibility of the hypothesized parameters, we report the mean of the posterior distribution, the High-Density Interval (HDI, covering 95% of the parameter distribution) as well as Bayes factors approximated by the Savage-Dickey-Density Ratio on the point 0 (Wagenmakers et al., 2010) to test for differences between conditions. The resulting Bayes Factor indicates the relative evidence for H_1 (e.g., $\delta \neq 0$) compared to H_0 (e.g., $\delta = 0$, Kass & Raftery, 1995).

ML estimations. We performed modeling using *fast-dm* (version 30.2, Voss & Voss, 2007) allowing drift rate and nondecision time to vary as a function of target threat (dangerous/harmless) and object (gun/phone), and starting point and boundary separation to vary as a function of target threat only. As recommended by Lerche and Voss (2016), the intertrial variabilities of the starting bias and the drift rate, but not of the non-decision time were fixed to zero. The Monte Carlo simulation process, as proposed by Voss et al. (2015), evaluated the model fit for each participant by entering simulated datasets into the diffusion model and using the 5% quantile of the resulting fit indices as critical value for the empirical data. The procedure identified participants with bad model fit resulting in the exclusion of about 1% of the participants (for model fit and model specifications, see *Supplements*). Trying to reduce the number of estimated parameters to further increase reliability of estimations because of the low number of trials, resulted in a worse model fit. The final model parameters of each participant were entered into frequentist analyses (ANOVA) with skin color and object as within-factors. Similar to the Bayesian procedure, negative drift rates in the gun condition have been recoded into absolute values for the calculation of the interaction effect.

We report a summary and interpretation of the parameters in the following *Results* section in the light of the previously stated hypotheses as well as some exploratory findings.

Results

Response Times.

The analysis of reaction times showed that the interaction effect between threat (associated with shirt color) and object was significant, $F(1, 86) = 26.60, p < .001, \eta_p^2 = .236$, confirming Hypothesis 1. Response times were shorter when the object matched the association with threat, this is, when the phone was presented in the harmless condition and the gun was presented in the dangerous condition (see Figure 3). The factor *threat* had no main effect on the reaction times, $F(1, 86) = 0.01, p = .915, \eta_p^2 < .001$; the factor *object* was significant, showing prolonged response times for phones compared to guns, $F(1, 86) = 751.45, p < .001, \eta_p^2 = .897$.

Figure 3. Response times (in ms) of all four condition $\times$ object combinations (Study 1). Results illustrate that reaction times were on average higher for the phone than the gun object in the dangerous condition where participants expected a higher gun probability compared to the harmless trials and vice versa for the gun object. Error bars represent ± 1 standard error of the mean.

Error rates.

Participants decided less frequently to shoot in the harmless (46%) than in the dangerous condition (54%). The exploratory analysis of error rates (see Figure 4) shows an interaction effect of both factors, $F(1, 86) = 28.74, p < .001, \eta_p^2 = .251$, with less errors in the harmless compared to the dangerous condition with a phone in hand, $M_{\text{Diff}} = -0.06, SE = 0.01, t(86) = -4.80, p < .001, d_z = -0.52$, and more errors for the harmless condition with a presented gun, $M_{\text{Diff}} = 0.05, SE = 0.01, t(86) = 4.18, p < .001, d_z = 0.45$.

Figure 4. Error rates of all condition $\times$ object combinations (Study 1). The error rates mirror the same pattern of results as the response times. Error bars represent ± 1 standard error of the mean.

Diffusion modeling.

The results of Bayesian diffusion modeling (see Figure 5) mainly support Hypothesis 2: the starting bias differed significantly between the harmless and dangerous conditions (mean posterior difference: $-0.1[-0.13, -0.07]$, $BF_{10} > 1000$). Participants were more prone to shoot in the dangerous ($\beta = .43 [.41, .45]$) than in the harmless condition ($\beta = .53 [.51, .55]$) as shown by the β parameter being closer to the 0-decision boundary (shoot decision). As predicted, there was no significant difference in the drift rate between the conditions (neither a condition main effect $BF_{10} = 0.03$, nor a condition $\times$ object interaction effect $BF_{10} = 0.12$). The boundary separation was not larger in the harmless than in the dangerous condition (mean posterior difference: $0.05[-0.01, 0.12]$, $BF_{10} = 0.02$).

The ML estimation procedure, employing a frequentist analysis, replicated the results of the Bayesian modeling concerning Hypotheses 2a, with a difference between harmless and dangerous condition on the starting bias, $t(85) = 5.53, p < .001, d_z = 0.60$, but—contrary to Hypothesis 2b—yielded a small interaction effect in the drift rate, $F(1, 85) = 4.31, p = .041, \eta_p^2 = .048$, and—confirming Hypothesis 2c—a larger boundary separation in the ‘harmless’ compared to the ‘dangerous’ condition, $t(85) = 2.03, p = .046, d_z = 0.22$.

Figure 5. Parameter results of the Bayesian diffusion modeling procedure (Study 1). The overview of possible effects contributing to biased shooting behavior illustrates that the connection between the learned stereotype (dangerous/harmless) and the object (gun/phone) is mirrored in the starting bias of the Bayesian model (with 0.5 as center between the decision thresholds). Error bars represent the 95% HDI.

Discussion

In the present study, our data support the notion of a frequency bias, this is, that expectations of higher threat occurrences in a group of virtual avatars can influence shooting decisions and that this bias predominantly maps on the starting bias parameter of the Bayesian drift diffusion model (as well as in the ML estimation). As expected, reaction times

(and error rates) depended on the interaction between the artificially induced threat and the object in hand. Additionally, participants held an initial starting bias depending on the induced dangerousness of the avatar. For example, when an avatar associated with (induced) threat was presented, participants started their information accumulation closer to the shooting boundary and were, therefore, faster to respond with a shoot than a no-shoot decision. Previous studies modeled stimulus proportion manipulations as means of validating diffusion models (e.g., Arnold et al., 2015; Ratcliff & McKoon, 2008) and also found effects predominantly on the starting bias. Importantly, these studies differed from the current study in the fact that participants were not given wrong initial contingency information.

After establishing that stereotypes about the frequency of dangerous exemplars can affect response times and a starting bias, the following studies tested whether similar patterns also appear in traditional race-based shooter biases.

Study 2: Cognitive processes in the primed shooter task

In the second study, the general methodology remained the same as in Study 1, with skin color now being the critical variable instead of a learned contingency with shirt color. The procedure was analogous to the basic FPST (Correll et al., 2002), but with fully standardized target avatars and backgrounds to exclude possible interfering influences. As a replication of the shooter bias, we predicted an interaction effect between skin color and object in the reaction time analysis (Hypothesis 3).

The video sequence showed the target avatar in the first step and the object in a subsequent step without making race explicitly salient. Thereby, we enabled participants to process skin color spontaneously without the object information acting as a confounding prime and allowing the diffusion model to adjust the decision starting point. As preregistered, we expected the following pattern of modeled parameters (Hypothesis 4): (a) the starting point differs significantly by race of the target, with the starting bias being closer to the shoot

decision threshold for the Black than the White avatar (frequency bias: expecting a gun in the hands of Black rather than White targets); (b) steeper drift rates for race-object combinations, supporting the implicit typicality Black-danger stereotype (e.g., Black-gun); (c) the threshold separation will be larger for the Black avatar than for the White avatar as participants may try to avoid making racist decisions and try to be more careful in the case of Black avatars (see Pleskac et al., 2018). We did not derive hypotheses concerning the decisions (e.g. higher error rates for unarmed Black targets) because of mixed findings in previous studies for the given reaction time window (Mekawi & Bresin, 2015).

Methods

Participants and design

138 German psychology students (118 females, 20 males; $M_{age} = 23.59$, $SD_{age} = 4.87$) participated in the study and received partial course credit as compensation. All participants declared that they completed the task seriously. The study used a 2×2 within-subjects design with object (gun or mobile phone) and skin color of the target avatar (Black or White) as within-subject factors.

Materials

The avatar was identical in the Black and White condition, only varying in skin color while retaining all other physical features. Therefore, we prefer the term skin color to race. As the scene lighting interacts with all rendered objects, we calculated the average color of skin tone using five corresponding points for the Black and the White avatar, creating a neutral grey color for both objects (gun and mobile phone) with equivalent luminance contrast ratios (1.9) relative to both skin colors. The comparability between the Black and the White avatars has been confirmed in a small separate study in terms of their perceived direction of gaze and their perceived attentiveness (see <https://osf.io/bg9kw/>).

The video sequence had the same background and settings as did the first study. Likewise, it consisted of a male avatar grabbing behind his back to take an object out of his pocket. The movement lasted 1100 ms until the first part of the object was visible, which was 100 ms longer than in the first study. The response time window started with the appearance of the object.

Behavioral analysis

Data analysis was the same as in the first experiment. We excluded 293 trials in which participants did not respond within the relevant time window, resulting in a total exclusion of one participant (with more than ¼ of excluded trials). Another 986 trials were excluded due to incorrect responses; no participant had an error rate above 40%. The final number of datasets was 137 for the subsequent behavioral analysis. The analysis had 80% power to detect an effect size of $\eta_p^2 = .055$ at $\alpha = .05$ for the interaction effect of interest.

Diffusion modeling

The general procedure of diffusion modeling to test Hypotheses 4a/b/c was conducted analogously to Study 1, with the independent variable now being skin color (rather than learned threat). We allowed the drift rate and non-decision time to vary as a function of skin color (Black/White) in addition to object (gun/phone), and the starting point and boundary separation to vary as a function of skin color only. Again, we modeled the data with a Bayesian and a ML-based estimation method. The evaluation of the model and model fit showed satisfactory results for both methods (see *Supplements*). The data of seven participants did not converge in the ML-based procedure because of too few data points per condition. No participants had to be excluded in the ML estimation due to bad model fit according to the MC simulation.

Results

Response Times.

Reaction times (see Figure 6) were faster for the gun than the phone condition, and the difference was on average more pronounced for the Black (143 ms) than the White (74 ms) avatar. This interaction effect between skin color and object was significant, resulting in large effect sizes (see Table 2), supporting Hypothesis 3 and replicating the shooter bias effect. Exploratory analyses show that the shooter bias was constant over the duration of the experiment, with similar average differences between the conditions in the first and in the second half (<https://osf.io/bg9kw/>). Additionally, we controlled for the random button assignment but found no systematic effect of the assigned side, $F(1, 135) = 0.46, p = .498, \eta_p^2 = .003$.

Figure 6. Response times (in ms) of all four condition $\times$ object combinations (Study 2). Results illustrate that reaction times were on average higher for the phone than the gun object when the avatar was Black compared to the White avatars and vice versa for the gun object. Error bars represents ± 1 standard error of the mean.

Table 2.

ANOVA results of reaction time analysis (Study 2). The object (gun/phone) and the interaction effect of phone and skin color (Black/White) contribute with high effect size to the reaction times.

Factor	<i>df</i>	<i>F</i>	<i>p</i>	η_p^2
Skin-color	1	0.08	.774	.001
Object	1	980.19	< .001	.878
Skin-color × Object	1	592.11	< .001	.813
Residuals	136			

Error rates.

The exploratory analysis of error rates (Figure 7) revealed a significant interaction effect skin color × object on the error rates, $F(1, 136) = 4.99$, $p = .027$, $\eta_p^2 = .035$. Error rates appeared to be generally higher for the White avatar but the difference was larger in the gun condition, $M_{\text{Diff}} = 0.04$, $SE = 0.01$, $t(136) = 5.31$, $p < .001$, $d_z = 0.45$, than the phone condition, $M_{\text{Diff}} = 0.02$, $SE = 0.01$, $t(136) = 3.58$, $p < .001$, $d_z = 0.31$.

Figure 7. Error rates of all condition × object combinations (Study 2). The error rates mirror a similar stereotypical pattern as the response times. Error bars represent ± 1 standard error of the mean.

Diffusion modeling.

At odds with Hypothesis 4a, skin color had no influence on the starting bias (mean posterior difference: $-0.02[-0.04, 0.00]$, $BF_{10} = 0.01$) of the hierarchical Bayesian diffusion model (Figure 8). This suggests that participants had no a-priori starting bias towards one of the decisions as a function of skin color. The starting bias was slightly shifted into the

direction to shoot for the Black *and* the White avatar, indicating a greater general readiness to shoot (vs. not shoot).

There was also no support for Hypothesis 4b as the interaction between skin color and object on the drift rate was not significant (mean posterior difference: $-0.14[-0.40, 0.13]$, $BF_{10} = 0.03$). Drift rates were generally faster for the Black than the White avatar (mean posterior difference: $0.65[0.50, 0.79]$, $BF_{10} > 1000$) in both object conditions. The study provides evidence in favor of Hypothesis 4c that skin color has a significant effect on the boundary separation α (mean posterior difference: $0.14[0.10, 0.19]$, $BF_{10} > 1000$).

Participants accumulated more information when the avatar was Black.

Figure 8. Parameter results of the Bayesian diffusion modeling procedure (Study 2). The overview of possible effects contributing to the biased shooting behavior illustrates that the connection between the skin color (Black/White) and the object (gun/phone) is mirrored in the drift rate and non-decision time of both models. Error bars represent the 95% HDI.

Further analyses suggest that the non-decision time parameter was especially relevant (mean posterior difference of the interaction effect: 0.06[0.04, 0.07], $BF_{10} > 1000$) as it mirrored the Black-violence stereotype with shorter times for the gun and longer times for the phone condition when comparing the Black to the White avatar.

The ML-based estimations confirmed the Bayesian results: no significant effect of skin color on the starting bias (Hypothesis 4a, although descriptively in the expected direction), $t(129) = 1.81, p = .072, d_z = 0.16$, but significant differences for the skin color on the boundary separation (Hypothesis 4c), $t(129) = 2.71, p = .008, d_z = 0.24$, as well as a significant interaction effect on the non-decision time (exploratory analysis), $F(1, 129) = 40.21, p < .001, \eta_p^2 = .238$, but not on the drift rate (Hypothesis 4b), $F(1, 129) = 0.53, p = .468, \eta_p^2 = .004$.

Discussion

In the second study, we successfully replicated the FPST under fully standardized conditions. Reaction times were prolonged when skin color and object conditions were counter-stereotypically combined and faster when they were stereotype-consistent. The stereotype-consistent pattern was also reflected in the error rates.

Furthermore, we aimed to investigate if the stereotypical expectation stigmatizing Black males as being armed and dangerous more often than their White counterparts (frequency bias) can serve as an explanation for the primed shooter bias, corresponding to the results of Study 1. Despite highly similar results for response times, however, diffusion modeling suggested different underlying processes behind these effects. The a-priori frequency assumption (starting bias) could account for the shooter bias.

The parameter of the diffusion model which mirrored the shooter interaction effect of the behavioral analysis was the non-decision time. This result implicates that participants were motorically better prepared to shoot on Black than on White targets, and vice versa for

the phone condition. If both sources of information compete with each other, the response execution is delayed. Such a motoric preparation has been associated with sequential priming effects (Voss et al., 2013) and might thus reflect an executive readiness, facilitating responses independent of information accumulation. As we did not pre-register the effect on this parameter, however, we sought to replicate this pattern in Study 3.

We found no evidence in favor of an accelerated information accumulation (i.e., drift rate) due to stereotype-object congruency. Participants were not faster in detecting handheld objects of a target person when they are stereotype-conform compared to objects that contradict the stereotype, consistent with a spreading-activation account of associative priming enhancing access to related activated concepts.

The results also suggest that processing was generally more efficient when the avatar was Black—independent of the object at hand—as information accumulation was faster, and participants were more careful, gathering more information before responding. The larger boundary separation for the Black avatar may reflect the motivation to control prejudice (Glaser & Knowles, 2008).

Study 3: Replication study

The third study aimed to replicate and validate the results of Study 2 using a different virtual setting and a more diverse pool of avatars as to increase generalizability. Additionally, we used twice as many trials as in the previous studies to consolidate the diffusion modeling procedure and prolonged stimuli presentation windows. We predicted an interaction effect between race and object in the reaction times (Hypothesis 5) and the following pattern of results concerning cognitive processes: an interaction effect on the non-decision parameter in terms of faster responses for Black targets with a gun and White targets with a phone compared to the other color-object combinations (Hypothesis 6a), which might reflect an execution bias hypothesis for preparing stereotype-consistent actions. In line with Study 2,

we expected more efficient processing for trials supporting the skin-color-violence stereotype resulting in steeper drift rates for Black than for White avatars (Hypothesis 6b, mistakenly described as main effect in the preregistration) as well as a generally larger threshold distance for Black avatars (Hypothesis 6c).

Methods

Participants and Design

111 German students (74 females, 37 males; $M_{age} = 22.04$, $SD_{age} = 3.00$) participated in the study and were compensated with 9 Euro per hour. The study used a 2×2 within-subjects design with object (gun or mobile phone) and skin color of the target avatar (Black or White) as within-subject factors.

Procedure

As this study was a replication of Study 2, the general methodology and materials were identical, with an increased number of trials (15 training trials and 240 experimental trials). We needed to make slight deviations from the pre-registered procedure due to technical problems in the lab: we could not implement the additional questionnaires and additional exclusion criteria, nor could we randomize the key assignment for the responses.

Materials

The video clips were like those in Study 2. This time we used a metro station as a background environment, rendered in a consistent shade of dark slate blue, to further focus participants on the avatar. We further varied more physiognomic parameters, resulting in more diverse avatars (see Figure 2). Specifically, in addition to the skin color (Black or White) and the object in hand (gun or phone), we varied the color of the shirt (bright or dark), as well as the physique of the target (stout or athletic), resulting in 16 different video clips, each with a duration of 1500 ms until the object was visible.

Data analysis

The method of data analysis was the identical to that used in both previous experiments. We excluded 750 trials in which participants did not respond within the appropriate time frame, resulting in a total exclusion of four participants (with more than ¼ of excluded trials). A total of 1300 trials was excluded due to wrong answers; no participant had an error rate above 40%. The final number of datasets was 107 for subsequent behavioral analysis. The analysis had 80% power to detect an effect size of $\eta_p^2 = .070$ at $\alpha = .05$ for the interaction effect of interest.

The diffusion model procedure to explore Hypotheses 6a/b/c was substantially the same as in the previous study, similarly incorporating incorrect trials into the model. Evaluation of model fit of both diffusion models showed satisfactory results (see *Supplements*). The data of three participants (3%) were excluded from further analysis of the estimated parameters based on poor ML model fit criteria following the MC model fit simulations.

Results

Response Times.

As predicted in Hypothesis 5, the interaction between skin color and object was significant (see Table 3) replicating the shooter bias effect in the reaction times. Stereotype-consistent skin color-object combinations accelerated the response times (see Figure 9), albeit with a noticeably smaller effect size than in Study 2 (raw descriptives exhibit a mean difference of 7 ms for Black vs. White avatars, compared to 69 ms in Study 2). The race differences per condition between the first and the second half were not significant (see <https://osf.io/bg9kw/>).

Table 3. ANOVA results of reaction time analysis (Study 3). The object (gun/phone) and the interaction between object and skin color had a significant effect on the reaction times.

Factor	<i>df</i>	<i>F</i>	<i>p</i>	η_p^2
Skin-color	1	3.21	.076	.029
Object	1	322.44	< .001	.753
Skin-color $\times$ Object	1	11.25	.001	.096
Residuals	106

Figure 9. Response times (in ms) of all four condition $\times$ object combinations (Study 3). Results illustrate that reaction times were on average higher for the phone than the gun object when the avatar was Black compared to the White avatars and vice versa for the gun object. Error bars represent ± 1 standard error of the mean.

Error rates.

We further explored whether the reduced effect size on response times was partially due to a shift of meaningful variance to error rates. Exploratory analyses of the error rates (Figure 10) demonstrate no interaction effect between skin color and object, $F(1, 106) = 1.79$, $p = .184$, $\eta_p^2 = .017$. The Black-violence stereotype was not significantly reflected in the error rates, with error rates generally being lower compared to the first and second study.

Figure 10. Error rates of all condition $\times$ object combinations (Study 2). The error rates do not mirror a similar stereotypical pattern as the response times. Error bars represent ± 1 standard error of the mean.

Diffusion modeling.

The analysis of parameters resulting from the Bayesian estimations (Figure 11) confirmed none of the Hypotheses 6a, 6b, nor 6c: There was neither a significant interaction effect of skin color and object for drift rate, mean posterior difference: $-0.17[-0.37, 0.05]$, $BF_{10} = 0.05$, nor for non-decision time, mean posterior difference: $0.02[0.01, 0.03]$, $BF_{10} = 0.47$. No main effect emerged for the boundary separation (mean posterior difference: $-0.01[-0.06, 0.05]$, $BF_{10} < 0.01$) and, once again, there was no evidence for the involvement of a starting bias (pattern of results is descriptively even contrary to the hypothesis, mean posterior difference: $0.03[0.01, 0.04]$, $BF_{10} = 0.21$). The ML modeling

approach also failed to unravel any differences in the cognitive processes underpinning the shooter bias in a significant way (see <https://osf.io/bg9kw/>).

Figure 11. Parameter results of the Bayesian diffusion modeling (Study 3). The overview of possible effects contributing to the biased shooting behavior illustrates that the shooter bias is not mirrored in any of the modeled parameters. Error bars represent the 95% HDI.

Discussion

The purpose of the present study was to validate the findings of the second study with a more diverse avatar pool and sample. We replicated the primed shooter bias in the response times with a noticeably smaller effect size than in Study 2. However, the diffusion model did not mirror this small effect on a cognitive level. Smaller videos on less powerful hardware, more unsystematic variance due to a more varied avatar set but especially the prolonged interstimulus interval between prime and target object likely limited the discovery of the cognitive effects of interest. Training effects, due to the increased number of trials, did not

reduce the effect in response times in the second compared to the first half. Nevertheless, the findings of Study 3 do not contradict the first two studies as no other competing parameter (i.e., the starting bias) was significant.

File drawer statement

After these studies, we conducted two more studies (studies S_1 and S_2) to further bolster the generalizability of the observed findings. As the shutdown of lab space due to the COVID-19 pandemic prohibited all in-person testing, we ran subsequent experiments online using participants recruited from Prolific academic. As expected, the greater variance in devices, display sizes, and settings compounded by uncontrollable distractions, resulting from unstandardized environments, likely created overall noisier data. A combined analysis of these two studies ($N = 426$) yielded the predicted interaction effect on response times but attenuated by an order of magnitude (raw effect size of the interaction being 3ms). Diffusion modeling accordingly did not map this small effect to any of the stated hypotheses. In addition to the already presented analysis, further detailed analyses are available in the *Supplements* to avoid file drawer issues and highlight the difficulties of obtaining reliable and meaningful response time effects in this paradigm in noisier online data.

General discussion

The objectives of the three studies (and two supplemental studies) were to replicate the shooter bias under standardized virtual conditions as well as to further decompose the underlying cognitive processes by applying diffusion modeling. Our studies on German samples support previous findings that the response-time pattern mirrors a racial bias connected to a Black-threat stereotype. The evidence of this racial bias in the motoric execution behavior is important for the rising public discussions about racial discrimination, e.g., in Germany where racial stereotypes have traditionally been treated as problems of countries with larger Black communities.

Furthermore, we found that the primed shooter bias mapped onto the non-decision time of the drift diffusion model supporting the role of motoric preparations (when the behavioral shooter bias effect was strong in the reaction times). This result is a pointer to the relevance of the *execution preparedness hypothesis* and, thus, that the skin color prime in the video sequence pre-activates a stereotype-consistent motor response program (Payne et al., 2005). The reason for this lies in congruency effects when prime and target are stereotypically related (Voss et al., 2013), i.e. the mechanism is a competition between responses, either facilitating the response when prepared response and correct response are congruent or interfering with the response execution if they are incongruent and competing. For example, Black avatars more often pre-activated a shooting response than White avatars resulting in faster responses in the gun than the phone condition. In a nutshell, skin-color primes can influence the motoric execution cascade.

Ruling out alternative hypotheses

Like the execution hypothesis, the frequency hypothesis assumes a response preparation as well, but on a decisional rather than a motoric level. In Study 1, we demonstrated how the frequency bias was manifested in behavioral and modeling results, yielding the expected effect on the starting bias of the diffusion model. Modeling the results of the second study, however, suggested different processes implied in spontaneous reactions to skin color in the racial shooter task. Expecting a gun a-priori more often in the hands of a Black avatar did not receive support as the ruling mechanism behind the shooter bias in our studies—although the dynamic video sequences with a sequential presentation of race and object information also allowed the starting bias to vary by skin color in advance of the decision. Liberating constraints of simultaneous processing by presenting the information sequentially suggests that race information is sometimes used to prepare a response execution but not the decision process.

This points to an interesting difference to another well-established experimental paradigm: the weapon identification task. Here, participants must identify an object as either a weapon or a tool as quickly as possible after being primed with either a Black or a White face. In a recent series of studies (Todd et al., 2020), the observed bias of an increased tendency to identify a weapon after Black faces consistently mapped on the starting bias. The crucial difference may lie in the experimental design of the studies relying on a typical priming task with explicit primes and raw object pictures without context scenes. The typical priming task possibly leads to a different processing style by *separately* presenting the race and object information sequentially. The videos in our studies provided race information in advance and then race and object *simultaneously*. Embedding cues in an information flow might limit the conscious processing of the skin color in advance of the main stimulus. A sterile study design without context reduced to the essential elements enhances the extraction of effects with high effect size but may undermine ecological validity. To our understanding, contextualized video shooting scenes can come closer to real-world shooting scenarios as they avoid focusing explicitly on target race and contain the sequential as well as parallel presentation of relevant information.

Previous studies that did not sequentially present race and object information (Correll et al., 2015; Johnson et al., 2018; Pleskac et al., 2018) point at the relevance of stereotype-consistent information processing facilitating the visual information uptake, i.e., that participants recognize a gun faster in the hands of Black than White avatars. In a simultaneous presentation of race and object, stereotypes can guide visual perception and object detection by activating overlapping stereotypical networks and thereby accelerating the object identification speed for concordant (e.g., Black – gun) compared to interfering stimuli. However, the information accumulation parameter (drift rate) was not significant in our studies speaking against this typicality hypothesis. By offering the skin color information in

advance, we allowed the diffusion model to additionally map the effect on parameters that are related to preparatory processes (i.e., starting bias and non-decision time). This experimental design might come closer to real world situations, but our implementation does not come without caveats.

Limitations

There are some limitations to this experimental approach. While we controlled for the contrast ratio between skin color and object, we could not control for the contrast between background and skin color. Furthermore, we did not use a multitude of avatars due to economic reasons and varied features of only one basic avatar. Moreover, objects always appeared at the same point on screen and a variation of the position of the focal point might have been experimentally cleaner.

The interpretation of the cognitive processes stands and falls with the validity of the diffusion model—in general and for this specific task. Interpretation of the parameters in terms of psychological processes may have some limitations and could also theoretically be explained by alternative processes. We consider the diffusion model as appropriate for this task due to sufficient model fit (see *Supplements*) and rely on previous research demonstrating the validity of the parameters (e.g., Johnson et al., 2018; Voss et al., 2004).

Evidence for the role of the non-decision time parameter might have been stronger in case of a larger bias effect size in Study 3, enabling the model to replicate the parameter findings as well. As it seems, the drastically reduced behavioral effect on response times left less systematic variance to be explained by diffusion model parameters.

Why exactly was the size of the racial bias in the reaction times small in Experiment 3 and especially in the additional online studies S₁ and S₂? Besides the previously named potential problems (i.e., differing, slower or smaller technical devices; online settings; prolonged interstimulus interval), the video sequences as stimuli may generally have made it

easier to ignore the skin color information compared to when they are presented suddenly together with the object (classical FPST) or when it is presented alone without context in advance (typical priming tasks). Maybe the skin color information was neither attention-grabbing in the first place among the many visual cues in the scene and lacking strong stereotypically driven preparations, nor did it become salient because the only new upcoming stimulus in the video sequence was the object. That said, it should also be noted that we only varied skin tone in our avatars and kept other features typically associated with race constant. Arguably, our experiments thus constitute comparably conservative tests of the effect of race, as experimental control was maximized by varying a single feature (albeit one that already children use as the strongest predictor of race; Dunham et al., 2015). Another explanation is based on the finding that the main effect of race was more pronounced in experiments where the interaction effect was small (with faster reaction times for Black avatars, but without making more errors). Maybe participants who were aware of the aim of these face-valid studies tried not to show a racial bias by paying more attention to the Black avatars relocating the stereotypical effects statistically on the factor of race.

Another limitation lies in the artificiality of the study. Real shooting situations may elicit different processes than laboratory studies. Although the unselected online samples as well as the student samples in Studies 1 to 3 may speak to the generalizability of the effect, these populations rarely come into situations where they are required to make shooting decisions. Furthermore, the lab studies were conducted only with German samples and some of the investigated processes might be specific for the country or culture.

Implications

The experimental game thus provides robust evidence for a common people's seemingly automatic behavioral bias due to (unintended) racial stereotypes impacting the motoric cascade. Despite this significant implication, it is not fully clear if the artificial FPST

paradigm is a valid approach to tackle real-world shooting situations (Cox et al., 2014) and the literature provides inconsistent findings concerning shooter bias effects among trained police officers (Cox & Devine, 2016).

However, the studies offer important insights into stereotype processes on the automatic cognitive level that are characterized by the shooter bias term: we assume that the bias is implicit and unintended due to the time pressure of the task. The implicit character may be the reason why the bias was unrelated to explicit prejudices (Glaser & Knowles, 2008). The stereotypes are not pre-decisional and pre-judicial but stereotypical pictures that pre-activate a stereotype-concordant motoric execution cascade. This finding highlights that the bias is mainly valid for fast first impressions and could be dominated or overridden by more elaborative processes and viewing times. Consequently, it may have only little or no problematic implications if a second, more elaborative system (see Frankish, 2010) corrects the first impressions of situations without acute time constraints. Training seems to reduce the bias as well: in general, studies suggest that practice effects reduce automatic racial biases (Calanchini et al., 2013; Plant et al., 2005) and, specifically for the shooter task, training appears to attenuate the stereotypic bias (Sim et al., 2013). Training could presumably affect the motoric execution control that we identified as relevant for the bias in the given studies. Nevertheless, training did not attenuate the effect in our prolonged Study 3 comparing the first and the second half of trials.

Future research

Among other questions, future research could investigate individual and intercultural differences of the bias and what constitutes implicit stereotypes, how they develop, and how their influence on first impressions can be reduced. Important for research on this paradigm is the awareness of the small effect size in units of absolute reaction time differences. As the supplementary studies 4 and 5 have demonstrated, the paradigm requires clean lab conditions

and accurate, comparable technical settings. Video sequences may be replaced by a sequence of screenshots producing less noisy data than possible video play delays. The position of the focal point should be varied, for instance, by also implementing visually mirrored stimuli. As realism and immersion are always reduced in standard questionnaire-type experiments, future studies could render the stimulus on a head-mounted display, possibly further improving external validity. Additionally, the response-time window and the interstimulus interval between prime and target have to be relatively short. Furthermore, the mean reaction time differences should not be treated as interindividually interpretable scales. The reliability of differences between Black and White avatars seems to be problematic, e.g., Study 3 yielded weak reliability coefficients of $\alpha = .26$ (gun) and $\alpha = .47$ (phone).

Conclusion

Our studies provide evidence of how unintended stereotyping of Black and White skin color primes can influence simulated (shooting) behavior. Although we cannot determine the impact of racial biases in real police shootings, the approach offers important insights into the social cognition of unintended, automatic stereotyping in untrained German samples. Demonstrating the existence of these implicit typicality stereotypes can contribute to sharpen the awareness for stereotypes—but should not be dramatized because of small absolute effect sizes. Training and less time pressure may reduce early racial biases in favor of more deliberative decisions.

Acknowledgments

We would like to thank Nina Winands, Jennifer Weyell, Jonas Knäble, and Felix Zimmer for assisting in data collection.

Open Practices

All relevant data, preregistrations, and stimuli as well as model fit calculations are accessible in an open archive (<https://osf.io/bg9kw/>). Model codes, full reports related to the

file drawer, Posterior Predictive Checks, Posterior Parameters, and Bayes Factors are available in the online *Supplement* of the paper.

References

- Arnold, N. R., Bröder, A., & Bayen, U. J. (2015). Empirical validation of the diffusion model for recognition memory and a comparison of parameter-estimation methods. *Psychological Research*, 79(5), 882–898. <https://doi.org/10.1007/s00426-014-0608-y>
- Boehm, U., Annis, J., Frank, M. J., Hawkins, G. E., Heathcote, A., Kellen, D., Kryptos, A.-M., Lerche, V., Logan, G. D., Palmeri, T. J., van Ravenzwaaij, D., Servant, M., Singmann, H., Starns, J. J., Voss, A., Wiecki, T. V., Matzke, D., & Wagenmakers, E.-J. (2018). Estimating across-trial variability parameters of the diffusion decision model: Expert advice and recommendations. *Journal of Mathematical Psychology*, 87, 46–75. <https://doi.org/10.1016/j.jmp.2018.09.004>
- Boehm, U., Steingroever, H., & Wagenmakers, E.-J. (2018). Using bayesian regression to test hypotheses about relationships between parameters and covariates in cognitive models. *Behavior Research Methods*, 50(3), 1248–1269. <https://doi.org/10.3758/s13428-017-0940-4>
- Calanchini, J., Gonsalkorale, K., Sherman, J. W., & Klauer, K. C. (2013). Counter-prejudicial training reduces activation of biased associations and enhances response monitoring. *European Journal of Social Psychology*, 43(5), 321–325. <https://doi.org/10.1002/ejsp.1941>
- Campbell, J. I. D., & Thompson, V. A. (2012). Morepower 6.0 for anova with relational confidence intervals and bayesian analysis. *Behavior Research Methods*, 44(4), 1255–1265. <https://doi.org/10.3758/s13428-012-0186-0>
- 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., Urland, G. R., & Ito, T. A. (2006). Event-related potentials and the decision to shoot: The role of threat perception and cognitive control. *Journal of Experimental Social Psychology*, 42(1), 120–128. <https://doi.org/10.1016/j.jesp.2005.02.006>
- 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>.
- Cox, W. T. L., & Devine, P. G. (2016). Experimental research on shooter bias: Ready (or relevant) for application in the courtroom? *Journal of Applied Research in Memory and Cognition*, 5(3), 236–238. <https://doi.org/10.1016/j.jarmac.2016.07.006>
- Cox, W. T. L., Devine, P. G., Plant, E. A., & Schwartz, L. L. (2014). Toward a comprehensive understanding of officers' shooting decisions: No simple answers to this complex problem. *Basic and Applied Social Psychology*, 36(4), 356–364. <https://doi.org/10.1080/01973533.2014.923312>
- Dunham, Y., Stepanova, E. V., Dotsch, R., & Todorov, A. (2015). The development of race-based perceptual categorization: Skin color dominates early category judgments. *Developmental Science*, 18(3), 469–483. <https://doi.org/10.1111/desc.12228>
- Eberhardt, J. L., Goff, P. A., Purdie, V. J., & Davies, P. G. (2004). Seeing black: Race, crime, and visual processing. *Journal of Personality and Social Psychology*, 87(6), 876–893. <https://doi.org/10.1037/0022-3514.87.6.876>
- Essien, I., Stelter, M., Kalbe, F., Koehler, A., Mangels, J., & Meli, S. (2017). The shooter bias: Replicating the classic effect and introducing a novel paradigm. *Journal of Experimental Social Psychology*, 70, 41–47. <https://doi.org/10.1016/j.jesp.2016.12.009>

- Fleming, K. K., Bandy, C. L., & Kimble, M. O. (2010). Decisions to shoot in a weapon identification task: The influence of cultural stereotypes and perceived threat on false positive errors. *Social Neuroscience*, 5(2), 201–220.
<https://doi.org/10.1080/17470910903268931>.
- Frankish, K. (2010). Dual-process and dual-system theories of reasoning. *Philosophy Compass*, 5(10), 914–926. <https://doi.org/10.1111/j.1747-9991.2010.00330.x>
- 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>
- UN Human Rights. (2017). *Germany racial profiling: UN experts highlight situation of people of african descent*.
<https://www.ohchr.org/en/NewsEvents/Pages/DisplayNews.aspx?NewsID=21239&LangID=E>
- Johnson, D. J., Cesario, J., & Pleskac, T. J. (2018). How prior information and police experience impact decisions to shoot. *Journal of Personality and Social Psychology*, 115(4), 601–623. <https://doi.org/10.1037/pspa0000130>
- Johnson, D. J., & Chopik, W. J. (2019). Geographic variation in the black-violence stereotype. *Social Psychological and Personality Science*, 10(3), 287–294.
<https://doi.org/10.1177/1948550617753522>
- Kahn, K. B., & Davies, P. G. (2017). What influences shooter bias? The effects of suspect race, neighborhood, and clothing on decisions to shoot. *Journal of Social Issues*, 73(4), 723–743. <https://doi.org/10.1111/josi.12245>
- Kahn, K. B., Goff, P. A., Lee, J. K., & Motamed, D. (2016). Protecting whiteness: White phenotypic racial stereotypicality reduces police use of force. *Social Psychological and Personality Science*, 7(5), 403–411. <https://doi.org/10.1177/1948550616633505>

- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795. <https://doi.org/10.1080/01621459.1995.10476572>
- Kenworthy, J. B., Barden, M. A., Diamond, S., & del Carmen, A. (2011). Ingroup identification as a moderator of racial bias in a shoot–no shoot decision task. *Group Processes & Intergroup Relations*, 14(3), 311–318. <https://doi.org/10.1177/1368430210392932>
- Knight, B. (2020). *Germany struggles to face its own police racism*. Deutsche Welle. <https://p.dw.com/p/3dIgi>
- Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2. ed.). AP Academic Press/Elsevier.
- Lerche, V., & Voss, A. (2016). Model complexity in diffusion modeling: Benefits of making the model more parsimonious. *Frontiers in Psychology*, 7, 1324. <https://doi.org/10.3389/fpsyg.2016.01324>
- Lundberg, G. J. W., Neel, R., Lassetter, B., & Todd, A. R. (2018). Racial bias in implicit danger associations generalizes to older male targets. *PloS One*, 13(6), e0197398. <https://doi.org/10.1371/journal.pone.0197398>
- Mange, J., Chun, W. Y., Sharvit, K., & Belanger, J. J. (2012). Thinking about arabs and muslims makes americans shoot faster: Effects of category accessibility on aggressive responses in a shooter paradigm. *European Journal of Social Psychology*, 42(5), 552–556. <https://doi.org/10.1002/ejsp.1883>
- Mange, J., Sharvit, K., Margas, N., & Sénémeaud, C. (2016). Do I shoot faster because I am thinking about an outgroup or a threatening outgroup? *Social Psychology*, 47(1), 29–37. <https://doi.org/10.1027/1864-9335/a000255>

- 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>
- Mekawi, Y., Bresin, K., & Hunter, C. D. (2019). Dehumanization of african-americans influences racial shooter biases. *Race and Social Problems*, 11(4), 299–307. <https://doi.org/10.1007/s12552-019-09267-y>
- Meltytech LLC. (2018). *Shotcut* [Computer software]. <https://shotcut.org/>
- Miller, S. L., Zielaskowski, K., & Plant, E. A. (2012). The basis of shooter biases: Beyond cultural stereotypes. *Personality & Social Psychology Bulletin*, 38(10), 1358–1366. <https://doi.org/10.1177/0146167212450516>
- 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>
- Perea, M., & Gomez, P. (2012). Increasing interletter spacing facilitates encoding of words. *Psychonomic Bulletin & Review*, 19(2), 332–338. <https://doi.org/10.3758/s13423-011-0214-6>
- Pickel, K. L., & Sneyd, D. E. (2018). The weapon focus effect is weaker with black versus white male perpetrators. *Memory*, 26(1), 29–41. <https://doi.org/10.1080/09658211.2017.1317814>
- Plant, E. A., Goplen, J., & Kunstman, J. W. (2011). Selective responses to threat: The roles of race and gender in decisions to shoot. *Personality & Social Psychology Bulletin*, 37(9), 1274–1281. <https://doi.org/10.1177/0146167211408617>

- 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>
- Pleskac, T. J., Cesario, J., & Johnson, D. J. (2018). How race affects evidence accumulation during the decision to shoot. *Psychonomic Bulletin & Review*, 25(4), 1301–1330. <https://doi.org/10.3758/s13423-017-1369-6>
- Plummer, M. (2003). Jags: A program for analysis of bayesian graphical models using gibbs sampling. *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, 125.
- R Core Team. (2021). *R: A language and environment for statistical computing* [Computer software].
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4), 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, 20(4), 260–281. <https://doi.org/10.1016/j.tics.2016.01.007>
- Rouder, J. N., Morey, R. D., & Pratte, M. S. (2017). Bayesian hierarchical models of cognition. In W. H. Batchelder, H. Colonius, E. N. Dzhafarov, & J. Myung (Eds.), *New handbook of mathematical psychology* (pp. 504–551). Cambridge University Press. <https://doi.org/10.1017/9781139245913.010>
- Sadler, M. S., Correll, J., Park, B., & Judd, C. M. (2012). The world is not black and white: Racial bias in the decision to shoot in a multiethnic context. *Journal of Social Issues*, 68(2), 286–313. <https://doi.org/10.1111/j.1540-4560.2012.01749.x>

Safronova, V. (2020). *Black germans say it's time to look inward*. The New York Times.

<https://www.nytimes.com/2020/10/04/style/black-germans-say-its-time-to-look-inward.html>

Sim, J. J., Correll, J., & Sadler, M. S. (2013). Understanding police and expert performance:

When training attenuates (vs. Exacerbates) stereotypic bias in the decision to shoot.

Personality & Social Psychology Bulletin, 39(3), 291–304.

<https://doi.org/10.1177/0146167212473157>

Sommers, S. R., Apfelbaum, E. P., Dukes, K. N., Toosi, N., & Wang, E. J. (2006). Race and

media coverage of hurricane katrina: Analysis, implications, and future research questions.

Analyses of Social Issues and Public Policy, 6(1), 39–55. [https://doi.org/10.1111/j.1530-](https://doi.org/10.1111/j.1530-2415.2006.00103.x)

[2415.2006.00103.x](https://doi.org/10.1111/j.1530-2415.2006.00103.x)

Todd, A. R., Johnson, D. J., Lassetter, B., Neel, R., Simpson, A. J., & Cesario, J. (2020).

Category salience and racial bias in weapon identification: A diffusion modeling approach.

Journal of Personality and Social Psychology. Advance online publication.

<https://doi.org/10.1037/pspi0000279>

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>

van Ravenzwaaij, D., Dutilh, G., & Wagenmakers, E.-J. (2012). A diffusion model

decomposition of the effects of alcohol on perceptual decision making.

Psychopharmacology, 219(4), 1017–1025. <https://doi.org/10.1007/s00213-011-2435-9>

Voss, A., Rothermund, K., Gast, A., & Wentura, D. (2013). Cognitive processes in associative

and categorical priming: A diffusion model analysis. *Journal of Experimental Psychology*.

General, 142(2), 536–559. <https://doi.org/10.1037/a0029459>

- Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, 32(7), 1206–1220.
<https://doi.org/10.3758/BF03196893>
- Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods*, 39(4), 767–775. <https://doi.org/10.3758/BF03192967>
- Voss, A., Voss, J., & Lerche, V. (2015). Assessing cognitive processes with diffusion model analyses: A tutorial based on fast-dm-30. *Frontiers in Psychology*, 6, 336.
<https://doi.org/10.3389/fpsyg.2015.00336>
- Wabersich, D., & Vandekerckhove, J. (2014). Extending jags: A tutorial on adding custom distributions to jags (with a diffusion model example). *Behavior Research Methods*, 46(1), 15–28. <https://doi.org/10.3758/s13428-013-0369-3>
- 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>
- Wandert, T., Ochsmann, R., Brug, P., Chybicka, A., Lacassagne, M.-F., & Verkuyten, M. (2009). Black german identities: Validating the multidimensional inventory of black identity. *Journal of Black Psychology*, 35(4), 456–484.
<https://doi.org/10.1177/0095798408329949>