Personality and Social
Psychology Bulletin
2015, Vol. 41(6) 769–778

DOI: 10.1177/0146167215579053
pspb.sagepub.com
SAGE

Semantic Information Influences Race Categorization From Faces

Konstantin O. Tskhay[1] and Nicholas O. Rule[1]

Abstract

It is well established that low-level visual features affect person categorization in a bottom-up fashion. Few studies have examined top-down influences, however, and have largely focused on how information recalled from memory or from motivation influences categorization. Here, we investigated how race categorizations are affected by the context in which targets are perceived by manipulating semantic information associated with the faces being categorized. We found that presenting faces that systematically varied in racial ambiguity with race-congruent (vs. incongruent) semantic labels shifted the threshold at which perceivers distinguished between racial groups. The semantic information offered by the labels therefore appeared to influence the categorization of race. These findings suggest that semantic information creates a context for the interpretation of perceptual cues during social categorization, highlighting an active role of top-down information in race perception.

Keywords

categorization, person perception, person construal, race, social cognition

Received July 15, 2014; revision accepted March 4, 2015

Despite great variability in skin tone and other distinctive features, people tend to think about racial groups as discrete entities with little perceptual overlap (Macrae & Bodenhausen, 2000; Maddox, 2004). Evidence suggests that individuals may integrate various pieces of information to arrive at these categorizations, however. Basic visual information (e.g., racially phenotypic facial features) and higher order cognitions (e.g., perceivers' motivations) can both influence how people categorize others into social groups (e.g., Blair, Judd, Sadler, & Jenkins, 2002; MacLin & Malpass, 2001, 2003; Rule, Garrett, & Ambady, 2010). Although much is understood about how perceptual cues contribute to social categorization bottom-up, less research has explored the influence of higher level information top-down. Unlike previous research that has studied top-down effects by manipulating individuals' recall of previously presented information (e.g., Eberhardt, Dasgupta, & Banaszynski, 2003) and internal states (e.g., motivations; Sacco & Hugenberg, 2012), we considered how semantic information might organize perception contemporaneously.

Research in cognitive psychology suggests that people's expectations, attention, and memory affect object perception and categorization (see Bar et al., 2006; Gilbert & Sigman, 2007, for reviews). Studies have also shown that semantic information can affect how people categorize and interpret stimuli. Tajfel and Wilkes (1963), for example, observed that participants perceived short lines labeled as *short* to be significantly shorter than long lines labeled as *long* when they viewed the lines with category labels (e.g., Categories A and B) compared with when they viewed the lines without labels. Presenting sets of short and long lines categorically therefore affected how perceivers saw them, suggesting that semantic information can organize and alter the perception.

The examination of top-down influences on categorization is not limited to object perception, however; social perception is also susceptible to the influence of higher order cognition (Freeman & Ambady, 2011). For example, researchers have demonstrated that people perceive others' race and sexual orientation in line with information that they previously encoded into memory (Eberhardt et al., 2003; Pauker et al., 2009; Rule, Tskhay, Freeman, & Ambady, 2014). Furthermore, Sacco and colleagues found that perceivers' motives to cooperate or compete affected their categorization of emotional expressions and that ostracized individuals showed a greater ability to discriminate between happy and angry facial expressions (Sacco & Hugenberg, 2012; Sacco, Wirth, Hugenberg, Chen, & Williams, 2011). Thus, individuals' memories, motivations, and internal

[1]University of Toronto, Ontario, Canada

Corresponding Author:

Konstantin O. Tskhay, Department of Psychology, University of Toronto, 100 St. George Street, Toronto, Ontario, Canada M5S 3G3.
Email: konstantin.tskhay@mail.utoronto.ca

feelings can all shape the way in which social stimuli are perceived and categorized, similar to object perception.

Further evidence suggests that semantic information might also influence person perception and categorization. For example, MacLin and Malpass (2001, 2003) demonstrated that racially ambiguous faces appeared "Black" when presented with stereotypically Black hairstyles but that the identical faces appeared "Hispanic" when paired with stereotypically Hispanic hairstyles. These studies suggest that changes to the targets' hairstyles (a basic visual feature of the stimuli) activated race-related concepts (an organizing semantic principle of the stimuli) to alter perception top-down. But the activation of racial categories in this paradigm is merely assumed. Direct evidence for whether semantic information modifies person categorization contemporaneously with the perception of basic visual features is therefore absent. Here, we sought to bridge this gap by connecting overt semantic representations of social categories to the perception of basic visual features in race categorization. We thus measured how semantic information (i.e., social category response options) changed participants' categorizations of racially ambiguous faces.

To manipulate semantic information, we focused on the well-established social-psychological finding that White is the default racial group in the minds of North Americans (Chen & Hamilton, 2012; Ho, Sidanius, Levin, & Banaji, 2011; Smith & Zárate, 1992). Specifically, Americans tend to perceive racially ambiguous faces as "Black" sooner in the context of White–Black categorization than they do in the context of two minority groups (e.g., Hispanic and Black). This effect tends to be robust and largely independent of the perceiver's own race. Building on this premise, we generated continua between prototypical Hispanic and White faces morphed with the same prototypical Black face (i.e., continua stretched from Hispanic to Black and White to Black, respectively). Independent groups of participants then viewed these continua and categorized the faces as either "White" and "Black" or "Hispanic" and "Black" in a 2 × 2 between-subjects design. Thus, participants' response options constrained the social categories that would be applicable for categorizing the targets, thereby creating a semantically informed context for judging the targets' race.

Consistent with previous findings illustrating White as the racial default group, we expected that participants would categorize faces as Black earlier in the White–Black continuum when "White" was a response option than when "Hispanic" was a response option (Chen & Hamilton, 2012; Ho et al., 2011; Smith & Zárate, 1992). Effects resulting from this simple change of category labels would help to illustrate how semantic information can alter race categorization.

Critically, we also swapped the category options, asking participants either to categorize the White–Black continuum faces as "Hispanic" and "Black" or to categorize the Hispanic–Black continuum faces as "White" and "Black." Under these conditions, we hypothesized that the semantic information provided by the category labels would again structure participants' categorizations. Specifically, we expected that participants categorizing faces from the Hispanic–Black continuum would categorize them as if they were viewing the White–Black continuum (categorizing the faces as Black sooner) compared with participants using the veridical labels "Hispanic" and "Black" to categorize the faces (control). Likewise, we expected that participants categorizing faces from the White–Black continuum as "Hispanic" and "Black" would categorize them as if they were viewing the Hispanic–Black continuum (i.e., categorizing the faces as Black later) than participants using the veridical labels "White" and "Black" (control).

Importantly, because we expected that the participants would assume that they are viewing the faces from the continuum suggested by the response options, we hypothesized that the low-level visual information from the faces would not interact with the response options during categorizations. In other words, similar to MacLin and Malpass's (2001) findings, we expected that higher order cognition would dominate the categorization of racially ambiguous faces. We tested these hypotheses in Experiments 1 and 2.

Finally, in Experiment 3, we wanted to ensure that the participants' categorizations in Experiments 1 and 2 were not simply the result of being "forced" to categorize the faces according to non-veridical labels (Chen & Hamilton, 2012) but, rather, were a product of the changes to the semantic information scaffolding their categorization. To test this, we asked the participants to categorize the faces from the White–Black continuum as "White" versus "Black," "White" versus "Not White," or "Black" versus "Not Black." In conditions with an unspecified contrast (i.e., "White" vs. "Not White" and "Black" vs. "Not Black"), we expected that the participants would be more conservative about assigning racially ambiguous faces to the specified categories ("White" and "Black") because the flexibility of the opponent option ("Not *X*") allows for greater accommodation of faces whose category membership is not obviously discrete. Specifically, the category "Not White" allows perceivers to reject not only Black faces but all faces that are not clearly White, and the category "Not Black" allows perceivers to reject not only White faces but all faces that are not clearly Black. Critically, because White is the default category and is therefore protected in the perceiver's mind (e.g., Ho et al., 2011), we expected that perceivers would be more conservative about categorizing racially ambiguous faces as White than they would be about categorizing racially ambiguous faces as Black. This conservative threshold for the default category would also apply to the third condition in which participants are constrained to choosing between the two specified labels "Black" and "White" (as in Experiment 1). More important, however, we expected that the threshold for White categorizations would be intermediate in this condition compared to the threshold in the other two, as all of the previously rejected racially ambiguous faces would have to be categorized as

Figure 1. Prototypes of the racial groups used in the experiments.

either "Black" or "White," thereby numerically increasing the acceptance levels within each category. Thus, the anticipated results of Experiment 3 would demonstrate that perception is indeed affected by semantic information, independent of constraints placed on the participants via labels.

Experiment 1

Method

Participants. We recruited 150 American participants (66 female; age range = 18-69 years; 112 White, 9 Black, 7 Hispanic/Latino, 15 Asian, 7 Other) from Amazon's Mechanical Turk (MTurk) to participate in a person perception study in exchange for financial compensation. We calculated the sample size needed to achieve at least 80% statistical power using the effect sizes reported in Ho et al. (2011).

Stimuli. To generate prototypes of White, Black, and Hispanic men, we averaged 12 neutral-expression faces of each race from an in-house database (Figure 1). We then morphed the prototype faces for each race in pairs (morph points: $M = 85.60$, $SD = 8.21$; see Schweinberger, Burton, & Kelly, 1999; Walker & Tanaka, 2003) to generate stimuli ranging from a 100% contribution from one prototype to a 100% contribution from the other. This produced three continua: White–Black, Hispanic–Black, and White–Hispanic (manipulation check condition). The morphs advanced in 5% increments to produce 21 faces within each continuum such that any given stimulus in the White–Black continuum, for example, would contain a% contribution from the White prototype and $(100 - a)$% contribution from the Black prototype. In other words, we linearly varied the degree to which the parent faces were represented in the face morphs in 21 steps with the 11th face in the continuum containing equal contributions from both parents.

Procedure. Approximately equal numbers of participants categorized the faces from the White–Black, Hispanic–Black, or White–Hispanic continuum using response labels that were congruent with the parent categories and presented below the stimuli. Additional participants categorized the faces from the Hispanic–Black or White–Black continuum as "White" and "Black" or "Hispanic" and "Black," respectively; that is, the labels for the White and Hispanic options were switched across these two conditions. Five participants' data could not be analyzed because they gave uniform responses to all of the stimuli (final $N = 145$). Each participant dichotomously categorized the race of each of the 21 morphs in their assigned continuum 6 times in random order (total trials = 126), which was necessary to ascertain the threshold at which participants' categorizations switched between the two response options (e.g., from White to Black; see below).

Results

Analytic strategy. We used the psychometric threshold as an index of the point along the continuum at which each participant began to categorize the faces as Black (see also Krosch & Amodio, 2014, for similar estimation). Thus, we fit sigmoidal psychometric functions modeling categorical perception as a function of incremental changes in the properties of the stimuli for each participant using logistic regression (Freeman, Rule, Adams, & Ambady, 2010; Macmillan & Creelman, 2005; Żychaluk & Foster, 2009). Using the predicted probabilities, we extracted the categorization threshold—the location on the morph continuum where the participants' perceptions and responses rapidly changed from one category to another. Specifically, we estimated a logistic regression equation whereby the predicted scores were expressed in logit units as a function of intensity (percent contribution from the Black prototype) for each participant. We then solved for the stimulus intensity level at a categorization probability level of p = .50 (logit = 0) and used this (threshold) value as our main dependent variable.

In other words, we estimated the stimulus level on the continuum at which each participant categorized the face to either category exactly 50% of the time. This measure is a more precise index of the perceptual boundary between categories than a mere approximation to the face that was actually presented to the participants would be (Żychaluk & Foster, 2009). Although the threshold marks a hypothetical face, it constitutes a meaningful boundary such that any face that appears before or after this face on the continuum should be categorized to one of the available categories with probabilities that are greater than 50%. We used these threshold values as the main dependent variable in a between-subjects ANOVA with continuum and label as factors. All simple effects analyses are accompanied by a Pearson product–moment correlation (r) as a measure of the effect size; we constructed the 95% confidence intervals (CIs) around the effect sizes using meta-analytic procedures (Rosenthal & Rosnow, 2008).

Manipulation check (White–Hispanic distinction). Prior to analysis, we verified our experimental manipulation by comparing how rapidly categorization changed across the conditions in which the race and label were congruent. We hypothesized that the White–Hispanic continuum would be significantly more ambiguous than the White–Black and Hispanic–Black continua. We therefore expected to see a lower rate of perceptual change (i.e., slope) along the White–Hispanic continuum compared with the other two continua, as the slope may serve as an index of ambiguity, confusion, or response variability in this context (Żychaluk & Foster, 2009).

As expected, we found a significant main effect of continuum on slope: $F(2, 85) = 23.89$, $p < .001$, $\eta^2 = .36$. Simple effects analyses confirmed that perceptual shifts in the White–Hispanic continuum ($M = .02$, $SD = .02$) occurred more slowly than in the White–Black, $M = .06$, $SD = .02$, $t(56) = 7.36$, $p < .001$, $r = .70$, 95% CI = [.54, .81], and Hispanic–Black, $M = .05$, $SD = .03$, $t(58) = 5.04$, $p < .001$, $r = .55$, 95% CI = [.34, .71], continua, which did not differ, $t(56) = 1.40$, $p = .17$, $r = .18$, 95% CI = [− .08, .42]. Thus, race was ambiguous between the White and Hispanic prototypes.

Main analysis. To test our hypothesis that perceptions of categories are affected by contemporaneously processed semantic information, we submitted the thresholds of the continua (i.e., the amount of perceptual information from the Black prototype needed to categorize the stimulus face as Black) to a 2 (continuum: White−Black, Hispanic−Black) × 2 (label: White−Black, Hispanic−Black) between-subjects ANOVA. This allowed us to examine the influence of the semantic information while controlling the influence of lower level perceptual features.

Consistent with previous findings suggesting that White is the default racial group (e.g., Ho et al., 2011), we observed a main effect of continuum: Participants categorizing faces from the White–Black continuum judged them as Black sooner than participants categorizing faces from the Hispanic–Black continuum, $F(1, 111) = 6.38$, $p = .013$, $\eta^2 = .04$. In addition, we observed a main effect of label: Participants categorizing the faces as “White” and “Black” (regardless of whether the faces were from the White–Black or Hispanic–Black continua) required less contribution from the Black prototype to categorize a face as Black compared with participants categorizing the faces as “Hispanic” and “Black,” $F(1, 111) = 36.62$, $p < .001$, $\eta^2 = .24$. Thus, the response labels influenced basic perception independent of visual information, supporting our main hypothesis. No significant Condition × Label interaction emerged: $F(1, 111) < 1$, $p = .45$, $\eta^2 < .001$.

Analyses within each continuum confirmed these effects. Participants categorizing faces from the White–Black continuum as “White” and “Black” ($M = 41.82$, $SD = 7.07$) required less contribution from the Black prototype to categorize a face as “Black” than did participants categorizing the identical faces using the “Hispanic” and “Black” response options: $M = 52.46$, $SD = 7.00$; $t(54) = 5.66$, $p < .001$, $r = .61$, 95% CI = [.41, .75] (Figure 2A). Similarly, participants categorizing faces from the Hispanic–Black continuum as “White” and “Black” ($M = 37.94$, $SD = 12.40$) required less contribution from the Black prototype to categorize them as “Black” than participants categorizing the identical faces using the “Hispanic” and “Black” response options: $M = 46.51$, $SD = 8.67$; $t(57) = 3.32$, $p = .001$, $r = .40$, 95% CI = [.16, .59] (Figure 2B).

Discussion

These results suggest that semantic information about social categories available to perceivers may influence how they categorize race. First, although participants largely did not distinguish White and Hispanic faces as different, they were more conservative about categorizing potentially White faces as Black than they were about categorizing potentially Hispanic faces as Black. This replicates past work showing that White is the default racial group in North America (e.g., Ho et al., 2011). More important, the thresholds at which participants perceived the White–Black and Hispanic–Black group boundaries were guided by semantic information embedded in their category response options. When the category labels were changed such that White–Black morphs were to be categorized as either “Hispanic” or “Black” and Hispanic–Black morphs were to be categorized as either “White” or “Black,” participants’ responses repeated the pattern of the labeled—not actual—groups. Importantly, the absence of an interaction suggests that the visual information from the faces was not integrated with the semantic information from the labels. Instead, participants’ perceptions aligned

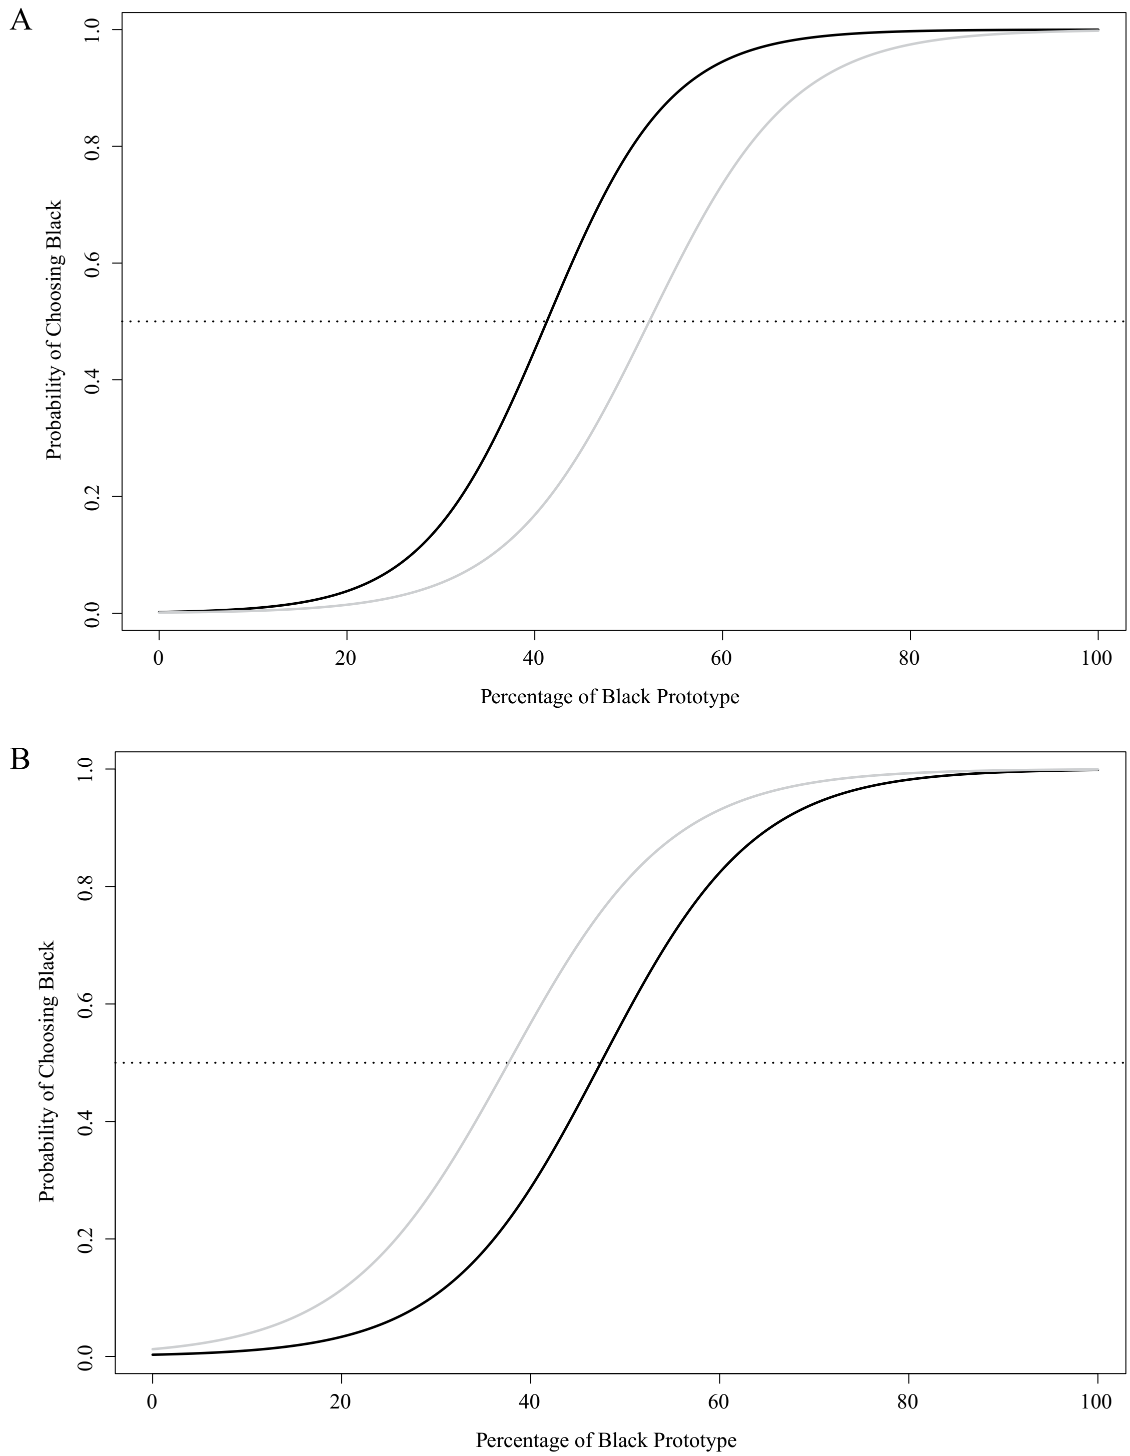

Figure 2. Shifts in the racial category boundary as a function of semantic information (i.e., response labels): (A) White–Black continuum and (B) Hispanic–Black continuum.

Note. The black line represents the aggregate function for participants categorizing the faces using continuum-congruent labels. The gray line represents categorizations of the same faces using continuum-incongruent labels. The horizontal dashed line marks the point of equiprobable assignment to either category, which intersects with the curves at their respective threshold points. The sigmoidal shape of the curve reflects the dichotomous nature of the categorizations. The steepness of each curve is described by its slope at the threshold level, and the changes in threshold position (illustrated by the distance between the curves at the point where they intersect with the dashed line) represent the different amount of Black features needed to distinguish the categories of faces when the continuum labels are varied.

with the response options even when the basic visual information from the faces clearly contradicted the labels. Thus, changes to the labels, which presumably constrained participants' interpretation about which social categories were viable, preferentially directed their categorizations.

These findings echo those of past studies suggesting that top-down cues can influence the perception of social groups. For instance, Eberhardt et al. (2003) found that participants' memory for racial labels affected their later perception and recall of faces. Here, we expanded on those findings to show that even simultaneously presented racial labels can affect categorization. Similarly relevant, MacLin and Malpass (2001, 2003) found that introducing race-diagnostic hairstyles to racially ambiguous faces caused participants to change their categorization of the faces' races, assuming that the perception of the visual cues in the hairstyles activated higher order concepts of race that then exerted top-down effects on the perception of the faces. Here, we co-presented low-level visual cues with semantic information and observed that semantic information changed the categorization of the faces somewhat similarly. Thus, the present data extend MacLin and Malpass's (2001, 2003) findings by providing empirical support for the influence of racial categories on visual perception top-down. These data therefore reflect a greater influence of semantic information upon social categorization than previously demonstrated.

Experiment 2

Manipulating semantic information changed participants' categorizations of race in Experiment 1. That is, the same faces were categorized differently when participants categorized them using one set of response labels (e.g., "White" and "Black") versus another (e.g., "Hispanic" and "Black"). Because the faces were identical, the change in categorization suggests that the labels simultaneously influenced the perception of the faces. There, we manipulated whether faces were allegedly White or Hispanic, in part because concomitant tests showed the Hispanic and White faces to be relatively interchangeable. In Experiment 2, we further examined how semantic information impacts social categorization by exchanging categories that are not so interchangeable. Specifically, we asked participants to categorize faces from the White–Hispanic morph continuum as either "White" versus "Black" or "Hispanic" versus "Black." That is, participants were asked to categorize faces as Black without the actual Black prototype face. Consistent with the results of Experiment 1 and previous research (e.g., Ho et al., 2011), we expected that participants would categorize the faces as Black sooner when judging the faces as "White" and "Black" versus "Hispanic" and "Black." Thus, the current experiment tested whether semantic information would affect perception even when the basic visual features of the faces clearly contradicted the visual information, highlighting the influence of semantic information on social categorization.

Method

We again recruited American participants from MTurk in exchange for financial compensation. Although we initially recruited 103 participants (62 female; age range = 18-73 years; 79 White, 14 Black, 2 Hispanic/Latino, 4 Asian, 4 Other), 31 categorized all of the faces into only one of the available categories, thereby rendering their data ineligible for analysis. We randomly assigned the participants to categorize the faces from the White–Hispanic continuum as either "White" and "Black" (n = 37) or as "Hispanic" and "Black" (n = 35). We presented each face 6 times in random order, as in Experiment 1.

Results

We again established categorization thresholds by fitting a logistic function to each participant's data. We extracted the thresholds from the predicted categorization probabilities for each participant (see Experiment 1). A planned comparison of the threshold values revealed that the participants who categorized the faces as "White" or "Black" indeed had a lower threshold than the participants who categorized the faces as "Hispanic" or "Black," as hypothesized: $t(70) = 2.10$, $p = .04$, $r = .24$, 95% CI = [.01, .45] (Figure 3).

Discussion

Participants' response options again influenced their apparent perception of the faces, even when one of the categorization labels was clearly incongruent with the faces' visual cues. The present findings thus provide a stronger demonstration of the extent to which semantic information can affect categorization than we observed in Experiment 1. There, we exchanged the labels of two groups that we had confirmed to be perceptually ambiguous (i.e., Hispanic and White). Here, however, we substituted an obviously incongruent label ("Black"). When doing so, we still observed the same pattern of results as would be expected when the faces matched the response options: "White" labels resulted in lower thresholds than "Hispanic" labels (see Experiment 1). However, perhaps it is possible that the presence of the specific labels here and in Experiment 1 coerced participants into responding in this way (Chen & Hamilton, 2012). We therefore asked participants to categorize only a single race in Experiment 3, rather than choosing between two specific groups.

Experiment 3

Although the results of Experiments 1 and 2 show robust and consistent changes in perception according to semantic information, perhaps those results are confounded by constraints on the response options (see Chen & Hamilton, 2012). In other words, maybe participants' categorizations were restricted by the use of non-veridical labels such that

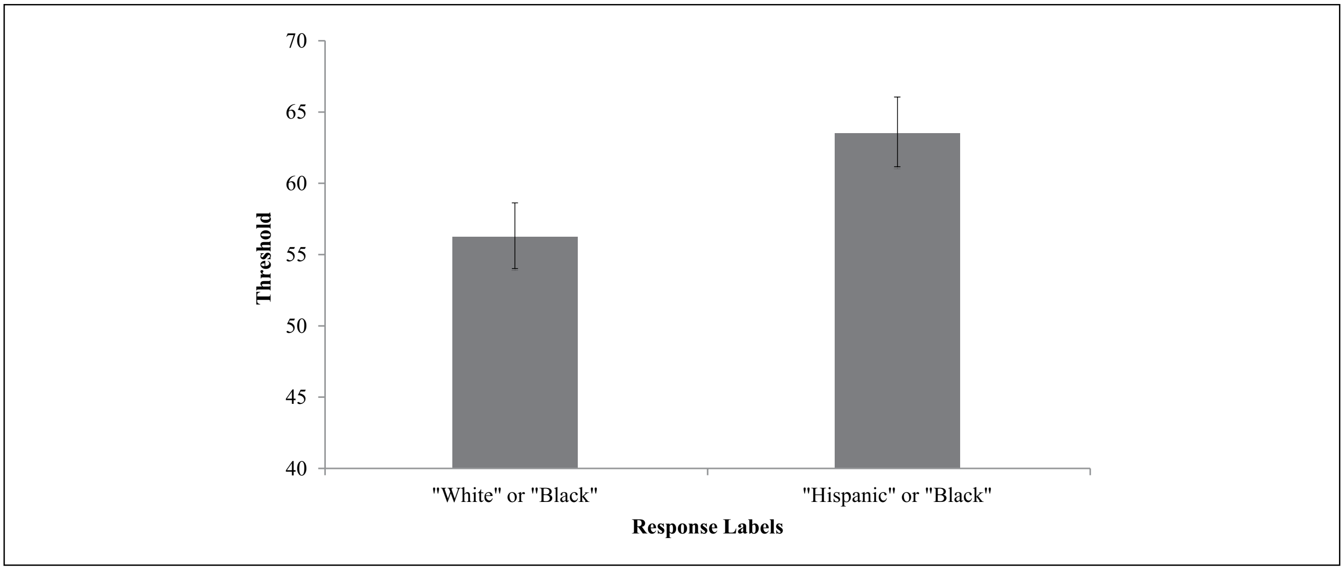

Figure 3. Thresholds for categorizing the faces from the White–Hispanic continuum as a function of semantic information (i.e., response labels).

they were compelled to categorize faces as “Hispanic,” for example, when the face contained more White features. We conducted Experiment 3 to eliminate this possibility and to expand our demonstration of how semantic information may guide social categorization.

To achieve this, we recruited participants to categorize the faces from the White–Black continuum as either “White” and “Black,” “White” and “Not White,” or “Black” and “Not Black.” This procedure therefore did not constrain the participants to using an incorrect category but still allowed us to observe how top-down information can affect their categorizations. Accordingly, we expected that the participants would categorize racially ambiguous faces as either “Not White” or “Not Black” in the conditions where one of the labels is unspecified. However, the same racially ambiguous faces would be categorized to specific groups when categorization is constrained by the two racial labels “White” and “Black,” yielding a greater probability of categorization to the White and Black categories. Furthermore, we expected that the frequency of categorizing the faces as White would be low across all conditions because White is the protected racial default group in North America (Ho et al., 2011). Thus, should we still find that the response labels augment the participants’ categorization thresholds, it would reinforce the top-down influence of semantic information on race categorization and increase our confidence that our effects were not simply the result of forcing participants to use particular category choices in Experiments 1 and 2.

Method

In total, 200 American MTurk Workers (104 female; age range = 18-65; 143 White, 20 Black, 15 Hispanic/Latino, 11 Asian, 11 Other) participated in the study for monetary compensation. Twelve participants categorized all of the faces to only one of the available categories, rendering their data inadmissible for analysis (final sample $N = 188$). We randomly assigned participants to categorize the faces from the White–Black continuum described in Experiment 1 as “White” and “Black,” “White” and “Not White,” or “Black” and “Not Black.” Again, each participant categorized each face from the continuum 6 times in random order.

Results

We analyzed the data using a one-way between-subjects ANOVA, which showed that participants’ categorization thresholds significantly differed as a function of their response options: $F(2, 185) = 13.67$, $p < .001$, $\eta^2 = .13$. Consistent with the results of Experiments 1 and 2, participants’ thresholds were lower when categorizing the faces in the “White” versus “Not White” condition compared with both the “White” versus “Black,” $t(123) = 3.35$, $p = .001$, $r = .29$, 95% CI = [.12, .44], and “Black” versus “Not Black,” $t(125) = 5.15$, $p < .001$, $r = .42$, 95% CI = [.27, .55], conditions. We observed a marginally significant difference between the “White” versus “Black” and “Black” versus “Not Black” conditions in the predicted direction: $t(122) = 1.79$, $p = .08$, $r = .16$, 95% CI = [−.02, .33]. Finally, we found a significant increasing linear trend from the “White” versus “Not White” to “White” versus “Black” to “Black” versus “Not Black” conditions: $t(186) = 5.16$, $p < .001$, $r = .35$, 95% CI = [.22, .47] (see Figure 4).

Discussion

Top-down semantic information continued to affect participants’ perceptions absent the use of two closed categories.

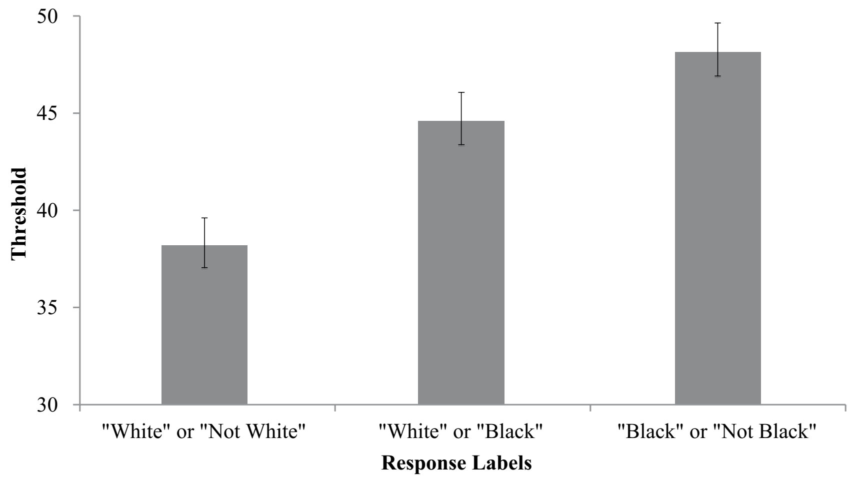

Figure 4. Participants' thresholds as a function of the response labels in Experiment 3.

Specifically, participants categorized the same faces differently when using the label pairs "White" versus "Not White," "White" versus "Black," and "Black" versus "Not Black." Importantly, all of these options are veridical and functionally equivalent. However, they emphasize the possible category options differently and the participants' categorizations reflected this. As in Experiments 1 and 2, participants exhibited a low threshold for categorizing faces as Black when given the options "White" and "Black." This threshold was even lower, however, when categorizing the faces as "White" versus anything else (i.e., "Not White"). Accordingly, the threshold was notably higher when categorizing the faces as "Black" versus anything else (i.e., "Not Black"). Thus, changing the category labels continued to influence the participants' categorizations even when all of the options were true and tenable. These results support those of Experiments 1 and 2, which replaced veridical with non-veridical labels. More important, they extend those findings by illustrating additional nuance in how semantic information can alter categorizations based on otherwise constant visual cues.

General Discussion

Across three experiments, we found consistent support for our hypothesis that semantic information can influence how targets are perceived. To review, we found that perceivers distinguished between White and Black faces and between Hispanic and Black faces in Experiment 1. Critically, they did not readily differentiate between White and Hispanic faces; rather, the boundary between the two groups was ambiguous. Furthermore, when participants categorized White and Black (Hispanic and Black) faces under the auspices that they were either Hispanic or Black (White or Black), their categorizations followed the pattern anticipated by the labels, not the faces. This suggested that the semantic information about the categories guided participants' categorizations independent of the faces' visual cues. In Experiment 2, we conducted a stronger test of this by swapping the Black label into categorizations of the White and Hispanic faces. Unlike in Experiment 1 (where the White and Hispanic faces looked hardly any different), introducing the Black label caused participants to mimic patterns akin to those of the White−Black and Hispanic−Black categorizations. Finally, in Experiment 3, we found that semantic information continued to affect participants' categorizations when we changed the labels in various veridical ways. Collectively, these data build a case for the key role of top-down semantic information on person categorization and perception.

These data are important because they demonstrate that perception is affected by semantic information. That is, although previous research demonstrated that higher order cognitions affect the perception and categorization of race (e.g., Eberhardt et al., 2003; Krosch & Amodio, 2014; MacLin & Malpass, 2001, 2003; Sacco et al., 2011), the current study is the first to demonstrate that contemporaneously perceived semantic information can affect perception. Importantly, some prior research suggested that this might be the case (MacLin & Malpass, 2001, 2003). However, even in that work, the effects of semantic information on categorization were merely assumed—the researchers suggested that manipulating visual cues activated categorical information that then affected perception top-down. Therefore, the current work explicitly tests and shows that semantic information has robust effects on categorizing race.

A number of studies over the last several decades have demonstrated how top-down processes can color the interpretation of visual information when presented prior to categorization (e.g., Eberhardt et al., 2003; Krosch & Amodio, 2014; Sacco et al., 2011; but see also MacLin & Malpass, 2001, 2003). What is distinct about the present work is that the top-down information affecting the perception of the bottom-up cues was perceived simultaneously but still externally induced. Thus, higher order information not only affects categorization upon retrieval from memory or as a function of fleeting individual conditions but also seems to modulate perception from one moment to the next. Thus, this work is important because it demonstrates the powerful effects that online changes to the context of perception can produce immediately upon stimulus presentation.

Despite its value, this work is not without limitations. Indeed, it remains unclear exactly how the introduction of new semantic information affected race categorization. It is possible that the participants' ingestion of the category labels affected their beliefs about the groups to which the faces might possibly belong. Although the data would support this, we do not have direct evidence that participants' beliefs were truly altered by the labels; hence, this remains a question that would require additional focused tests. Moreover, it is possible that the response labels could have directed participants' attention to different visual features of the same stimuli. For example, participants may have paid greater attention to stereotypically White features when the face was said to be White (vs. when it was said to be Hispanic, for instance). We did not explore the exact source of influence (e.g., visual, semantic information) on race perception in the current study, however. Although this limitation does not affect the reliability of our findings, further exploration of how semantic and visual information may simultaneously affect the perception of specific visual cues during race categorization would be a fruitful avenue for future research.

Along these lines, although we controlled for the basic visual features of the faces in our experiments by using the same faces across different conditions, the faces were still somewhat racially ambiguous. Because the faces were morphs, thereby engendering some degree of racial ambiguity, participants may have been more willing to accept that the response options provided meaningful information. Specifically, if participants regarded the Hispanic category as intermediate in skin tone between White and Black, they may have shifted their threshold according to the labels in Experiments 1 and 2 to reflect this difference. However, the fact that replacing the Hispanic label with "Black" in Experiment 2 turned a perceptually ambiguous distinction between the Hispanic and White faces (as observed in Experiment 1) into a discrete boundary similar to that observed for the White–Black comparison in Experiments 1 and 3 suggests that skin tone would have only become a distinctive characteristic when the category labels effectively requested it. Thus, this reinforces our central point that top-down semantic information exerts considerable influence on the perception of visual cues, and that it seems to do so potently and with immediacy during social perception and categorization.

Conclusion

In sum, we found that semantic information influenced how individuals categorized race. Furthermore, this work highlights one instance in which conceptual category information is integrated into perception during the categorization process. In addition, the current work demonstrates that external information may continuously mold the lens through which people interpret incoming visual information, changing how individuals categorize others into social groups. Thus, higher order inputs and categorization processes, in general, might be driven by interactions between information that is inherent to stimuli (their low-level features) and information that is external to stimuli (cues provided by context).

Acknowledgments

We would like to thank Paolo Palma for his help with data collection. We also would like to thank William Cunningham and John Paul Wilson for their insightful comments on earlier drafts.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada to Nicholas O. Rule.

Supplemental Material

The online supplemental material is available at http://pspb.sagepub.com/supplemental.

References

Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmid, A. M., Dale, A. M., & Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences of the United States of America*, *103*, 449-454.

Blair, I. V., Judd, C. M., Sadler, M. S., & Jenkins, C. (2002). The role of Afrocentric features in person perception: Judging by features and categories. *Journal of Personality and Social Psychology*, *83*, 5-25.

Chen, J. M., & Hamilton, D. L. (2012). Natural ambiguities: Racial categorization of multiracial individuals. *Journal of Experimental Social Psychology*, *48*, 152-164.

Eberhardt, J. L., Dasgupta, N., & Banaszynski, T. L. (2003). Believing is seeing: The effects of racial labels and implicit beliefs on face perception. *Personality and Social Psychology Bulletin*, *39*, 360-370.

Freeman, J. B., & Ambady, N. (2011). A dynamic interactive theory of person construal. *Psychological Review*, *118*, 247-279.

Freeman, J. B., Rule, N. O., Adams, R. B., Jr., & Ambady, N. (2010). The neural basis of categorical face perception: Graded representations of face gender in fusiform and orbitofrontal cortices. *Cerebral Cortex*, *20*, 1314-1322.

Gilbert, C. D., & Sigman, M. (2007). Brain states: Top-down influences in sensory processing. *Neuron*, *54*, 677-696.

Ho, A. K., Sidanius, J., Levin, D. T., & Banaji, M. R. (2011). Evidence for hypodescent and racial hierarchy in the categorization and perception of biracial individuals. *Journal of Personality and Social Psychology*, *100*, 492-506.

Krosch, A. R., & Amodio, D. M. (2014). Economic scarcity alters the perception of race. *Proceedings of the National Academy of Sciences*, *111*, 9079-9084.

MacLin, O. H., & Malpass, R. S. (2001). Racial categorization of faces: The ambiguous race face effect. *Psychology, Public Policy, and Law*, *7*, 98-118.

MacLin, O. H., & Malpass, R. S. (2003). The ambiguous-race face illusion. *Perception*, *32*, 249-252.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide*. Mahwah, NJ: Lawrence Erlbaum.

Macrae, C. N., & Bodenhausen, G. V. (2000). Social cognition: Thinking categorically about others. *Annual Review of Psychology*, *51*, 93-120.

Maddox, K. B. (2004). Perspectives on racial phenotypicality bias. *Personality and Social Psychology Review*, *8*, 383-401.

Pauker, K., Weisbuch, M., Ambady, N., Sommers, S. R., Adams, R. B., Jr., & Ivcevic, Z. (2009). Not so Black and White: Memory for ambiguous group members. *Journal of Personality and Social Psychology*, *96*, 795-810.

Rosenthal, R., & Rosnow, R. L. (2008). *Essentials of behavioral research: Methods and data analysis* (3rd ed.). New York, NY: McGraw-Hill.

Rule, N. O., Garrett, J. V., & Ambady, N. (2010). Faces and places: Geographic environment influences the ingroup memory advantage. *Journal of Personality and Social Psychology*, *98*, 343-355.

Rule, N. O., Tskhay, K. O., Freeman, J. B., & Ambady, N. (2014). On the interactive influence of facial appearance and explicit knowledge in social categorization. *European Journal of Social Psychology*, *44*, 529-535.

Sacco, D. F., & Hugenberg, K. (2012). Cooperative and competitive motives enhance perceptual sensitivity to angry and happy facial expressions. *Motivation and Emotion*, *36*, 382-395.

Sacco, D. F., Wirth, J. H., Hugenberg, K., Chen, Z., & Williams, K. D. (2011). The world in black and white: Ostracism enhances the categorical perception of social information. *Journal of Experimental Social Psychology*, *47*, 836-842.

Schweinberger, S. R., Burton, A. M., & Kelly, S. W. (1999). Asymmetric dependencies in perceiving identity and emotion: Experiments with morphed faces. *Perception & Psychophysics*, *61*, 1102-1115.

Smith, E. R., & Zárate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, *99*, 3-21.

Tajfel, H., & Wilkes, A. L. (1963). Classification and quantitative judgement. *British Journal of Psychology*, *54*, 101-114.

Walker, P. M., & Tanaka, J. W. (2003). An encoding advantage for own-race versus other-race faces. *Perception*, *32*, 1117-1126.

Żychaluk, K., & Foster, D. H. (2009). Model-free estimation of the psychometric function. *Attention, Perception, & Psychophysics*, *71*, 1414-1425.