Behavioural Science Institute

The Influence of Group Membership on Prioritization for Visual Awareness

Désirée Kleverwal

Master of Science Thesis

Supervised by:
1. Dr. Gijs Bijlstra
2. Prof. Dr. Rob Holland

July 3rd, 2018

Radboud University

Abstract

As our ability to detect everything in our visual environment is limited, an important question focuses on how visual detection mechanisms are influenced by the social saliency of stimuli. In the current experiments, we investigated whether in-group information is preferentially detected for visual awareness as opposed to out-group information. Utilizing a breaking – continuous flash suppression (b-CFS) task, we measured the duration of detection of in- versus out-group associated symbols (Experiments 1 and 2), but did not find any differences between in- versus out-group information on visual awareness. In Experiment 2, we additionally measured the duration of detection of social group *faces*, and found evidence that Western perceivers preferentially detect Western (over Arabic) faces. However, for Arabic perceivers, we did not show such in-group prioritization effects. Rather than being dependent on low-level visual characteristics, visual detection mechanisms may be influenced by one's own social group. These findings may have implications for understanding the origins of social biases.

Keywords: social identity, self-relevance, face detection, visual awareness, breaking - continuous flash suppression

The Influence of Group Membership on Prioritization for Visual Awareness

People seem to have a notable tendency to sympathize and identify with groups. Based on their race, age or gender, but also sports teams or religious affiliations, people belong to multiple social groups. Key to this is the concept of social identity – deriving a sense of who you are, based on your group membership (Tajfel, 1970). According to social identity theory (SIT; Tajfel, Billig, Bundy, & Flament, 1971), people understand the social world as consisting of in-groups – those groups to which the individual belongs – and out-groups – those groups to which the individual does not belong. Accordingly, identification with one's in-group motivates the individual to distinguish their group from others, which is an important source of collective pride and self-esteem (Brewer, 1999). One of the most common ways of doing this involves in-group favoritism: the systematic tendency to evaluate one's own group or its members more favorably than an out-group or its members (Tajfel, 1982). This in-group favoritism has been shown to occur as the preferential treatment of in-group members, rather than the direct detrimental treatment of out-group members (Allport, 1954; Brewer, 1999). In line with this, studies have shown that mere categorization into arbitrary groups by means of the minimal group paradigm, immediately leads to in-group favoritism (Tajfel et al., 1971), and more positive affect towards in-groups (Otten & Wentura, 1999).

Categorizing other people into groups can have consequences for a variety of processes, including perception (Bijlstra, Holland, & Wigboldus, 2010), choices, and behavior (Amodio & Devine, 2006; see also Kawakami, Friesen, & Vingilis-Jaremko, 2018). That is, social group information not only influences how we behave towards another person (Amodio & Devine, 2006), but also how we attend to that person's face. For instance, already at 3 months of age, infants demonstrate a looking preference for own- versus other-race faces (Bar-Haim, Ziv, Lamy, & Hodes, 2006), and also in adulthood we seem to attend more to faces of our own race (Lovén et al., 2012) and age (He, Ebner, & Johnson, 2011). At the same time, under some conditions attention also seems to be directed to out-group members. For example, out of vigilance (Bean et al., 2012) or when the out-group is perceived as threatening or dangerous (Maner & Miller, 2013; Trawalter, Todd, Baird, & Richeson, 2008). Thus, at least in non-threatening social contexts, people seem to have preferential attention for in- as opposed to out-group faces. Results like these are not to be taken lightly, as attentional biases can considerably affect both interpersonal interactions and intergroup attitudes (Kawakami et al., 2018).

Face Detection

Attending to faces can provide us with a great deal of social information. However, before we can acquire this information, the presence of a face in our visual environment needs to be detected. While a lot of studies have been conducted to investigate the preferential attention for own-group faces, surprisingly little is known about the perceptual mechanisms underlying the *initial* detection of in- versus out-group faces. That is, do in- versus out-group differences also play a role in what we (do not) attend to in the first place?

Recent studies employing a breaking - continuous flash suppression (b-CFS) paradigm seem to suggest so. In b-CFS (Tsuchiya & Koch, 2005; also see Stein, Hebart, & Sterzer, 2011), dynamic Mondrian patterns flashed to one eye render stimuli presented to the other eye invisible for prolonged periods of time, making it possible to sensitively measure detection threshold differences of visual stimuli (Stein et al., 2011). That is, to ensure that the Mondrian patterns gain visual dominance first, the patterns are presented immediately at full contrast, yet the contrast of the stimulus of interest is gradually ramped up. Consequently, one can measure the time it takes for the stimulus to overcome suppression of the dynamic patterns. This measure, suppression duration, is seen as an indicator of how different stimuli have an advantage in attaining perceptual dominance. This advantage in perceptual dominance can consequently be seen as preferential processing of that stimulus (Stein et al., 2011). Thus, with regard to face detection, by measuring the suppression durations of different faces, b-CFS can investigate how different faces are differentially processed, and thereby investigate how initial visual processing can be influenced by the social group a target face belongs to.

In doing so, Stein, End and Sterzer (2014) employed this b-CFS paradigm and investigated whether pictures of in-group faces were detected more quickly than pictures of out-group faces. That is, they showed that pictures of own-race and own-age faces were detected more quickly compared to pictures of other-race and other-age faces. Similarly, Gobbini and colleagues (Gobbini et al., 2013) showed that, in a b-CFS paradigm, pictures of friends' faces were detected more quickly than pictures of strangers' faces. Together, these studies seem to indicate a preferential detection of in- as opposed to out-group faces with regard to visual awareness. However, although we learned a great deal from these studies, it is yet unknown what seems to cause this preferential processing of in-group faces. Several explanations are plausible for the underlying effects, including variability in low-level visual characteristics of the stimulus, familiarity of the stimulus, or the relevance of the stimulus for the self. We will now outline each of these possibilities.

First, stimuli with different low-level visual characteristics influence access to awareness differently in b-CFS tasks (for an overview see: Gayet, Van der Stigchel, & Paffen, 2014). For example, studies have shown that higher contrast stimuli break into awareness more readily than lower contrast stimuli (Tsuchiya and Koch, 2005). As faces from different social groups are found to differ in these low-level visual characteristics (e.g., face shape; Berry and McArthur, 1986), this may have influenced previously found preferential processing of in-group faces effects. That is, whereas Stein et al., (2014) showed that Caucasian perceivers preferentially detect Caucasian (compared to Black) faces for visual awareness, or that young perceivers preferentially detect young (compared to elder) faces, these effects were only shown for Caucasian, young perceivers. Therefore, these findings are open to the possibility that detection mechanisms are tuned to young faces with light skin color (i.e., low-level visual characteristics), regardless of the perceiver's own social group (Stein et al., 2014). If the latter is the case, a Black perceiver would also preferentially detect a Caucasian face over Black face. However, if in-group information is indeed prioritized in visual awareness, then a Black perceiver would preferentially detect a Black face over a White face, and prioritization effects in visual awareness would thus depend on the perceivers' in-group and not necessarily on the low-level visual characteristics of the perceived face.

Second, in b-CFS, even when low-level visual characteristics are controlled for, higher-level visual characteristics also influence access to awareness differently in b-CFS tasks. That is, studies have shown that more familiar stimuli are prioritized in visual awareness over more unfamiliar stimuli. For instance, images of both human faces (Jiang, Costello, & He, 2007; Yang, Zald, & Blake, 2007; Stein et al., 2011), and human bodies (Stein, Sterzer, & Peelen, 2012) are detected faster when presented upright as opposed to inverted. Thus, with regard to social groups, own- race and age (Stein et al., 2014), or friends' (Gobbini et al., 2013) faces can be considered more familiar than other- race and age, or strangers' faces, respectively. Therefore, it may be in-group facial familiarity that drives in-group prioritization effects.

Third, however, recent studies suggest that it may be the self-relevance of in-group information, rather than facial characteristics, or facial familiarity, driving these preferential processing effects. That is, an even higher level of preferential in-group processing. Namely, Macrae and colleagues (Macrae, Visokomogilski, Golubickis, Cunningham, & Sahraie, 2017) hypothesized that, as self-relevant information is prioritized during various stages of processing, this information may also be prioritized in visual awareness. By associating arbitrary geometric shapes with person-related labels (i.e., self, friend, stranger) and consequently suppressing these shapes by use of b-CFS, it was shown that participants could report the identity of self-relevant shapes faster than the identity of other-related shapes. Note that by using geometric shapes instead of facial stimuli, the authors were able to rule out both visual characteristics and facial familiarity effects: the shapes are perceptually arbitrary. Thus, as the authors concluded, self-relevance may enhance visual awareness (Macrae et al., 2017). Conversely, however, Stein and colleagues (Stein, Siebold, & van Zoest, 2016) simultaneously conducted a similar b-CFS study in which they measured the detection threshold for self- and other-associated figures and were not able to find any self-related prioritization effects. Note however, that this non-replication could be explained by several methodological differences. Whereas Macrae et al. (2017) had perceivers report the *identity* of the figures, Stein et al. (2016) had perceivers report the *location* of the figures. This difference between identification and detection may be reflected in differences in processing stages. That is, reporting location may reflect an earlier processing stage than reporting identity, as for reporting a stimulus' location it is not yet necessary to be aware of its identity. Nevertheless, as of yet, no firm conclusions on how initial visual processing could be influenced by the personal relevance of information, can be drawn.

Present Research

In the current, pre-registered, experiments (see https://osf.io/rxztj for both pre-registrations), we aimed to further investigate the role of self-relevance in the prioritization of in- versus out-group information in visual awareness, while employing a design suitable to control for both familiarity and visual characteristics effects. Experiment 1 employed a b-CFS paradigm in which the suppression duration of geometric figures (i.e., symbols) previously associated with in- versus out-group information was measured. As mentioned above, using these symbols (see also: Sui, He, & Humphreys, 2012; Stein et al., 2016; Macrae et al., 2017) enabled us to rule out both in-group familiarity effects and mere visual characteristics effects, as the associated symbols were both equally familiar to the participant, and perceptually similar to one another. Additionally, to rule out in-group familiarity and other confounds like out-group threat (Maner & Miller, 2013; Trawalter et al., 2008) the symbols were associated with a *novel* in- versus out-group, created by means of the minimal-group paradigm (Tajfel et al., 1971). In this paradigm, people are assigned to groups based on factors that are *not* typically expected to account for intergroup bias (e.g., stereotypes, prior conflict, or competition for resources), yet assignment to these groups can be sufficient to evoke identification with in-group members (Otten & Bar-Tal, 2002; Otten & Wentura, 1999; Van Bavel, Packer, & Cunningham, 2011; for a review see Brewer, 1979). Thus, in line with social identity theory, even these minimal groups become relevant to the self (Tajfel et al., 1971). By utilizing minimal groups, we are able to study with additional experimental control whether it is only the self-relevance of in-group information that prioritizes access into visual awareness.

In Experiment 2, to further investigate whether prioritization in visual awareness can be influenced by a perceivers' own social group, we used *existing* as opposed to novel groups, and included participants from both Arabic and Western origins.

Hypotheses

Theoretically, our main hypothesis entailed that in-group information is preferentially selected for visual awareness as opposed to out-group information (i.e., Gobbini et al., 2013; Stein et al., 2014). With regard to the b-CFS paradigm, this resulted in the following two-fold hypothesis: if group associations with symbols can influence detection time of these symbols in a b-CFS paradigm, then in-group associated symbols are detected faster compared to out-group associated symbols (i.e., a self-relevance effect of symbol associations on detection time). However, if group associations with symbols are not sufficient to influence their detection time in a b-CFS paradigm, then visual input of group information is necessary to influence detection time (i.e., either a familiarity or visual characteristics effect of group information on detection time).

Method Experiment 1

Participants

In total, 65 participants recruited from the Radboud University repository system for participation in research, took part in the experiment. Only participants with normal eye-sight were eligible to participate. In line with our pre-registered exclusion criteria, twelve participants were excluded from analyses because they had more than 10% of the trials in the b-CFS task incorrect. An additional five participants were excluded because they were unable to see one square in the b-CFS task[1]. Therefore, in the main analyses, data of 48 participants was used (72.9% female, M_{age} = 23.8, SD_{age} = 4.8). Participants were offered course credit or a fixed fee of €5,- for their participation.

Overview and Design

First, participants were assigned to groups based on a minimal-group categorization procedure (i.e., dot-estimation task). Next, both their own group and the other group were coupled with geometric figures (i.e., symbols). After playing a competitive game to boost the strength of in-group feelings, a learning task was performed to strengthen group-symbol associations. Next, to measure the suppression duration of these in- versus out-group associated symbols, the symbols were presented within the b-CFS paradigm. The experiment had a within-subjects design, investigating differences in suppression duration for in- versus out-group associated target symbols.

Materials & Procedure

The experiment was programmed using PsychoPy (Peirce, 2007). A computer mouse and keyboard were used for collecting responses. During the whole experiment the background color of the screen was gray, and all text was presented in white.

[1] Although this was not a pre-registered criterion for participants to be excluded, it is a crucial aspect of b-CFS that participants see one square (i.e., that the two visual fields are placed on top of one another). If not, the b-CFS measure of suppression duration would not be valid, as a stimulus would not actually be suppressed from awareness.

Minimal group categorization. Participants were welcomed into the lab, seated in a cubicle, and asked to provide demographic information (sex, age, nationality). Then, following Tajfel (1970) participants underwent a minimal-group categorization based on a dot-estimation task. In this task, varying numbers of dots (ranging between 10 and 50) were projected on the screen and the task of the participant was to estimate how many dots were shown. On every trial, the dots were visible for 500ms, after which participants were asked to type in an estimation of the number. After two practice trials, participants performed five trials for estimation. In order to avoid deception as much as possible, participant's actual performance in this task determined to which group they were consequently assigned. If participants tended to overestimate the number of dots shown - i.e., in three or more cases - they were assigned to the 'overestimators' group, whereas if they tended to underestimate the number of dots shown, they were assigned to the 'underestimators' group. In total, 29 participants were classified as ‘underestimators’, and 19 participants were classified as ‘overestimators’.

Symbol assignment. Next, participants were told that "from now on, both groups will be represented by a logo". Symbol assignment was based on participant number, with uneven versus even participant numbers getting assigned the square versus diamond, respectively. To avoid and control for perceptual differences between groups, target symbols consisted of a square with two successively smaller squares inside, and its rotated version: a diamond (see Figure 1). These symbols were chosen to be perceptually similar, and the smaller versions inside were added to slightly increase the symbols' captivation. Consequently, participants got to see 'my group: over/underestimators' and 'other group: under/ overestimators' presented at the top of the screen, with both logos presented underneath.

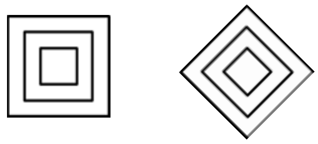

a. b.

Figure 1. Target symbols presented in the b-CFS task: a square (a) and a diamond (b).

Snake-game. As groups were novel to participants, we aimed to strengthen group feelings by having groups compete against one another (Ashmore, Jussim, & Wilder, 2001), by having participants play a computerized snake-game. In this game, the task was to move the snake in order to eat as many treats as possible. By using the arrow keypad the snake could be moved throughout the field, yet not through itself nor through the walls of the field. With every treat, the snake grew longer, and with every five treats, the speed of the snake increased by 20%. The game consisted of a practice-round (automatically terminated after 5 points) and two ‘actual’ rounds. During the practice-round, a score-counter at the top-left edge of the field displayed "you earned xx points", whereas during the actual rounds this score-counter displayed "you earned xx points for your group". The latter was added to emphasize the possible in-group gain.

Every round would terminate when the snake bumped into the wall or itself. After every round, if game-over, it would display "Game over! In this round you earned xx points for your group. Press Esc to close the window to go to the next round, or to return to the experiment.". The snake-game was played individually, yet by playing the game participants won points for their own group. Participants were told to win as many points for their own group, so that together with their group, they will earn more points than the other group. Additionally, participants were told that at the end of the experiment (when all data from all participants is collected), a prize would be raffled among all members of the winning group, and that the round in which they scored the most points would count to their group total.

Learning task. To create a stronger association between the symbols and groups we included a learning task (see Stein et al., 2016). After the snake-game, participants were told that we “are interested in seeing how fast you can categorize the logos associated with the groups”. Participants saw the screen with the logos and groups again, and were told to memorize which logo was represented which group, after which they performed the learning task.

In the task, participants were presented with one of the symbols on each trial and asked to indicate whether the logo belonged to the overestimators or underestimators. Participants were instructed to press one designated key if the symbol belonged to the overestimators (A), and another designated key if the symbol belonged to the underestimators (L). During this task, the top-left corner of the screen displayed 'Press A for over/underestimators', and the top-right corner of the screen displayed 'Press L for under/overestimators', with response keys counterbalanced between blocks.

Each symbol was presented for a maximum of 3 seconds. If participants responded correctly to the symbol within this timeframe, the word ‘correct!’ was presented for 1 second, whereas if they responded incorrect, or not within the timeframe, the presented word was ‘wrong!’ or ‘too late!’, respectively. After this feedback signal, the next symbol was presented. Participants first completed a practice block of four trials, with both symbols presented twice. The learning task consisted of 40 trials: two blocks of 20 trials each, with both symbols presented ten times per block. When finished, participants were instructed to get the experimenter, who guided them to another cubicle where the b-CFS task was performed and the minimal group questions were answered.

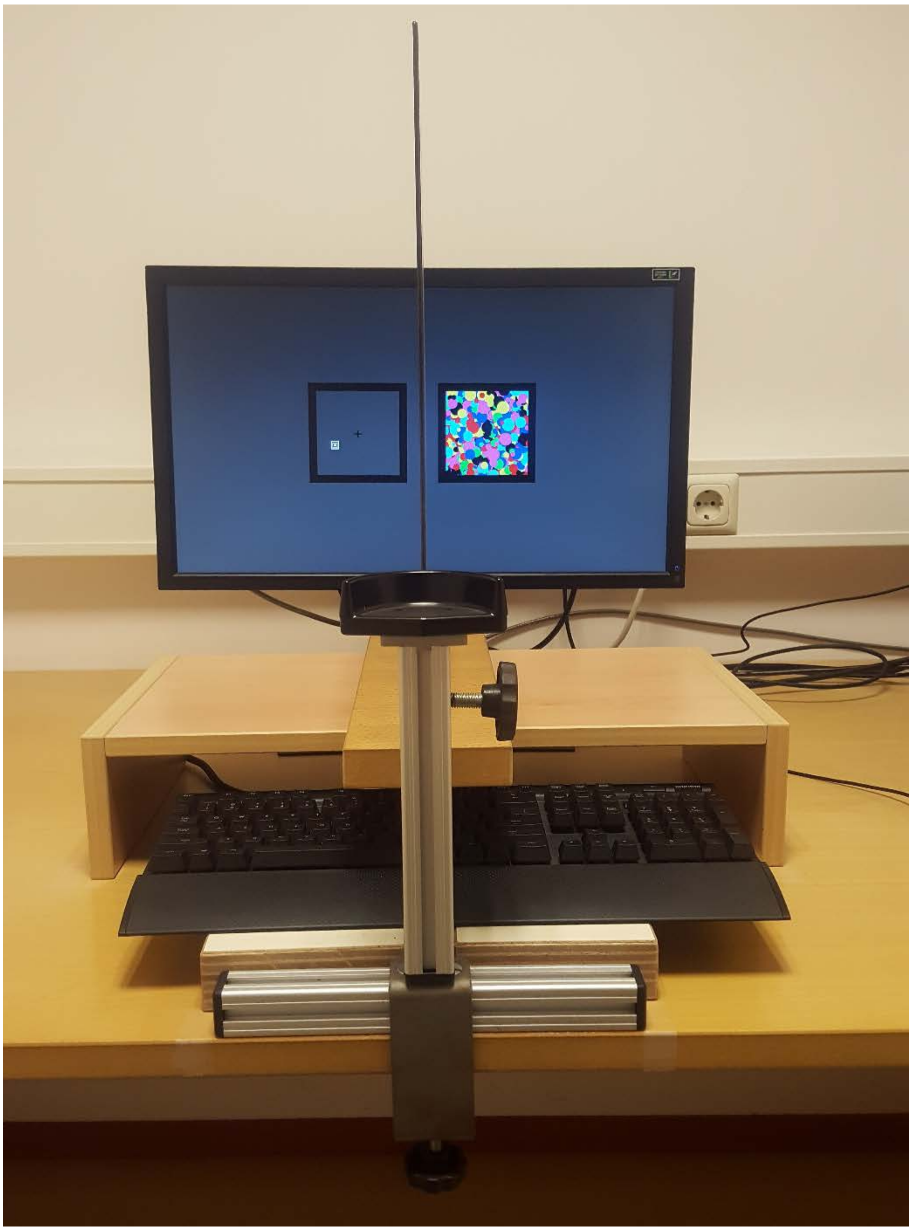

Figure 2. The b-CFS set-up. Participants observe the display, split by a screen divider, while wearing prism glasses. The screen divider and prism glasses ensure that the two visual fields are placed on top of one another. To stabilize fixation, a chinrest is used.

Breaking – continuous flash suppression task. In the b-CFS task (see Tsuchiya & Koch, 2005; Stein et al., 2011), two black frames with both a black fixation cross in the middle were presented on a gray background. This remained on the screen for the whole b-CFS procedure. With use of a screen divider, a chin rest and prism glasses (see Figure 2), participants' two visual fields were placed on top of one another as to look like one (i.e., participants see one black frame with one fixation cross).

In every b-CFS trial, presented to one - randomly determined - eye were dynamic Mondrian patterns, consisting of many differently colored and sized circles placed at random locations within the frame. Almost simultaneously, the other eye was presented with one of the previously learned symbols, appearing either on the left or the right of the fixation cross. Each trial started with the dynamic Mondrian patterns presented at 100% of their contrast. Then, after 500ms, the contrast of the geometric figure was slowly ramped up (from 0 to 96% in steps of 1,6% every screen refresh – once every ~16.67ms on a 60 Hz screen). When reaching 96% after 1000ms, the contrast of the dynamic Mondrian patterns started to decrease from 100% to 0%, reaching 0% after 8500ms (i.e., the end of the trial). A complete trial lasted for a maximum of 10 seconds. See Figure 3 for a schematic example of a b-CFS trial. By slowly ramping up the contrast of the geometric figure, the target figure is temporarily suppressed from awareness. The dependent variable (suppression duration) is the time it takes from the onset of the trial to the moment participants press a response button indicating the location of the stimulus. Thus, longer reaction times mean longer suppression durations, indicating that the suppressed stimulus takes longer to reach awareness.

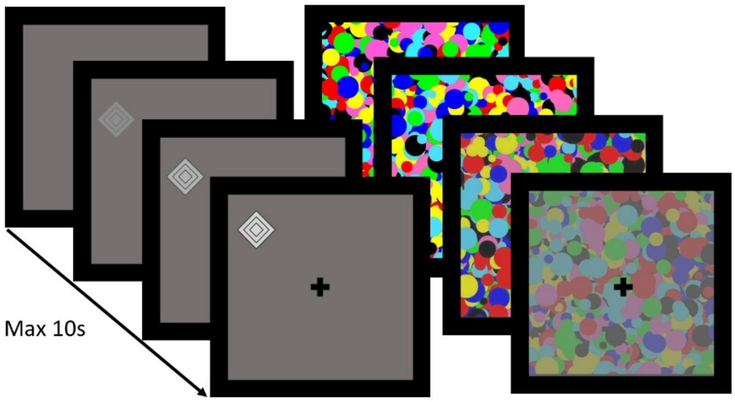

Figure 3. Schematic example of a b-CFS trial. Over the course of 10 seconds, the contrast of the target stimulus is gradually ramped up, whereas the contrast of the dynamic Mondrian patterns gradually decrease.

The task of the participant was to indicate the location (left or right from the fixation cross) of the figure by pressing either the S on the left side of the keyboard (to indicate left) or the 5 on the right side of the keyboard (to indicate right). Having participants indicate the side on the keyboard on which the image appears allowed us to exclude data from those participants who were merely guessing and pressing the response buttons without being aware of the right position of the stimulus.

In the b-CFS task, both target symbols (diamond and square) appeared 34 times, of which 17 times on both the left and right side of the fixation cross, leading to a total of 68 trials per participant.

Minimal group questions. At the end of the experiment, based on Otten & Moskowitz (2000), eight items were added as a manipulation check for the minimal group paradigm and for additional exploratory purposes. All items were answered on a 5-point Likert scale ranging from 1-5 (not at all – neutral – very much).

The first four items were added as a manipulation check for the minimal group paradigm, and entailed group identification: (1) Were you happy to belong to your group?, (2) Would you have preferred to belong to the other group?, (3) I felt connected to my group, and (4) I felt connected to the other group. That is, we wanted to see whether the minimal group paradigm worked sufficiently in that it created classic in-group favoritism effects (i.e., Allport, 1954; Brewer, 1999; Tajfel, 1982).

The last four items were added in an explorative manner. First, we were interested whether in-group effects would also manifest itself in the group-symbols. That is, if arbitrary assignment to groups leads to in-group favoritism then perhaps arbitrary assignment to group-logos also leads to some sort of in-group logo favoritism. Thus, the next two items entailed symbol identification: (5) Were you happy with the logo that represented your group?, and (6) Would you have preferred the other logo to represent your group?.

Finally, we were interested in seeing whether participants would assign meaning to the minimal groups. That is, whether they believe the groups to be real and relevant. Thus, the last two items entailed perceived relevance: (7) Do you think belonging to either the overestimators or underestimators is also important in other domains of life? and (8) Do you think overestimators and underestimators also differ significantly in other respects?.

When finished filling in the items, participants were thanked, paid, and debriefed.

Data-analysis

As the first six items regarding the minimal group paradigm consisted of pairs with one question mirroring the other question (e.g., happy with own-group versus happy with other group), the data were analyzed for each pair separately, conducting a one-sample *t*-test on the average scores of the pair to see whether they differed significantly from the scales' midpoint 3 (with the scores of one question per pair reversed). As the last two items on perceived relevance (item seven and eight) were not each other's opposites, these items were analyzed separately by two one sample *t*-tests investigating whether scores differed significantly from the scales' midpoint.

Regarding the b-CFS task, reaction times (i.e., suppression durations) of correct trials were the main outcome measure of interest. In line with Yang and colleagues (Yang et al., 2007), and our pre-registration, participants with errors in more than 10% of the b-CFS trials were excluded from analyses (*N* = 12), and for the remaining participants all incorrect trials were excluded (2.11%). Additionally, reaction times deviating more than 3 *SD* from the average were regarded as outliers, and therefore excluded (3.42%).

The b-CFS data were examined using three different methods: Bayesian Inference, Frequentist *t*-tests, and Linear Mixed Models. As all three methods yielded the same results, only the Bayesian Inference methods are reported here (for the latter two see the Appendix). Whereas Bayesian analyses were performed in JASP (JASP Team, 2018), all other analyses were performed in R (R Core Team, 2016). Because the (difference scores of) suppression durations were not normally distributed, all analyses were conducted using log-transformed suppression durations. For interpreting purposes, only the non-transformed durations are reported.

By estimating Bayes Factors (BFs; Jeffreys, 1961; Rouder, Speckman, Sun, Morey, & Iverson, 2009) we were able to quantify the relative evidence for two models, typically a null hypothesis and an alternative hypothesis. BF_{10} denotes a ratio: the likelihood of the data occurring under the alternative model, divided by the likelihood of the data occurring under the null model, whereas BF_{01} denotes its' inverse ($BF_{01} = 1 / BF_{10}$). For example, if $BF_{10} = 5$, then the observed data are five times more likely to have occurred under the alternative compared to the null model. If $BF_{10} = 0.5$, then $BF_{01} = 1/ 0.5 = 2$, meaning that the observed data are two times more likely to occur under the null compared to the alternative model. Thus, BFs allow for quantification of support for either the alternative or the null hypothesis. With regard to the strength of evidence gathered in favor of either model, Jeffreys (1961) developed guidelines for language that may be used to interpret Bayes Factors. See Table 1 for the suggested terminology.

Table 1

Jeffreys (1961) classification scheme for interpreting Bayes Factors

Bayes Factor BF_{10}	Support for H1
1 - 3	Anecdotal
3 – 10	Substantial
10 – 30	Strong
30 – 100	Very Strong
> 100	Decisive

Results Experiment 1

Manipulation Check Minimal Group Paradigm

The minimal group paradigm created an in-group preference that was more positive than the midpoint of the scale. Namely, both with regard to the own- versus other-group preference items (M = 3.40, SD = 0.80), $t(47)$ = 4.37, p < .001, and with regard to the own- versus other-group feelings of connectedness items (M = 3.66, SD = 1.19), $t(46)$ = 6.76, p <.001, participants' scores significantly deviated from the scales' midpoint 3.

The explorative items on symbol identification and perceived relevance of groups, yielded mixed results. First, scores on the symbol-identification items significantly deviated from the scales' midpoint (M = 3.44, SD = 0.96), $t(47)$ = 3.50, p = .001, indicating that the minimal group paradigm not only worked sufficiently to evoke in-group preference, but also to evoke in-group symbol preferences. However, with regard to the perceived relevance of groups, scores on item seven (i.e., the importance of over/underestimators in other domains of life) did not differ significantly from the scales midpoint, $t(47) = 1.14$, $p = .262$, whereas scores on item eight (i.e., other differences between over/underestimators) marginally significantly differed from the scales' midpoint (M = 3.38), t(47) = 2.11, p = .04. These latter two items were added in an exploratory manner to investigate whether participants would assign meaning to the minimal groups. However, no firm conclusions about this can be drawn.

Breaking - Continuous Flash Suppression

The mean suppression duration for in-group associated symbols was 1.35s (*SD* = 0.62), whereas the mean suppression duration for out-group associated symbols was 1.38s (*SD* = 0.73). To quantify the evidence for our hypothesis that in-group associated symbols yield shorter suppression durations compared to out-group associated symbols, the log-transformed suppression durations were subjected to a one-sided Bayesian independent samples *t*-test with the default Cauchy prior width selected in JASP. An estimated Bayes Factor ($BF_{01} = 8.90$), comparing the fit of the data under the null hypothesis and the alternative hypothesis, suggested that the data were 1:8.90 in favor of the null hypothesis, or rather 8.90 times more likely to occur under a model postulating the absence of the effect. This is substantial evidence for the null hypothesis (Jeffreys, 1961).

Exploratory Analyses

In an exploratory manner, we further analyzed the suppression duration data to check for possible differences between symbols. This additional dependent *t*-test showed a significant effect of symbol, with the square ($M = 1.27$, $SD = 0.89$) on average yielding shorter suppression durations than the diamond ($M = 1.35$, $SD = 0.96$), $t(47) = 2.60$, $p = .013$.

Additionally, we decided to check for possible differences in suppression durations between minimal groups. This dependent *t*-test showed a marginally significant effect of in-group, with overestimators (*M* = 1.17, *SD* = 0.79) having shorter suppression durations than underestimators (*M* = 1.41, *SD* = 1.0), *t*(47) = 1.96, *p* = .056.

Discussion Experiment 1

In Experiment 1, we hypothesized that, as in-group information can be considered more relevant to the self than out-group information (Tajfel & Turner, 1979), in-group information is preferentially detected in visual awareness as opposed to out-group information (i.e., Gobbini et al., 2013; Stein et al., 2014). Utilizing a b-CFS paradigm where novel groups were associated with symbols, we did not find evidence for an effect of in- versus out-group information on prioritization for visual awareness. Instead, we found more evidence for the null hypothesis stating the absence of an effect.

Several factors might explain this finding. First, on a methodological level, the symbols we used elicited relatively short suppression durations compared to other studies utilizing b-CFS. Whereas b-CFS should be able to suppress stimuli up till 10 seconds (e.g., Stein et al. (2014) used parameters similar to ours for presenting the stimuli, and their average suppression duration was around 5 seconds), the average suppression duration in the current experiment was only around one second. A reason for this difference could be the spatial frequency of the used symbols. Namely, because the symbols had sharp lines and edges, their spatial frequency was relatively high. However, Yang and Blake (2012) recently showed that suppression durations in b-CFS are relatively shorter for high as opposed to low spatial frequency stimuli. Consequently, in our task, both symbols may have been detected too quickly, leaving little room for variance in suppression durations between symbols.

Second, maybe *new* associations, especially between *new* groups and *new* symbols, are not strong enough to elicit in-group prioritization in visual awareness. That is, although previous work did show prioritization effects of abstract figure associations in b-CFS (i.e., Macrae et al., 2017), this was done using self-associated symbols versus other-associated symbols. Whereas the distinction between self and others has had time to form over a lifetime, the current groups were only created just now. Thus, although we did find evidence indicating in-group preference, participants may not had have enough time to associate this in-group with the self, and therefore the subsequent group-symbol associations may have been too weak to elicit any prioritization effects. Perhaps if people have more time to identify with their minimal group, or if symbols are associated with *existing* groups they already have strong associations with, in-group prioritization effects would occur.

Additionally, the Macrae et al. (2017) work may have yielded different results because of the used b-CFS task. Namely, whereas Macrae et al. (2017) instructed participants to respond to the *identity* of the symbols and consequently found self-related prioritization effects, another similar study instructed participants to respond to the *location* of the symbols, and was not able to find any self-related prioritization effects (i.e., Stein et al., 2016). As mentioned in the introduction, this difference between identification and detection may reflect different processing stages, with stimulus identification representing later processing stages than stimulus detection. It follows that the Macrae et al. (2017) study may not have measured the *initial* detection of self-associated symbols, whereas the Stein et al. (2016) study did, possibly explaining their, and our, mixed results.

As we were interested in the initial detection of group information, the null findings of Stein et al. (2016) are still relevant to the current discussion. Namely, their study may indicate that associations with abstract figures are not sufficient to influence suppression durations in the b-CFS paradigm specifically. That is, it may well be the case that the paradigm is not equipped to measure these abstract group-association with symbols, as the underlying mechanisms of b-CFS are relatively unknown (see Stein & Sterzer, 2014). Instead, visual input of group information might be needed, since other studies that investigated group information while using pictures of faces (as opposed to group-associations), were able to yield differences for in- versus out-group suppression durations (i.e., Gobbini et al., 2013; Stein et al., 2014).

In sum, even though participants reported a preference both for their own group and for the symbol associated with their own group, we found more evidence for the null hypothesis stating that in- versus out-group associated symbols do not differ in their detection times. This may be because of the spatial frequency of the symbols, because group associations were not strong enough, or because b-CFS is not capable of measuring associations, and instead needs visual input of group information. In Experiment 2, we tried to account for (and investigate) these possibilities in several ways.

Experiment 2

First, in Experiment 2, we reduced the spatial frequency of the stimuli by using Gabor patches as opposed to geometric figures as symbols for the groups (i.e., Stein et al., 2016). As Gabor patches are known for their low spatial frequency (i.e., blurred lines), we aimed to increase suppression durations in the b-CFS paradigm, and thereby allowing for more variance between symbols. Second, to ensure that participants' group-associations were already formed, we used existing as opposed to novel groups, by recruiting participants that were either from Arabic or Western origins. Third, to account for the possibility that b-CFS is not sensitive to group-symbol associations, we added another b-CFS block that instead utilized pictures of faces.

Thus, the b-CFS task consisted of two blocks: one utilizing Gabor patches associated with an Arabic and Western group, and one utilizing pictures of faces belonging to those groups. By including participants from both groups, we were able to explore whether in-group prioritization effects are two-sided, and exist regardless of lower-level visual characteristics. That is, Stein et al. (2014) showed that Caucasian and young participants, prioritize Caucasian (as opposed to Black) and young (as opposed to elderly) target faces for visual awareness, respectively, yet did not include Black or elderly participants to investigate whether in-group prioritization effects also occur for them. Therefore, it is not possible to differentiate between prioritization effects due to low-level visual characteristics of Caucasian and young faces, or prioritization effects based on the perceivers' own in-group. The latter has never been investigated using a full design where both groups are represented in both the task and the participant sample. Thus, we aimed to conceptually replicate Stein et al. (2014) by showing that White participants prioritize White target faces for visual awareness, and we simultaneously aimed to extend their findings by investigating these effects for an Arabic in- and out-group.

With regard to the group-associated symbols, we hypothesized that *if* group-associations with symbols influence detection times in the b-CFS paradigm, then in-group associated symbols would be detected faster than out-group associated symbols. With regard to the pictures of faces, we hypothesized that the ethnicity of the target face (Arabic x Western) would influence the detection time of these faces in the b-CFS paradigm, with Western (compared to Arabic) faces being detected faster by Western participants.

With regard to Arabic target faces and participants we accounted for several outcomes. Namely, if face detection mechanisms are generally tuned to the low-level visual characteristics of a face, then both the Arabic and Western subgroup should prioritize Western faces in visual awareness (i.e., in line with Stein et al., 2014). However, if facial detection processes can be influenced by the perceiver's own social group, either self-relevance or familiarity driven, then differences between the subgroups are likely to arise, as for both groups, their in-group can be considered more familiar, and more relevant to the self. Since these effects have never been investigated including an Arabic group, our hypotheses regarding this group were exploratory.

Method Experiment 2

Participants

In total, 47 participants recruited either through informal and social media channels, or from the Radboud University repository system for participation in research, took part in the study. In order to be eligible to participate in Experiment 2, subjects had to have a) normal eye-sight, and b) fall under either one of two categories: Arabic or Western. In the study, subjects were considered ‘Arabic’ when the participant himself, or (at least) one of his parents, was born in the MENA-region (Middle-East – North-Africa, including Turkey), whereas subjects were considered ‘Western’ when the participant himself, and both of his parents, were born in the Netherlands or Germany[2]. Depending on the participants’ native language, all instructions were written and given either in English or Dutch.

Out of 47 participants, two were excluded because they were unable to see one square in the b-CFS task, leading to a total of 30 participants of Western ethnicity (56% female, $M_{\text{age}} = 24.2$, $SD_{\text{age}} = 4.77$) and 15 participants of Arabic ethnicity (57% female: 1 missing, $M_{\text{age}} = 24.8$, $SD_{\text{age}} = 4.43$). Participants were offered course credit or a fixed fee of €5,- for joining the study.

Overview and Design

Thus, participants were either from Arabic or Western origins. For every participant, both these groups got coupled with symbols, after which a learning task was performed to strengthen group-symbol associations. Next, to measure the suppression duration of both in- and out-group associated symbols, and of pictures of Arabic and Western faces, two separate

[2] The 'Arabic' operationalisation was chosen in order to still have a pool of participants to recruit from (as participants had to come to the lab, and people born in the MENA-region themselves are scarce in the Netherlands), and simultaneously have a pool of participants that still feel connected to their Arabic in-group (i.e., participants who's grandparents were born in the MENA-region were not included, as their Arabic in-group feelings may be decreased). Germans were included as enrolled in our university are students mainly coming from the Netherlands or Germany.

b-CFS blocks were performed: a symbol- and faces-block. Both the symbol- and faces-block yielded a mixed design investigating effects of the dependent variable suppression duration, with Ethnicity (Arabic x Western) as a between-subjects factor, and Target Symbol (in-group associated x out-group associated), or Target Face (in-group x out-group) as within-subjects factor, respectively.

Materials and Procedure

Experiment 2 utilized a similar learning task and b-CFS task as in Experiment 1. However, in Experiment 2, no competitive game between groups was included, nor were any questions about group identification asked.

First, participants were welcomed into the lab, seated in a cubicle, and asked to provide demographic information: sex, age, nationality and ethnicity. In Experiment 2, group-symbols differed from Experiment 1 in that the symbols were Gabor patches as opposed to geometric figures. Namely, the target symbols were Gabor patches either *vertically* or *horizontally* oriented (see Figure 4), with the same mean luminance as the gray area against which they were presented (in line with Stein et al., 2016). Participants were told that “in the next part of the experiment, both groups (Arabic / Western) are coupled to a symbol (horizontal or vertical stripes)”. Again, symbol assignment was based on participant number, with uneven versus even participant numbers getting assigned the horizontal versus vertical Gabor patches, respectively. Next, participants got to see the words ‘Arabic’ and ‘Western’ presented at the top of the screen, with the respective Gabor patches presented underneath.

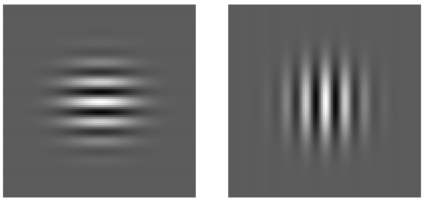

a. b.

Figure 4. Target symbols presented in the b-CFS task: Gabor patches oriented either horizontally (a) or vertically (b).

Next, participants were told that “in the next part, to learn accurately which symbol belongs to which group”, they would get a learning task. For each trial, participants had to indicate as accurately as possible whether the Gabor patch belonged to the Arabic or Western group. Participants were asked to press one designated key if the symbol belonged to the one group (S), and another designated key if the symbol belonged to the other group (5 on the number-pad), with response keys counterbalanced between blocks. During the whole task, depending on the designated response keys, the top-left corner of the screen displayed 'Press S for Arabic (Western)’, and the top-right corner of the screen displayed 'Press 5 for Western (Arabic)’. Stimulus presentation times, feedback screens and the number of trials were identical to the learning task in Experiment 1.

After performing the task, participants were instructed to get the experimenter, who then instructed them on the b-CFS paradigm. The b-CFS task had the same properties and procedures as in Experiment 1, with the only exception being the stimuli: Gabor patches instead of geometric figures, and the addition of facial pictures. In the symbol-block, similar to Experiment 1, both Gabor patches were presented 34 times (17 times on both the right and left side of the fixation cross), leading to a total of 68 trials. In the faces-block, target faces were 34 male faces, selected from the Radboud Faces Database (Langner et al., 2010), of which 17 were Caucasian male, and 17 were Moroccan male. In line with Stein et al. (2014), the images were cropped by an oval-shape, deleting all non-facial features yet keeping variability in face size (see Figure 5). Each face was presented two times (one time on both the left and right side of the fixation cross), leading to a total of 68 trials.

After finishing the two b-CFS blocks, participants were thanked, paid, and debriefed.

a. b.

Figure 5. Example target faces presented in the b-CFS task: oval-cropped (a) Moroccan, or (b) Caucasian faces.

Data-analysis b-CFS

Again regarding the b-CFS task, reaction times (i.e., suppression durations) of correct trials were the main outcome measure of interest. In line with our pre-registered criteria, data of three participants (two Arabic, one Western) were excluded from the both the symbol- and faces-block analyses as they had more than 10% of the trials incorrect. Concerning the same criteria, data of another three participants (all Arabic) were excluded from the symbol-block only. Thus, symbol-block data was analyzed using data of 39 participants ($N_{\text{Western}} = 29$, $N_{\text{Arabic}} = 10$), whereas the faces-block data was analyzed using data of 42 participants ($N_{\text{Western}} = 30$, $N_{\text{Arabic}} = 12$). For the remaining participants, all incorrect trials were excluded from analyses (2.47% in the symbol-block; 2.18% in the faces-block). Finally, per block, reaction times deviating more than 3 *SD* from the average were regarded as outliers, and therefore excluded (0% in the symbol-block; 3.90% in the faces-block).

To test for effects of participant Ethnicity (Arabic x Western), Target Symbol (in-group associated x out-group associated), and Target Face (in-group x out-group) on the dependent variable suppression duration, a Bayesian repeated measures ANOVA was conducted for the two blocks separately in JASP (JASP Team, 2018). Bayes Factors (BF_{10}, or its inverse BF_{01}) were calculated comparing alternative models to the null model. Additionally, to test for the interaction effect, a Bayes Factor was calculated comparing the full model to the model excluding only the interaction effect (i.e., $BF_{inclusion}$). Similar to Experiment 1, because the (difference scores of) suppression durations were not normally distributed per group, all analyses were conducted using log-transformed suppression durations (although for interpreting purposes, only the non-transformed durations are reported).

Results Experiment 2

Symbol-block

To test whether both groups show prioritization effects for in-group (compared to out-group) associated symbols, of interest here is the main effect of Target Symbol. That is, whether in- versus out-group associated symbols differ in their suppression duration. The estimated Bayes Factor ($BF_{01} = 3.72$), comparing the fit of the data under the null hypothesis and the alternative hypothesis of an effect of Target Symbol, suggested that the data were 1:3.72 in favor of the null hypothesis, or rather around 4 times more likely to occur under a model not including the effect of Target Symbol. This is substantial evidence in favor of the null hypothesis without an effect of Target Symbol (Jeffreys, 1961). On average, in-group associated symbols ($M = 3.23$, $SD = 1.66$) did not yield different suppression durations than out-group associated symbols ($M = 3.23$, $SD = 1.52$).

See Table 2 for the estimated Bayes Factors comparing all alternative models to the null model. All the inversed estimated Bayes Factors (BF_{01}) ranged between 1 and 12, indicating evidence ranging from anecdotal to strong, in favor of the null hypothesis without any effects (Jeffreys, 1961). Thus, none of the alternative models were able to improve the null model.

Lastly, to test for the interaction effect, a Bayes Factor was calculated comparing the full model to the model excluding only the interaction effect. The estimated Bayes Factor ($BF_{inclusion}$ = 0.401), suggested that the data were 1/0.401 = 2.49 times more likely to occur under a model not including the Target Symbol x Ethnicity interaction effect, rather than a model with it. This is anecdotal evidence in favor of the model excluding the Target Symbol x Ethnicity interaction (Jeffreys, 1961).

Table 2

Model Comparison Symbol-Block

Models	P(M)	P(M|data)	BF_M	BF_{10}	BF_{01}	error %
Null model (incl. subject)	0.200	0.440	3.149	1.000	1.000	
Target Symbol	0.200	0.118	0.537	0.269	3.720	1.359
Ethnicity	0.200	0.319	1.878	0.725	1.379	0.856
Target Symbol + Ethnicity	0.200	0.087	0.380	0.197	5.076	2.579
Target Symbol + Ethnicity + Target Symbol ✻ Ethnicity	0.200	0.035	0.144	0.079	12.637	2.355

Note. All models include subject.

Faces-block

To quantify the evidence for our hypothesis that in-group faces yield shorter suppression durations compared to out-group faces, the log-transformed suppression durations were subjected to a Bayesian repeated measures ANOVA with participants' Ethnicity as a between-subjects factor (Arabic x Western), and Target Face (in-group x out-group) as a within-subjects factor. See table 3 for the estimated Bayes Factors comparing all alternative models to the null model. Of interest here are the models that include the main effect of Target Face, the main effect of Ethnicity, and the full model including both main effects and their interaction term.

Table 3

Model Comparison Faces-Block

Models	P(M)	P(M|data)	BF_M	BF_{10}	BF_{01}	error %
Null model (incl. subject)	0.200	1.033e -4	4.131e -4	1.000	1.000	
Target Face	0.200	0.004	0.015	36.768	0.027	1.102
Ethnicity	0.200	6.440e -5	2.576e -4	0.624	1.603	2.855
Target Face + Ethnicity	0.200	0.003	0.011	25.839	0.039	3.320
Target Face + Ethnicity + Target Face ✻ Ethnicity	0.200	0.993	599.107	9620.331	1.039e -4	4.188

Note. All models include subject.

On average, suppression times were shorter for in-group (M = 2.01, SD = 0.79) compared to out-group (M = 2.15, SD = 0.85) Target Faces. The estimated Bayes Factor (BF_{10} = 36.79) suggested that the data were around 36 times more likely to occur under a model including, rather than excluding, the main effect of Target Face. According to Jeffreys' (1961) classification scheme, this is very strong evidence for the effect of Target Face.

With regard to the main effect of Ethnicity, an estimated Bayes factor (BF_{01} = 1.60) suggested that the data were around 1.6 times more likely to occur under a model without including an effect of Ethnicity, rather than a model with it. This is anecdotal evidence (Jeffreys, 1961) for a model stating that Western and Arabic participants do not differ in their average suppression durations (Western: M = 2.03, SD = 1.24; Arabic: M = 2.01, SD = 1.22).

To test for the interaction effect, a Bayes Factor was calculated comparing the full model to the model excluding only the interaction effect. The estimated Bayes Factor ($BF_{inclusion}$ = 339.68), suggested that the data were around 340 times more likely to occur under a model including the Target Face x Ethnicity interaction effect, rather than a model without it. See Figure 6 for the mean suppression durations per Ethnicity and Target Face.

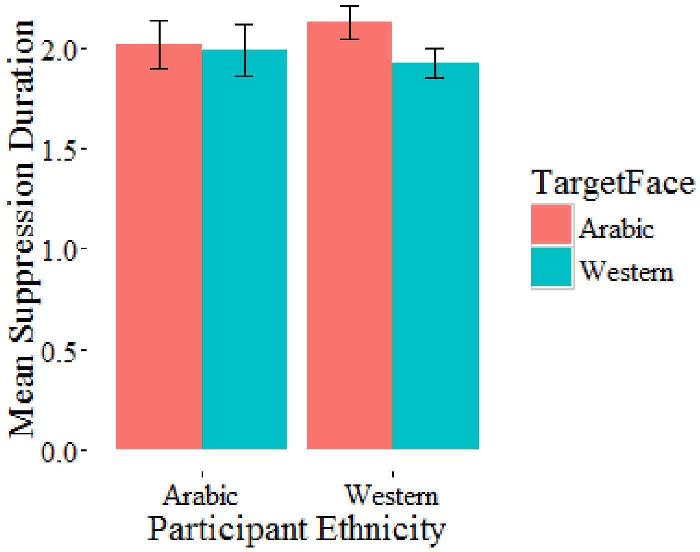

Figure 6. Mean suppression durations faces-block.

To further investigate this interaction effect, two additional Bayesian independent *t*-tests were performed to test for the simple slopes of Ethnicity (Arabic x Western) on Target Face (in-group x out-group). As we had specific hypotheses about the Western but not the Arabic subgroup, the log-transformed suppression durations of the Western participants were subjected to a one-sided Bayesian paired samples *t*-test (in-group target face < out-group target face), whereas the log-transformed suppression durations of the Arabic participants were subjected to a two-sided Bayesian paired samples *t*-test.

The estimated Bayes Factor (BF_{10} = 156224) suggested that, for the Western subgroup, the data were around 150000 times more likely to occur under the model including an effect of Target Face, rather than the model without it. This is decisive evidence in favor of the alternative hypothesis (Jeffreys, 1961), stating that Western participants have shorter suppression durations for in-group target faces (*M* = 1.99, *SD* = 0.86), than out-group target faces (*M* = 2.21, *SD* = 0.90). For the Arabic subgroup, the estimated Bayes Factor (BF_{01} = 2.12) suggested that the data were around two times more likely to occur under the null compared to the alternative model. This is weak or anecdotal evidence in favor of the null model (Jeffreys, 1961), stating that Arabic participants have equal suppression durations for in-group faces (*M* = 2.04, *SD* = 0.60), and out-group faces (*M* = 2.00, *SD* = 0.73). See Figure 7 for the individual subject data of both the Western and Arabic perceivers, additionally separated by the mean suppression duration for in- versus out-group target faces.

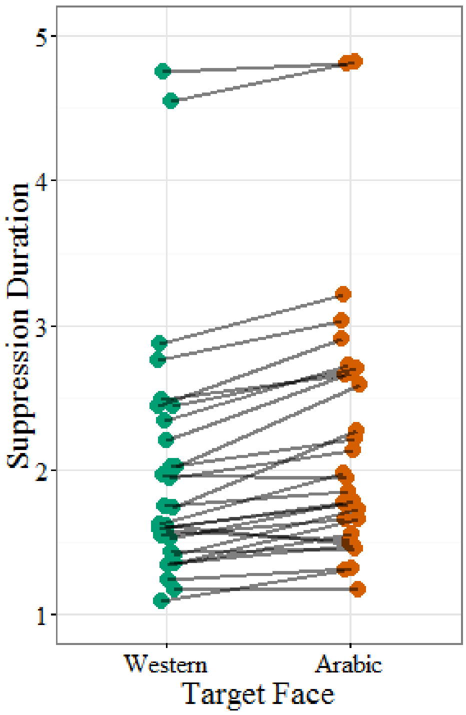

a. Western perceivers

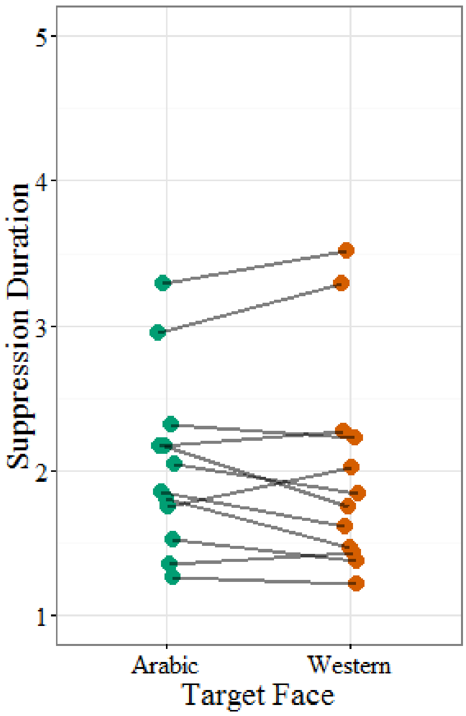

b. Arabic perceivers

Figure 7. Individual subject data for a) Western perceivers and b) Arabic perceivers. Most Western participants have shorter suppression durations for Western than for Arabic faces. Around half of the Arabic participants, show the same pattern as the Western participants.

To further investigate the Ethnicity x Target Face interaction, we tested again the simple slopes of Ethnicity (Arabic x Western), yet now on suppression durations regarding

Target Face (Arabic x Western), rather than Target Face (in-group x out-group). As such, we were able to investigate separately whether Arabic and Western participants differ in their detection times of Arabic faces, and whether they differ in their detection times of Western faces. As such, we conducted two additional Bayesian independent samples *t*-tests, with participants' Ethnicity (Arabic x Western) as a between-subjects factor, and with the (log-transformed) mean suppression durations for either the Arabic or the Western faces as the dependent variable. The estimated Bayes Factors for both Arabic faces ($BF_{01} = 2.75$) and Western faces ($BF_{01} = 3.04$), suggested that the data were more likely to occur under the null hypothesis, indicating that we found weak evidence for the model stating that Arabic and Western participants do not differ in suppression durations, neither for Arabic faces, nor for Western faces.

Exploratory Analyses

In an exploratory manner, we further analyzed the symbol-block suppression duration data to check for effects of low-level visual characteristics. As such, to test for the effect of participant Ethnicity and Gabor patch (horizontally oriented x vertically oriented) on suppression duration, a Bayesian repeated measures ANOVA was conducted. First, to investigate possible differences in suppression durations between Gabor patches, of interest here is the main effect of Gabor patch. The estimated Bayes Factor ($BF_{10} = 7.27$) suggested that the data were around seven times more likely to occur under a model stating that Gabor patches differ in their suppression durations. This is substantial evidence for this model (Jeffreys, 1961). The mean suppression duration for horizontally oriented Gabor patches was 3.39s (*SD* = 1.61), whereas the mean for vertically oriented Gabor patches was 3.07 (*SD* = 1.55). To test for the interaction effect, a Bayes Factor was calculated comparing the full model to the model excluding only the interaction effect. The estimated Bayes Factor ($BF_{inclusion} = 0.80$), suggested that the data were around 1/0.80 = 1.25 times more likely to occur under a model not including the participant Ethnicity x Gabor patch interaction. Thus, although suppression durations did differ between Gabor patches, this difference did not differ between Western and Arabic participants.

General Discussion

In two studies, we hypothesized that in-group information is preferentially selected for visual awareness as opposed to out-group information (i.e., Gobbini et al., 2013; Stein et al., 2014). Both in Experiment 1 and Experiment 2, when groups were associated with symbols, we did not find evidence for the effect of in- versus out-group information on prioritization for visual awareness. However, when pictures of in- versus out-group faces were presented (Experiment 2), we found evidence for an interaction effect, with Western perceivers prioritizing Western over Arabic faces in visual awareness. Conversely, however, we did not find such in-group prioritization effects for Arabic perceivers.

By using in- versus out-group associated *symbols*, we aimed to investigate whether only the self-relevance of an in-group (Tajfel & Turner, 1979) rather than the low-level features or mere increased familiarity of in-group faces (i.e., Gobbini et al., 2013), can drive in-group prioritization effects. While employing a design controlling for the latter two, we were not able to demonstrate any differences between in- and out-group information in prioritization for visual awareness. Therefore, as of now, we cannot conclude that self-relevance is the only factor driving in-group prioritization effects.

Instead, we may begin by exploring why, while using symbols, we did not find any in-group prioritization in the first place. One explanation for this null finding may pertain to the sensitivity of the b-CFS paradigm. That is, by associating group information with symbols, we implicitly stated that these higher-level group-symbol associations would be sufficient to influence suppression durations in b-CFS (i.e., Macrae et al., 2017). However, it remains to be seen whether stimuli can be analyzed up to the semantic level (i.e., further than just visual characteristics) when they are suppressed from awareness using b-CFS (Rabovsky, Stein, & Abdel Rahman, 2016). That is, although in general the affective or emotional value of a stimulus can attract attention (for a review, see Vuilleumier, 2005), it is unclear whether this can be observed when investigating detection mechanisms within b-CFS. Instead, when investigating influences of social group information on detection processes, visual input of these groups may be necessary in b-CFS. In line with this, several studies have shown that, affective person knowledge associated with pictures of faces, does not influence access to visual awareness, whereas those faces itself do (i.e., Stein, Grubb, Bertrand, Suh, & Verosky, 2017; Rabovsky, et al., 2016). In similar vein, whereas emotional facial expressions have been found to influence suppression durations in b-CFS (Yang et al., 2007), others have also suggested these effects to be driven by visual factors, rather than emotional valence per se (Gray, Adams, Hedger, Newton, & Garner, 2013; Stein & Sterzer, 2012). Thus, it is unclear whether the b-CFS paradigm is capable of reliably measuring differences in higher-level associations between stimuli, regardless of lower-level visual features. As the symbols only differed in their higher-level associations, this may explain why we did not find any in-group prioritization effects using symbols in b-CFS.

Simultaneously, this explanation relates to the fact that we *did* find prioritization effects when we used pictures of faces, rather than symbols, as group-associated information in b-CFS. That is, in Experiment 2, by using pictures of in- versus out-group faces, we conceptually replicated Stein et al. (2014) by showing an in-group prioritization effect for Western perceivers. We found evidence for an interaction effect between perceiver ethnicity and perceived face, indicating that facial detection mechanisms for Western and Arabic faces, are different for Western and Arabic perceivers. This may suggest that initial perception processes can be influenced by one's own social group (e.g., through self-relevance or familiarity), and are not necessarily dependent on the visual characteristics of the perceived face. Namely, even though low-level stimulus features can influence access to awareness (for an overview see Gayet et al., 2014), this explanation alone cannot account for our findings that Western and Arabic perceivers differ in how they perceive Western and Arabic faces. In line with this, several studies have shown that faces can influence awareness regardless of low-level stimulus features. More specifically, in a b-CFS paradigm, upright faces are detected faster than inverted faces (Jiang et al., 2007; Yang et al., 2007; Stein et al., 2011). As in these studies the facial stimuli are kept constant, it can be concluded that facial detection mechanisms are influenced by something other than just the low-level features of faces. That is, either facial familiarity (i.e., upright faces are more familiar than inverted faces), or the self-relevance of the social group the target person belongs to.

As the results regarding our Arabic subgroup were inconclusive, it seems fruitful to speculate about possible explanations for these outcomes. First, it is important to keep in mind that the sample size of the Arabic group was small (N = 12 for the faces block). Even though Stein et al. (2014) had a similar sample size (N = 14), increasing the sample size in future studies could help in gathering stronger evidence in favor of either the null or the alternative hypothesis. Second, we recruited participants from both the Western and Arabic in-group, operating under the assumption that the Arabic group would be more relevant to the Arabic perceivers' self, and that an Arabic perceiver would be more familiar with Arabic as opposed to Western faces. However, to be included in the Arabic in-group, participants were to have at least one parent that was born in the MENA-region (Middle East – North Africa, including Turkey), resulting in a sample of which more than half of the participants were born in the Netherlands themselves. Growing up in the Netherlands may have caused Arabic participants to be more familiar with Western faces than Western participants are with Arabic faces, possibly leading to indistinguishable face-prioritization effects. Additionally, for Arabic participants, living in the Netherlands may have caused both the Western and Arabic group to be similarly relevant to the self. Furthermore, the MENA-region is a very large region consisting of many different countries, ranging over two continents, and consisting of people coming from many different origins, cultures and ethnicities. Hence, although done for practical reasons, the operationalization of the ‘Arabic’ group may have been too broad, and instead, people might be more inclined to identify with a country-, or ethnicity-specific in-group (e.g., ‘Moroccans’).

To account for these possibilities, future studies could operationalize the groups more concisely by, for example, conducting the study *in* one of the MENA-countries, and thereby include participants that were born in, and identify with, the country of interest. In doing so, one can simultaneously account for in-group familiarity effects, as then it can be assumed more safely that the in-group is indeed more familiar with its’ in-group faces (compared to out-group faces). Taking this a step further, an interesting avenue for future research would be to investigate these familiarity effects, by including for example participants that were born in and still live in the Netherlands and Morocco, yet to also include Dutch participants that now live in Morocco, and Moroccan participants that now live in the Netherlands. As such, all groups are fully represented, enabling researchers to systematically disentangle effects of facial familiarity on facial detection. Additionally, to further investigate the possible effects of self-relevance on in-group facial detection, future studies could include a measure of group-identification. That is, if the self-relevance of an in-group can indeed influence perceptual detection mechanisms, then more strongly group-identifying individuals (i.e., those to whom the group is more relevant to the self) should show stronger in-group prioritization effects.

Furthermore, in future studies, the ecological validity of our work could be extended. That is, in the current studies, we had participants look at a computer screen while wearing prism glasses to place participants' visual fields on top of one another. However, this method did not always yield the desired results, in that some participants were still unable to see one square in the b-CFS task (meaning that we as researchers are dependent on participants to report on their ability to see one square). Additionally, the images presented are small, two-dimensional images. Very recently, however, a new version of the b-CFS paradigm was developed - termed ‘real-life CFS’ - in which augmented reality goggles are used to present CFS patterns to one eye, while leaving the other eye exposed to the real world (Korisky, Hirschhorn, Mudrik, 2018). This method allows researchers to exercise more experimental control over the paradigm, and simultaneously opens up possibilities to increase the ecological validity of b-CFS. ‘Real-life CFS’ may open up exciting new avenues by getting us one step closer to investigating stimuli we encounter in real, everyday life.

Implications

Attentional biases based on social groups can considerably impact interpersonal interactions and intergroup attitudes (Kawakami et al., 2018). For example, preferential attention can influence approach motivations (Kawakami et al., 2014) or discriminatory behavior (Granot, Balcetis, Schneider, & Tyles, 2014). The current research contributes to our understanding of preferential in-group attention effects, by indicating that perhaps these biases originate when we prioritize on what (not) to attend to. Simultaneously, this may have implications for factors underlying processes such as in-group favoritism (Tajfel, 1982). For that reason, it would be worthwhile to further investigate how- and whether initial face detection relates to subsequent behavior. That is, does initial face detection have predictive value for other (behavioral) processes? For example, is preferential treatment of an in-group member, or detrimental treatment of an out-group member, based on whether or not we initially detect their face?

Overall, the current experiments pioneered in trying to disentangle effects of low-level stimulus features, in-group familiarity, and in-group self-relevance on prioritization for visual awareness. At least for now, it seems that self-relevance alone, is not capable of influencing prioritization processes. Albeit the fact that this line of research studies relatively core principles and possibilities for direct applications of findings are beyond the present research, the importance for society still stands. The question whether social salience can influence perceptual processes is a fundamental one, aiding us in understanding how people are able to successfully adapt to complex and dynamic situations and thereby optimize self-survival. As well as having implications for understanding perceptual processes, the line of research also has more general implications for understanding the factors underlying in-group favoritism. Given the problems in society made up from in- versus out-group processes (i.e., stereotypes, prejudice, discrimination), it is necessary and legitimate to investigate the underlying principles maintaining us-and-them behavior. Pinpointing factors underlying these social phenomena is important as to eventually understand their causes and thereby obliterate them or prevent them from happening in the first place.

References

Allport, G.W. (1954). The nature of prejudice. Oxford, England: Addison-Wesley.

Amodio, D.M., & Devine, P.G. (2006). Stereotyping and evaluation in implicit race bias: Evidence for independent constructs and unique effects on behavior. *Journal of Personality and Social Psychology, 91*(4), 652-661. DOI: 10.1037/0022-3514.91.4.652

Ashmore, R. D., Jussim, L., & Wilder, D. (Eds.). (2001). *Rutgers series on self and social identity; Vol. 3. Social identity, intergroup conflict, and conflict reduction.* New York, NY, US: Oxford University Press.

Bar-Haim, Y., Ziv, T., Lamy, D., & Hodes, R. M. (2006). Nature and nurture in own-race face processing. *Psychological Science*, *17*, 159–163. DOI:10.1111/j.14679280.2006.01679.x

Bean, M. G., Slaten, D. G., Horton, W. S., Murphy, M. C., Todd, A. R., et al. (2012). Prejudice concerns and race-based attentional bias: New evidence from eyetracking. *Social Psychological and Personality Science, 3*, 722–729. DOI:10.1177/1948550612436983

Berry, D.S., & McArthur, L.Z. (1986). Perceiving character in faces: the impact of age-related craniofacial changes on social perception. *Psychological Bulletin, 100,* 3-18. DOI: 10.1037//0033-2909.100.1.3

Bijlstra, G., Holland, R.W., & Wigboldus, D.H.J. (2010). The social face of emotion recognition: Evaluations versus stereotypes. *Journal of Experimental Social Psychology, 46*, 657-663. DOI:10.1016/j.jesp.2010.03.006

Brewer, M.B. (1979). In-group bias in the minimal intergroup situation: A cognitive-motivational analysis. *Psychological Bulletin, 86*(2), 307-324. DOI:

Brewer, M.B. (1999). The psychology of prejudice: Ingroup love or outgroup hate? *Journal of Social Issues, 55*(3), 429-444. DOI:10.1111/0022-4537.00126

Gayet, S., Van der Stigchel, S., & Paffen, C.L.E. (2014). Breaking continuous flash suppression: Competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology, 5*, 460. DOI:10.3389/fpsyg.2014.00460

Gobbini, I., Gors, J.D., Halchenko, Y.O., Rogers, C., Guntupalli, J.S., Hughes, H., & Cipolli, C. (2013). Prioritized detection of personally familiar faces. *PLoS ONE, 8*(6): e66620. DOI: 10.1371/journal.pone.0066620

Granot, Y., Balcetis, E., Schneider, K. E., & Tyler, T. R. (2014). Justice is not blind: Visual attention exaggerates effects of group identification on legal punishment. *Journal of Experimental Psychology: General, 143*, 2196–2208. DOI:10.1037/a0037893

Gray, K.L., Adams, W.J., Hedger, N., Newton, K.E., & Garner, M. (2013). Faces and awareness: low-level, not emotional factors determine perceptual dominance. *Emotion, 13*, 537-544. DOI:10.1037/a0031403

He, Y., Ebner, N. C., & Johnson, M. K. (2011). What predicts the own-age bias in face recognition memory? *Social Cognition*, 29(1), 97-109. DOI10.1521/soco.2011.29.1.97

JASP Team (2018). JASP (Version 0.9)[Computer software].

Jeffreys, H. (1961). *Theory of probability*. Oxford, UK: Oxford University Press.

Jiang, Y., Costello, P., & He, S. (2007). Processing of invisible stimuli: Advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological Science, 18*, 349-355. DOI: 10.1111/j.1467-9280.2007.01902.x

Kawakami, K., Friesen, J., & Vingilis-Jaremko, L. (2018). Visual attention to members of own and other groups: Preferences, determinants, and consequences. *Social and Personality Psychology Compass, 12*(4). DOI: 10.1111/spc3.12380

Korisky, U., Hirschhorn, R., & Mudrik, L. (2018). Real-life continuous flash suppression - Rendering the real world unconscious using augmented reality goggles. Pre-print: 10.31234/osf.io/y3fcz

Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D.H.J., Hawk, S.K., & Van Knippenberg, A. (2010). Presentation and validation of the Radboud faces database. *Cognition and emotion, 24*(8), 1377-1388. DOI: 10.1080/02699930903485076

Lovén, J., Rehnman, J., Wiens, S., Lindholm, T., Peira, N., et al. (2012). Who are you looking at? The influence of face gender on visual attention and memory for own- and other-race faces. *Memory*, *20,* 321–331. DOI: 10.1080/09658211.2012.658064

Macrae, C.N., Visokomogilski, A., Golubickis, M., Cunningham, W.A., & Sahraie, A. (2017). Self-relevance prioritizes access to visual awareness. *Journal of Experimental Psychology: Human Perception and Performance, 43*, 438-443. DOI:10.1037/xhp0000361

Maner, J. K., & Miller, S. L. (2013). Adaptive attentional attunement: Perceptions of danger and attention to outgroup men. *Social Cognition*, *31*, 733–744. DOI:10.1521/soco.2013.31.6.733

Otten, S., & Bar-Tal, Y. (2002). Self-anchoring in the minimal group paradigm: The impact of need and ability to achieve cognitive structure. *Group Processes & Intergroup Relations, 5*(4), 167-284. DOI: 10.1177/1368430202005004001

Otten, S. & Moskowitz, G.B. (2000). Evidence for implicit evaluative in-group bias: Affect-biased spontaneous trait inference in a minimal group paradigm. *Journal of Experimental Social Psychology, 36*, 77-89. DOI: 10.1006/jesp.1999.1399

Otten, S. & Wentura, D. (1999). About the impact of automaticity in the minimal group paradigm: Evidence from affective priming tasks. *European Journal of Social*

Psychology, 29, 1049-1071. DOI: 10.1002/(SICI)1099-0992(199912)29:8<1049::AID-EJSP985>3.0.CO;2-Q

Peirce, J.W. (2007). PsychoPy - Psychophysics software in Python. *Journal of Neuroscience Methods, 162*(1-2):8-13. DOI: 10.1016/j.jneumeth.2006.11.017

Rabovsky, M., Stein, T., & Abdel Rahman, R. (2016). Access to awareness for faces during continuous flash suppression is not modulated by affective knowledge. *PLoS ONE, 11*(4): e0150931. DOI: 10.1371/journal.pone.0150931

R Core Team (2016). R: A language and environment for statistical computing [Computer Software]. Vienna, Austria: Foundation for Statistical Computing.

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review, 16*, 225–237. DOI:10.3758/PBR.16.2.225

Stein, T., End, A., & Sterzer, P. (2014). Own-race and own-age biases facilitate visual awareness of faces under interocular suppression. *Frontiers in Human Neuroscience, 8*, 582. DOI: 10.3389/fnhum.2014.00582

Stein, T., Grubb, C., Bertrand, M., Suh, S. M., & Verosky, S. C. (2017). No impact of affective person knowledge on visual awareness: Evidence from binocular rivalry and continuous flash suppression. *Emotion, 17*(8), 1199-1207. DOI: 10.1037/emo0000305

Stein, T., Hebart, M. N., & Sterzer, P. (2011). Breaking continuous flash suppression: A new measure of unconscious processing during interocular suppression? *Frontiers in Human Neuroscience, 5*, 167. DOI: 10.3389/fnhum.2011.00167

Stein, T., Siebold, A., & van Zoest, W. (2016). Testing the idea of privileged awareness of self-relevant information. *Journal of Experimental Psychology: Human Perception and Performance, 42,* 303-307. Advance online publication. DOI:10.1037/xhp0000197

Stein, T., & Sterzer, P. (2012). Not just another face in the crowd: detecting emotional schematic faces during continuous flash suppression. *Emotion, 12,* 988–996. DOI: 10.1037/a0026944

Stein, T., & Sterzer, P. (2014). Unconscious processing under interocular suppression: getting the right measure. *Frontiers in Psychology, 5*, 387, DOI: 10.3389/fpsyg.2014.00387

Stein, T., Sterzer, P., & Peelen, M. V. (2012). Privileged detection of conspecifics: evidence from inversion effects during continuous flash suppression. *Cognition, 125*, 64–79. DOI: 10.1016/j.cognition.2012.06.005

Sui, J., He, X., & Humphreys, G. W. (2012). Perceptual effects of social salience: Evidence from self-prioritization effects on perceptual matching. *Journal of Experimental Psychology: Human Perception and Performance, 38*, 1105–1117. DOI:10.1037/a0029792

Tajfel, H. (1970). Experiments in intergroup discrimination. *Scientific American, 223,* 96-102. DOI: 10.1038/scientificamerican1170-96

Tajfel, H. (1982). Social Psychology of Intergroup Relations. *Annual Review of Psychology, 33*, 1-39. DOI:10.1146/annurev.ps.33.020182.000245.

Tajfel, H., Billig, M.G., Bundy, R.P., & Flament, C. (1971). Social categorization and intergroup behavior. *European Journal of Social Psychology, 1*, 149-178. DOI: 10.1002/ejsp.2420010202

Tajfel, H., & Turner, J.C. (1979). An integrative theory of intergroup conflict. In W.G. Austin & S. Worchel (Eds.), The social psychology of intergroup relations (pp. 33-47). Monterey, CA: Brooks/Cole

Trawalter, S., Todd, A. R., Baird, A. A., & Richeson, J. A. (2008). Attending to threat: Race-based patterns of selective attention. *Journal of Experimental Social Psychology*, *44,* 1322–1327. DOI:10.1016/j.jesp.2008.03.006

Tsuchiya, N. & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience, 8*, 1096-1101. DOI:10.1038/nn1500

Van Bavel, J.J., Packer, D.J., Cunningham, W.A. (2011). Modulation of the fusiform face area following minimal exposure to motivationally relevant faces: Evidence of in-group enhancement (not out-group disregard). *Journal of Cognitive Neuroscience, 23*(11), 3343-3354. DOI: 10.1162/jocn_a_00016

Vuilleumier, P., Armony, J.L., Driver, J., & Dolan, R.J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature Neuroscience, 6*, 624–631. DOI: 10.1038/nn1057

Yang, E., & Blake, R. (2012). Deconstructing continuous flash suppression. *Journal of Vision, 12*(3), 1–14 (article no. 8). DOI: 10.1167/12.3.8

Yang, E., Zald, D.H., Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion, 7,* 882–886. DOI: 10.1037/1528-3542.7.4.882

Appendix

Frequentist Paired Samples *t*-test

A paired-samples *t*-test was conducted to compare suppression durations of in- and out-group associated symbols. This test showed no significant difference between log-transformed durations for in- versus out-group associated symbols, $t(47) = 0.48$, $p = 0.6$

Linear Mixed Models

A linear mixed effects analysis of the relationship between suppression durations and group-associated symbols was performed. In R, the *lmer* function of the lme4 package was used (version 1.1.12; Bates, Maechler, Bolker, & Walker, 2015). This method entails estimating both fixed and random effects. To avoid inflated Type 1 errors, Barr, Levy, Scheepers, and Tily (2013) advise to use a maximal random-effects structure when possible. That is, a per-participant random adjustment to the fixed intercept ("random intercept"), as well as per-participant random adjustments to the slopes of predictors varying within-subject ("random slopes") were modeled. In addition, all possible random correlation terms among the random effects were included in the model.

The maximal model included the fixed effect of group (in- versus out-group). Keeping a maximal random-effects structure as suggested by Barr et al. (2013), a per-participant adjustment to the fixed intercept (i.e., a random intercept of participant) was included, as well as per-participant random adjustments to the group-slopes (i.e., random slopes for group). Additionally, a random intercept of trial was included. Including all this, as well as all possible random correlations terms among the random effects, the main maximal model in R (R Core Team, 2016) looked as follows:

Suppression Duration ~ Group + (1 + Group | Participant) + (1 | Trial)

For factors, in this case group, sum-to-zero coding was used, with the in-group condition coded as -1. Following Judd, Westfall & Kenny (2012), to determine *p*-values for the overall effects, we computed Type III conditional *F* tests with Kenward-Roger approximation for degrees of freedom using the mixed function of the *afex* package (Singmann, Bolker, Westfall, & Aust, 2016), which in turn calls on the KRmodcomp function of the *pbkrtest* package (Halekoh & Højsgaard, 2014). According to Judd et al. (2012) this procedure results in a better approximation of the *F* distribution. Additionally, yielded *p*-values are close to nominal values, accounting for small sample sizes (as is the case in the current study).

Model diagnostics of the main maximal model were examined and showed reasons for concern regarding influential cases: suppression durations of two participants could be influential. Therefore, we ran the same model again, excluding data of these participants. In doing so, the model seemed robust and the pattern of results did not change notably. Therefore, the reported results on the main model are based on the analysis including all 48 participants.

The linear mixed-effects model yielded no significant differences between in- versus out-group suppression durations (Estimate = -0.005, *se* = 0.01; $F(1, 46.8) = 0.25$, $p = .621$). That is, no evidence was found that the suppression duration differed between in- versus out-group associated symbols.

References Appendix

Barr, D.J., Levy, R., Scheepers, C., & Tily, H.J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language, 68*(3), 255-278.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1-48.

Halekoh, U., & Højsgaard, S. (2014). A Kenward-Roger approximation and parametric bootstrap methods for tests in linear mixed models - The R package pbkrtest. *Journal of Statistical Software, 59*(9), 1-30.

Judd, C.M., Westfall, J., & Kenny, D.A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology, 103*(1), 54-69.

R Core Team (2016). R: A language and environment for statistical computing[Computer Software]. Vienna, Austria: Foundation for Statistical Computing.

Singman, H., Bolker, B. Westfall, & Aust, F. (2016). Afex: Analysis of factoral experiments (version 0.16-1) [R package].