Original Manuscript

Social Psychological and Personality Science
1–9

DOI: 10.1177/19485506211062285
journals.sagepub.com/home/spp
SAGE

Facial Stereotyping Drives Judgments of Perceptually Ambiguous Social Groups

Maryam Bin Meshar[1], Ryan M. Stolier[2], and Jonathan B. Freeman[1]

Abstract

When seeing a face, people form judgments of perceptually ambiguous social categories (PASCs), for example, gun-owners, gay people, or alcoholics. Previous research has assumed that PASC judgments arise from the statistical learning of facial features in social encounters. We propose, instead, that perceivers associate facial features with traits (e.g., extroverted) and then infer PASC membership via learned stereotype associations with those traits. Across three studies, we show that when any PASC is more stereotypically associated with a trait (e.g., alcoholics = extroverted), perceivers are more likely to infer PASC membership from faces conveying that trait (Study 1). Furthermore, we demonstrate that individual differences in trait–PASC stereotypes predict face-based judgments of PASC membership (Study 2) and have a causal role in these judgments (Study 3). Together, our findings imply that people can form any number of PASC judgments from facial appearance alone by drawing on their learned social–conceptual associations.

Keywords

face perception, social categories, impression formation, stereotyping

There seems no limit to the inferences we can make from others' facial appearance. Considerable research has explored the perception of social characteristics that tend to be perceptually obvious, such as gender, race, and age, which often occurs automatically and outside awareness (Freeman & Ambady, 2011; Macrae & Bodenhausen, 2000). But there are many social characteristics whose perceptual basis is far less clear, and yet perceivers nevertheless can infer them with ease.

Indeed, previous studies suggest that perceivers are able to make inferences about any number of perceptually ambiguous social categories (PASCs),[1] such as education attainment (Olivola & Todorov, 2010), sexual orientation (Rule et al., 2008), mental health (Daros et al., 2016; Giacomin & Rule, 2018; Kleiman & Rule, 2013), political orientation (Rule & Ambady, 2010), religion (Rule et al., 2010), social class (Bjornsdottir & Rule, 2017), criminality (Funk et al., 2017; Wilson & Rule, 2015), or occupation (Oldmeadow et al., 2013). Evidence supporting the ability to make PASC inferences has typically relied on demonstrations that perceivers show high agreement in their face-based judgments, and in limited cases, their judgments exhibit a modest correspondence with targets' actual category memberships (Tskhay & Rule, 2013). Although perceptually ambiguous, PASC judgments nevertheless strongly guide our social behavior and predict real-world outcomes (Olivola et al., 2018; Re & Rule, 2016). They may occur automatically and outside awareness (Rule et al., 2009), despite the fact that perceivers generally do not believe they can make these judgments (Daros et al., 2016; Rule et al., 2008, 2010).

A common assumption in the literature is that perceivers are able to form PASC judgments from having learned the statistical regularities between specific facial features and group membership (for reviews, see Rule & Sutherland, 2017; Tskhay & Rule, 2013). This is supported by findings showing that the accuracy of PASC judgments increases with greater familiarity with the PASC (Brambilla et al., 2013). Other researchers have criticized such a focus on accuracy on methodological grounds (Olivola & Todorov, 2010; Todorov et al., 2015). Yet accuracy aside, a central issue with the direct statistical learning perspective is that it fails to explain how perceivers could form judgments of PASCs that lack any physical basis or have never been encountered before.

In the current work, we argue for a different perspective that relies on indirect conceptual associations, that is, stereotypes, to drive PASC judgments. Specifically, we propose that perceivers hold visual associations with lower-level personality traits (e.g., eyelid-openness → unintelligent; Talamas et al., 2016), and then via a set of

[1]New York University, New York City, USA
[2]Columbia University, New York, NY, USA

Corresponding Author:
Maryam Bin Meshar, Department of Psychology, New York University, 6 Washington Place, New York, NY 10003, USA.
Email: maryam.meshar@nyu.edu

stereotype associations with those traits infer PASC membership (e.g., unintelligent → alcoholic). Indeed, it is well-established that perceivers consistently map specific facial features to numerous personality traits (Hehman et al., 2019; Oosterhof & Todorov, 2008), and these feature–trait mappings are highly consistent across the world (Jones et al., 2021). In fact, past studies have provided some evidence that such intermediary traits can be correlated with PASC judgments, as in perceived religion (facial skin → healthy → Mormon; Rule et al., 2010), sexual orientation (facial masculinity/femininity → gender-atypical → gay/lesbian; Freeman et al., 2010), or occupation (large jaw → dominant → banker; Oldmeadow et al., 2013). Here, we aim to comprehensively test whether such intermediary traits and PASC stereotypes play a broad and causal role in PASC judgments.

Across three studies, we provide evidence in support of this social–conceptual account. In Study 1, we show that the conceptual similarity between any given trait and PASC (e.g., the stereotype strength between "alcoholic" and "neurotic") predicts a corresponding similarity in how those attributes are perceived from faces (e.g., correlation between "neurotic" and "alcoholic" in face judgments). In Study 2, we demonstrate that perceivers exhibit unique individual differences in these stereotypical trait–PASC associations, which predict corresponding variability in face judgments. Finally, by manipulating these trait–PASC stereotype associations in Study 3 directly, we provide evidence for their causal role in shaping PASC judgments based on facial appearance. All stimuli, data, and analysis scripts are available on OSF (https://osf.io/uvzgt/).

Study 1

In Study 1, we tested whether stereotype associations between traits and PASCs predict the extent to which the facial appearance of a given trait is related to PASC judgments. Specifically, we hypothesized that when a PASC and a trait are more stereotypically associated, face-based PASC judgments would correlate more strongly with facial appearance conveying that trait. For instance, if perceivers believe that alcoholics tend to be unintelligent, then facial features judged as unintelligent should evoke perceptions of alcoholic category membership. To test this correspondence between conceptual similarity and perceptual similarity in a comprehensive manner across all pairs of PASCs and traits, we used a representational similarity analysis (RSA) approach (Kriegeskorte et al., 2008).

Method

Using an RSA framework, we captured similarity between all PASC–trait pairs at both the conceptual level (perceivers' stereotype associations) and perceptual level (perceivers' face judgments) and tested the correspondence between these two similarity models. To generate the conceptual similarity model, a group of participants provided data on conceptual associations; and to generate the perceptual similarity model, two independent groups of participants provided data on judgments of PASCs or traits from face stimuli.

Participants. All participants in this study and the studies that follow provided informed consent, completed demographic questions at the end, and were financially compensated. Participants were all US residents and primary English speakers.

For the conceptual similarity model, all conceptual associations could be assessed within subject, and thus, we based target sample size on prior work estimating similar models (Stolier et al., 2018), seeking a target sample of $n = 100$. For the perceptual similarity model, due to time constraints each participant could only judge the face stimuli on one PASC or trait. Thus, our target sample was 25 participants per each PASC/trait being judged, consistent with current recommendations regarding interrater reliability in face-based judgments (Hehman et al., 2018) and prior work estimating similar models (Stolier et al., 2018). This totaled a target sample of $n = 425$ across all 11 PASCs and six traits.

All participants were recruited from Amazon Mechanical Turk. Conceptual association data were collected from 95 participants (original $n = 115$, 20 subjects removed for failing attention checks; $M_{age} = 37.03$ years, $SD_{age} = 11.21$ years; 42 females, 45 males, 1 other; 66 Whites, 13 Blacks, 1 Asian, 8 others; seven participants did not complete demographic questions). Face PASCs judgments were collected from 272 participants ($M_{age} = 37.06$ years, $SD_{age} = 11.45$ years; 124 females, 129 males, 1 other; 192 Whites, 34 Blacks, 9 Asians, 19 others; 18 participants did not complete demographic questions). Face traits judgments were collected from 174 participants ($M_{age} = 35.96$ years, $SD_{age} = 12.00$ years; 92 females, 77 males, 1 other; 129 Whites, 19 Blacks, 3 Asians, 19 others; four participants did not complete demographic questions). Participants were randomly divided roughly equally (average of ˜25–30 participants per PASC/trait).

Stimuli. For face-based PASC judgments, we chose several PASCs used in prior research on face judgments (Olivola & Todorov, 2010; Rule & Sutherland, 2017; see Table 1). For face-based trait judgments, we selected a subset of trait stimuli based on the primary dimensions of face impressions (Oosterhof & Todorov, 2008) and Big-Five factors of personality (Goldberg, 1999), so as to be maximally independent from one another. These traits were "dominant," "extroverted," "intelligent," "neurotic," "trustworthy," and "attractive." Face stimuli were taken from the Chicago Face Database (Ma et al., 2015) and included 93 portrait photographs of young to middle-aged White male adult

Table 1 PASCs Used in Study 1.

	Face-based judgment task	
Conceptual similarity task	PASC question	PASC response options
Homosexual	What is this person's sexual orientation?	Heterosexual/homosexual
Use drugs	Does this person use drugs?	Yes/no
Have attended public schools	Did this person go to public school?	Yes/no
Have been arrested	Has this person ever been arrested?	Yes/no
A virgin	Is this person a virgin?	Yes/no
Drink alcohol a lot	Does this person drink alcohol a lot?	Yes/no
Own a gun	Does this person own a gun?	Yes/no
Their parents are divorced	Are this person's parents divorced?	Yes/no
Have gotten into a fist fight	Has this person ever gotten into a fist fight?	Yes/no
In a long-term relationship	Is this person in a long-term relationship?	Yes/no
Have a college degree	Does this person have a college degree?	Yes/no

Note. Conceptual similarity task refers to perceivers' stereotype associations between traits and PASCs; face-based judgment task refers to participants' categorizations of faces as to whether they belong to different PASCs. PASCs = perceptually ambiguous social categories.

faces with neutral facial expressions. Stimuli were vignetted to preserve the internal face and obscure peripheral features (e.g., hairstyles).

Procedure

Conceptual Similarity Task. Participants were informed they would partake in a study on how different personality traits correspond to social characteristics in the world. Each trial item asked either: "If someone is a/an [TRAIT] person, how likely are they to [PASC]?" or "How likely is someone who [PASC] to be [TRAIT]?" For example, "If someone is a [trustworthy] person, how likely are they to [own a gun]?" Participants evaluated the conceptual relationship of each PASC–trait pair for the 6 trait × 11 PASC stimuli (1–7 Likert-type scale, 1—"Not at all likely"–7—"Very likely"), presented in both orders given the wording of the item question (e.g., "trustworthy–gun owner" and "gun owner–trustworthy"). Therefore, there were a total of 132 randomized trials for each participant. An additional six attention check trials were included, where participants who failed to select a specific response option (as instructed) on more than 50% of trials were excluded for not following instructions.

Face-Based Judgments Task. Participants were randomly assigned to evaluate either one of the 11 PASC stimuli or one of the six personality trait stimuli in faces. In the task, participants rated each of the 93 face stimuli (randomized order) on the PASC (see Table 1 for PASC questions and response options) or the trait they were assigned ("How [TRAIT] is this person?" on a Likert-type scale: "1—Not at all [TRAIT]" to "7—Very [TRAIT]").

Results and Discussion

Taking an RSA approach, a conceptual similarity model and perceptual similarity model may be represented each as a matrix of all pair-wise similarities between PASCs and traits, that is, a "similarity matrix" (a total of 66 possible unique pairwise combinations of all PASC–trait stimuli). The conceptual similarity model comprises cells for each PASC–trait pair denoting the average conceptual similarity rating for that pair. The perceptual similarity model comprises cells for each PASC–trait pair denoting the point-biserial Pearson correlation between binary judgments of that PASC and Likert-type ratings of that trait for the 93 face stimuli. Each matrix is then flattened into a vector of unique pair-wise PASC–trait similarities, and the correspondence between these conceptual and perceptual similarity vectors is tested using Spearman-rank correlation (rank-ordering is preferable when comparing similarity matrices from different measures as it does not assume a linear relation; Kriegeskorte et al., 2008). This analysis therefore tests whether the extent to which a PASC–trait pair is conceptually related (e.g., "drug user"—"neurotic" more strongly related than "drug user"—"intelligent") predicts the extent to which that PASC and trait are also correlated in face judgments (e.g., faces categorized as "drug users" appear more "neurotic" than "intelligent"; see Figure 1). Consistent with our predictions, the conceptual and perceptual similarity models for PASC–trait relations were strongly related (Spearman's $\rho(64) = .761$, $\rho^2(64) = .579$, $p < .0001$; 95% CI = [0.636, 0.847]; Figure 1C). These findings demonstrate that, on average across perceivers, the relationship between PASC judgments and facial appearance of a trait strongly resembles perceivers' stereotype associations between PASCs and traits. Additional analyses accounting for different sources of variability (i.e.,

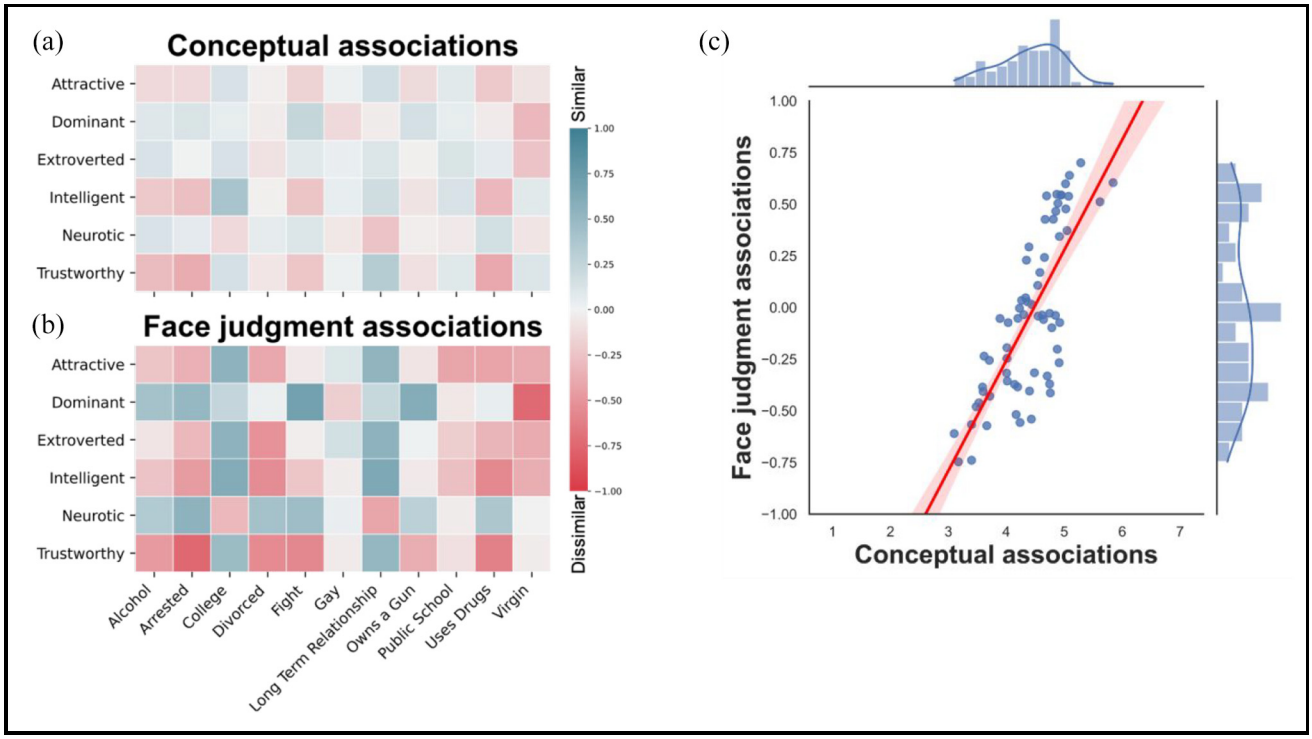

Figure 1. Study 1 Results *Note*. (A) The conceptual similarity model is shown as a similarity matrix, with each cell describing the conceptual association between each PASC–trait pair in the form of an average similarity rating (panel A; *x*-axis, panel C). Blue colors indicate similarity (more positive ratings), and red colors indicate dissimilarity (more negative ratings). (B) The perceptual similarity model is shown as a similarity matrix, with each cell describing the Pearson correlation between face judgments of a PASC and face judgments of a trait performed by independent participants (panel B; *y*-axis, panel C). Blue colors indicate more positive correlations, and red colors indicate more negative correlations. (C) The conceptual and perceptual similarity matrices were then flattened into vectors and submitted to RSA in which their correspondence was tested using a Spearman rank-order correlation. Each data point is one PASC–trait pair. A significant correlation was observed, where the extent to which a PASC and trait are more conceptually related was associated with a stronger correlation in face judgments for that PASC and trait. PASC = perceptually ambiguous social categories; RSA = representational similarity analysis.

participants, PASCs, and traits) using mixed-effects models converged on the same pattern of results (see Supplementary Materials).

Study 2

In Study 1, our approach provided a bird's eye view of how PASC judgments correspond to face trait appearance to the extent those PASCs are stereotypically associated with a trait. In Study 2, we measure individual differences in stereotype associations and face judgments by modifying the paradigm to allow for a measurement of within-subject correspondences between conceptual and perceptual similarity. To better establish the generality of our effects, we also expanded the set of PASCs (a total of 15) and the set of trait appearances (a total of 8) tested.

Method

Data were collected from two independent groups of participants. In the first group, whose individual differences we aimed to measure, participants were randomly assigned to one of the 15 PASCs. Next, they made PASC judgments of face stimuli and subsequently reported on their conceptual association between their assigned PASC and all eight traits. Traits were judged by the second independent group of participants so as to reduce demand characteristics. Because six of the trait stimuli overlap with those used in Study 1, these face ratings were obtained from Study 1 and participants in Study 2 were assigned to make face ratings on one of the two remaining traits.

Participants. For the perceptual similarity model, as in Study 1, our target sample was 25 participants per each PASC/trait being judged. Across the 15 PASCs and eight traits (see the "Stimuli" section), this totaled a target sample of $n = 575$ for the perceptual similarity model.

Participants were recruited from Mechanical Turk. We collected data on PASC judgments of faces and PASC–trait conceptual associations from 414 subjects, original ($n = 451$); we excluded subjects due to task incompletion ($n = 2$); failing attention checks ($n = 25$); or constant responses ($n = 10$); $M_{age} = 33.36$ years, $SD_{age} = 6.35$ years; 192 females, 222 males; 295 Whites, 53 Blacks, 50 Asians, and

16 others. We collected data on face judgments of traits from 234 subjects (n_{Study1} = 174; demographics reported in Study 1; n_{Study2} = 60; M_{age} = 36.88 years, SD_{age} = 11.54 years; 25 females, 33 males; 44 Whites, 3 Blacks, 7 Asians, 4 others, demographic data for two participants were missing).

Stimuli. To increase the generality of our effects, we expanded the set of PASCs and traits. The set of PASCs included others used in previous research: "CEO," "criminal," "democrat," "drug user," "homeless person," "lawyer," "military member," "musician," "nerd," "politician," "republican," "scientist," "sports fan," "stoner," "teacher" (Koch et al., 2016; Rule & Sutherland, 2017). Trait stimuli approximately followed the Big-Five factors of personality. In addition to the six traits tested in Study 1 (see the "Stimuli" section), we added "friendly" and "open-minded." Face stimuli were a subset of the Chicago Face Database face stimuli used in Study 1. As some of the PASCs in Study 2 do not plausibly apply to younger targets (e.g., "politician"), we only used the White male adult faces of Study 1 that were middle-aged (ages 25–45), resulting in 59 face stimuli. Trait ratings on the six traits obtained from Study 1 here correspond to this subset of the 59 face stimuli.

Procedure. Unlike in Study 1, here each of our primary participants making PASC judgments were also assessed on their own conceptual association between the PASC in question and all eight traits. First, each participant was randomly assigned to one of 15 PASCs and they categorized all face stimuli on their assigned PASC (yes/no response options). Next, using an identical conceptual similarity task to that of Study 1, participants reported their conceptual association between their assigned PASC and all eight traits with 16 reversed ordered items. Each PASC–trait pair then was averaged together, such that, each participant had eight scores corresponding to their conceptual associations between their assigned PASC and the eight traits.

An independent group of participants provided face-based judgments of traits. Six out of the eight traits were obtained from Study 1 (see the "Method" section). Participants in Study 2 were randomly assigned to judge one of the two remaining trait stimuli on an identical face-based trait judgments task to that of Study 1. These ratings were then averaged across participants.

Results and Discussion

For each of our primary participants (who completed the PASC judgment and conceptual association tasks), we calculated the point-biserial Pearson correlation between their binary judgments of the assigned PASC (0 = *no*; 1 = *yes*) and mean ratings of a specific trait across the 59 face stimuli. This procedure was repeated for all traits resulting in eight similarity values between each subject's assigned PASC and the eight traits. This correlation coefficient per participant therefore represents the extent to which their PASC judgments of faces were related to the facial appearance of the trait. We then predicted perceivers' perceptual similarity between their PASC judgments and traits from their corresponding conceptual similarity using a mixed-effects model to account for variation between participants, traits, and PASCs.

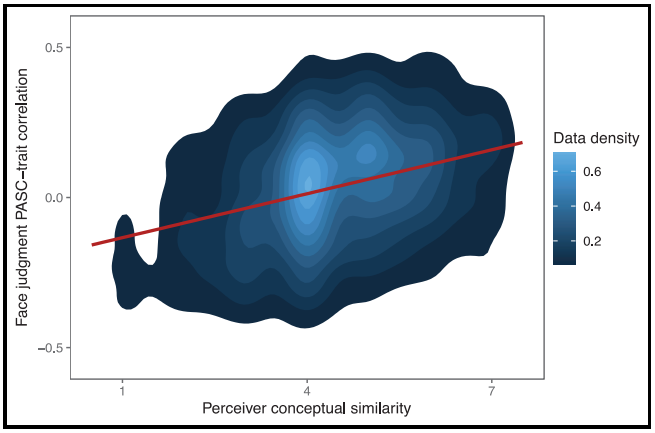

Figure 2. Study 2 Results *Note*. Participants' correlation coefficient between their face judgments of PASCs and the trait appearance of faces is plotted (*y*-axis) as a function of their PASC–trait conceptual association (*x*-axis). We observed a positive relationship between perceivers' conceptual association between a PASC and trait (*x*-axis) and the extent to which the facial appearance of that trait was related to their PASC judgments of faces (*y*-axis). For illustrative purposes, we show a density plot due to the large quantity of data (n = 414; total of 3,312 data points). Each data point is one PASC–trait pair (a total of eight pairs) unique to each subject. Light colors of the density plot represent higher probability of each value given the range of possible values. The red line denotes the predicted values estimated through an ordinary least-squares linear fit (note that actual analyses were conducted using mixed-effects regression model). PASCs = perceptually ambiguous social categories.

The perceptual similarity of participants' trait–PASC pairs was regressed onto their conceptual similarity, with random intercepts for participants, traits, and PASCs.[2] The model revealed a strong positive relationship, b = .041, SE = .003, $t(3{,}228.900)$ = 16.367, $p < .0001$, 95% CI = [0.036, 0.046], marginal R^2 = .07, conditional R^2 = .28 (Figure 2). These findings extend those of Study 1, showing that individual differences in participants' stereotypical beliefs in how any given PASC is similar to a trait predicts a corresponding similarity in how these PASC judgments are related to the facial appearance of that trait. For example, individuals believing criminals to be more neurotic (in comparison to those who believe criminals are

less neurotic) are more likely to categorize neurotic faces as criminals.

Study 3

Thus far, the evidence presented that stereotype associations linking PASCs to traits set the stage for how PASCs are judged has been correlational in nature. Aiming to provide causal evidence, in Study 3, we manipulated participants' stereotype associations by shifting their conceptual beliefs about PASCs. If PASC judgments rely on PASC–trait stereotype associations, then when those associations are strengthened or weakened, we would predict a corresponding effect on PASC judgments.

Method

Participants were randomly assigned to one PASC–trait pair out of six possible pairs: three PASCs ("has been arrested," "in a long-term relationship," and "lawyers") × two traits ("extraversion" and "neuroticism"). They were also randomly assigned to one of two between-subject conditions wherein they were led to believe that their assigned PASC was either positively or negatively correlated with their assigned trait using a faux scientific article. Participants then completed PASC judgments of faces and a manipulation check.

Participants. We aimed to collect ˜120 participants per between-subject association condition and oversampled to account for excluding those who may fail the manipulation check (˜25% of participants; Stolier et al., 2020). Our final sample included 370 Mechanical Turk participants (original $n = 377$; five participants did not consent to using their data, two participants were excluded for having constant responses; $M_{age} = 33.54$ years, $SD_{age} = 16.98$ years; 187 females, 179 males, 3 declined, 1 other; 277 Whites, 54 Blacks, 21 Asians, and 18 others).

Stimuli. Three representative PASCs were used in this study: "having been arrested," "being a lawyer," "being in a long-term relationship" (Koch et al., 2016; Oldmeadow et al., 2013; Olivola & Todorov, 2010). Two traits reflecting two of the Big-Five factors used in the previous studies were paired with these PASC stimuli: "extroverted" and "neurotic." The face stimuli were identical to those used in Study 2.

To manipulate participants' conceptual associations between traits and PASCs, we created faux scientific articles which explained how scientists had presumably discovered a relationship between the participant's assigned PASC and trait. This relationship was described as either positive (e.g., lawyers tend to be neurotic) or negative (e.g., lawyers tend to be less neurotic) depending on the participant's randomly assigned association condition. These articles were adapted from prior research manipulating lay theories and conceptual associations of personality (Coleman & Hong, 2008; Stolier et al., 2020).

Procedure. First, participants were presented with the faux scientific article which appeared on the screen for 2 min before allowing participants to proceed. After reading the article, participants were instructed to summarize their thoughts to encourage engagement with the information.

Next, participants judged the face stimuli on their assigned PASC using binary Yes/No responses in a manner identical to PASC judgment task of Study 2. As in Study 2, these participants did not make trait judgments of faces (trait judgments were derived from independent participants), to avoid issues of demand characteristics. Trait judgment data for the face stimuli for the two traits used in the current study were taken from Study 1.

Finally, participants completed two questions about their PASC–trait conceptual association that served as a manipulation check, and then were debriefed. This was identical to the conceptual similarity task of Study 1 and these two ratings were averaged together. This measurement method has been modeled from prior research assessing the effectiveness of manipulating lay theories (e.g., associations between personality traits; Coleman & Hong, 2008; Stolier et al., 2020).

Results and Discussion

Manipulation Check. Overall, the manipulation was successful, with those who were led to believe the trait and PASC are positively related reporting a stronger average trait–PASC conceptual association ($M = 5.41$, $SD = 1.12$) than those led to believe they were negatively related ($M = 3.64$, $SD = 1.70$), $t(368) = -11.821$, $p < .0001$, $d = 1.23$, 95% CI = [–2.058, –1.471]. Additional analyses showed that the associations between conditions were different from the scale midpoint and that the manipulation was equally successful across all PASCs and traits tested (see Supplementary Materials).

To test the primary hypothesis, we examined participants who successfully responded to the manipulation, that is, reported an average conceptual association that was equal to or above/below 4 ("Neutral") for participants assigned to the positive/negative condition, respectively, resulting in 80 exclusions (21.2%). A chi-square test of independence showed that this exclusion of participants was equally distributed across the 12 (three PASCs × two traits × two association conditions), $\chi^2(5) = 5.016$, $p = .414$.

Main Analyses. Using the trait data from Study 1, we obtained the mean trait rating for each of the 59 faces on the two traits used in the current study, which we grand-mean centered. We ran a generalized linear mixed-effects model to predict each of Study 3 participants' PASC face judgments (0 = *no*, 1 = *yes*) from their assigned

association condition (positive or negative), the average face appearance of their assigned trait (derived from Study 1 data), the average face appearance of their non-assigned trait (derived from Study 1 data), and association condition × assigned trait and association condition × non-assigned trait interactions. We included random intercepts for individual participants and faces. For example, for a participant assigned to positively associate lawyers with extraversion, we predicted their face "lawyer" judgment ratings from their association condition, faces' extraversion appearance (assigned trait), faces' neuroticism appearance (non-assigned trait), and the association condition × assigned trait and association condition × non-assigned trait interactions. The central prediction was the association condition × assigned trait interaction, whereby the relationship between face appearance of the assigned trait and judgments of a PASC is strengthened in the positive-association condition relative to the negative-association condition. Odds ratio (OR) and Wald Z are reported as a measure of effect size.

There were no main effects of association condition (OR = 1.099, SE = .110, z = .941, p = .346, 95% CI = [0.903, 1.337]), assigned-trait rating (OR = .890, SE = .063, z = −1.638, p = .101, 95% CI = [0.774, 1.023]), or non-assigned trait rating (OR = 1.106, SE = .078, z = 1.425, p = .154, 95% CI = [0.963, 1.270]). More importantly, the association condition × assigned trait interaction was significant (OR = 1.249, SE = .073, z = 3.807, p < .001, 95% CI = [1.114, 1.401]). The probability of categorizing faces with high trait appearance as the assigned PASC increased by 24.9% in the positive-association condition relative to the negative-association condition. For instance, participants who were led to believe lawyers tend to be extroverted were more likely to categorize extroverted-looking faces as lawyers, in comparison to participants who were led to believe that they were negatively associated. Note that, given participants already hold strong prior PASC–trait associations (consistent with the results of Studies 1 and 2), of interest here is the relative change in the PASC–trait association due to shifting participants' beliefs, rather than the absolute level of the association as the absolute level largely reflects their priors.

Interestingly, there was also a significant association condition × non-assigned trait interaction (OR = .809, SE = .047, z = −3.663, p < .001, 95% CI = [0.723, 0.906]) which followed the opposite pattern. The probability of categorizing the face as the assigned PASC *decreased* by 19.1% for participants in the positive-association condition relative to those in the negative-association condition. For instance, participants who were led to believe lawyers tend to be more extroverted were less likely to categorize *neurotic* looking faces as lawyers. This result is not surprising given that the appearance of the two traits (extraversion and neuroticism) in the faces used in this study are negatively correlated (r = −.223), as is typically observed in face impressions research (Oosterhof & Todorov, 2008; Stolier et al., 2020). Thus, if the positive-association condition increases the tendency for extroverted-looking faces to be categorized as a given PASC, it is likely the same condition will decrease the tendency for neurotic-looking faces to be categorized as that PASC. Additional regression models were run to more fully account for variability related to specific PASCs or traits. While the effects were more/less pronounced for certain PASCs, the overall pattern of results was not meaningfully changed (see Supplementary Materials).

General Discussion

Overall, we provide evidence that PASC judgments of faces are driven by stereotype associations with trait-related facial appearances. The degree to which a perceiver stereotypically associated any given PASC and trait predicted a greater tendency to infer PASC membership when a target's facial appearance conveyed that trait (Study 1). Moreover, we found that individual differences in perceivers' own unique stereotype associations predicted to what extent their pattern of PASC judgments was related to trait-related facial appearances (Study 2). Finally, by manipulating perceivers' PASC–trait stereotype associations, we demonstrated a causal role for these associations in driving how PASCs are perceived from faces (Study 3).

Previous research has proposed that people form PASC inferences of faces from having detected the statistical regularities between specific facial features and PASC membership via prior encounters with category members (Brambilla et al., 2013; Rule & Sutherland, 2017; Tskhay & Rule, 2013). This requires there to be some physical basis to category membership as well as a degree of familiarity with category members. While our results do not refute this possibility for certain PASCs, they do suggest that direct statistical learning cannot be the only mechanism driving PASC judgments. Instead, our findings suggest that perceivers draw on facial stereotyping and social–conceptual associations to extrapolate about any number of PASC judgments from facial appearance alone. It should be noted, however, that statistical learning is by no means inconsistent with our perspective. For instance, perceivers could detect the statistical regularities in the personality traits expressed by PASC members, whether from their own direct observations or from third-party relaying by others, and then use these trait–PASC associations to infer PASC membership. As these trait–PASC associations would thus arise from preconceived notions about PASC members, perceivers could infer the membership of virtually any PASC—even if it lacks any physical basis or has never been encountered—because perceivers could draw on such stereotype associations (which may be

statistically learned). Future research could measure PASC–trait stereotypes in tandem with exposure to PASCs to yield a fuller understanding of how facial stereotyping and direct statistical learning may shape PASC judgments.

Our findings may also speak to theoretically important questions regarding the accuracy and real-world consequences of PASC judgments, which have been the focus of considerable social psychological research (for review, see Rule & Sutherland, 2017). For example, finding that inter-individual variability in the associative strength of PASC stereotypes affects how perceivers judge PASC membership raises new questions about potential individual differences and important moderators in the accuracy of PASC judgments and these judgments' correspondence with real-world consequences.

This work is not without its limitations. The use of White male faces was important for avoiding confounds, such as individual differences in gender and racial bias (see Xie et al., 2021). Future studies should test the generality of our findings to faces varying in gender and race and explore how gender and racial stereotypes might interact with trait–PASC stereotypes. Another limitation is that our studies do not directly measure the temporal dynamics involved in feature trait PASC activation. The results of Study 3 suggest that trait representations can have a causal influence on PASC judgments, but feature, trait, and PASC representations are likely all interactive and bidirectionally related (Freeman et al., 2020; Kunda & Thagard, 1996). It is also likely that context and task goals could amplify or attenuate the activation of such representations (Freeman & Ambady, 2011; Freeman et al., 2020). While previous work suggests that faces' trait representations (Freeman et al., 2014; Kidder et al., 2018; Macrae & Martin, 2007) and PASC representations (Rule et al., 2009) can be automatically activated, future research could test whether such a feature–trait–PASC cascade occurs automatically or only in relevant task contexts.

In sum, our findings suggest that judging ambiguous group membership from faces need not require any direct statistical learning or previous exposure to category members, but instead can broadly depend on learned stereotypes via intermediary trait representations. Such results bring to light how social-conceptual knowledge lets us see an entire social world in others' faces.

Acknowledgment

The authors thank John Andrew Chwe and Clodagh Cogley for their assistance in materials' development and data collection.

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 work was supported in part by research grants NIH-F31-MH114505 (R.M.S.) and NSF-BCS-1654731 (J.B.F.).

ORCID iDs

Maryam Bin Meshar https://orcid.org/0000-0002-3220-4249
Jonathan B. Freeman https://orcid.org/0000-0002-2061-8460

Supplemental Material

Supplemental material is available in the online version of the article.

Notes

1. Following prior work (Olivola & Todorov, 2010), perceptually ambiguous social categories (PASCs) were defined as social characteristics not easily inferable from facial appearance but still defined by a clear category in reality (i.e., individuals could self-label into the category with high confidence). For instance, "alcoholics" would be considered a PASC, while "people who work more than 80 hours per week" would not.
2. An additional model accounting for random intercepts and slopes for participants and trait–PASC pairs also converged on the same results (see Supplementary Materials).

References

Bjornsdottir, R. T., & Rule, N. O. (2017). The visibility of social class from facial cues. *Journal of Personality and Social Psychology*, *113*(4), 530–546.

Brambilla, M., Riva, P., & Rule, N. O. (2013). Familiarity increases the accuracy of categorizing male sexual orientation. *Personality and Individual Differences*, *55*(2), 193–195.

Coleman, J. M., & Hong, Y.-Y. (2008). Beyond nature and nurture: The influence of lay gender theories on self-stereotyping. *Self and Identity*, *7*(1), 34–53.

Daros, A. R., Ruocco, A. C., & Rule, N. O. (2016). Identifying mental disorder from the faces of women with borderline personality disorder. *Journal of Nonverbal Behavior*, *40*(4), 255–281.

Freeman, J. B., & Ambady, N. (2011). A dynamic interactive theory of person construal. *Psychological Review*, *118*(2), 247–279. https://doi.org/10.1037/a0022327

Freeman, J. B., Johnson, K. L., Ambady, N., & Rule, N. O. (2010). Sexual orientation perception involves gendered facial cues. *Personality and Social Psychology Bulletin*, *36*, 1318–1331.

Freeman, J. B., Stolier, R. M., & Brooks, J. A. (2020). Dynamic interactive theory as a domain-general account of social perception. In B. Gawronski (Ed.), *Advances in experimental social psychology* (Vol. 61, pp. 237–287). Academic Press.

Freeman, J. B., Stolier, R. M., Ingbretsen, Z. A., & Hehman, E. A. (2014). Amygdala responsivity to high-level social information from unseen faces. *The Journal of Neuroscience*, *34*(32), 10573–10581.

Funk, F., Walker, M., & Todorov, A. (2017). Modelling perceptions of criminality and remorse from faces using a data-driven computational approach. *Cognition and Emotion*, *31*(7), 1431–1443.

Giacomin, M., & Rule, N. O. (2018). Eyebrows cue grandiose narcissism. *Journal of Personality*, *87*(2), 373–385.

Goldberg, L. R. (1999). A broad-bandwidth, public domain, personality inventory measuring the lower-level facets of several five-factor models. *Personality Psychology in Europe*, *7*(1), 7–28.

Hehman, E., Stolier, R. M., Freeman, J. B., Flake, J. K., & Xie, S. Y. (2019). Toward a comprehensive model of face impressions: What we know, what we do not, and paths forward. *Social and Personality Psychology Compass*, *13*(2), e12431.

Hehman, E., Xie, S. Y., Ofosu, E. K., & Nespoli, G. A. (2018). *Assessing the point at which averages are stable: A tool illustrated in the context of person perception*. https://osf.io/mwtuz/

Jones, B. C., DeBruine, L. M., Flake, J. K., Liuzza, M. T., Antfolk, J., Arinze, N. C., Ndukaihe, I. L. G., Bloxsom, N. G., Lewis, S. C., Foroni, F., Willis, M. L., Cubillas, C. P., Vadillo, M. A., Turiegano, E., Gilead, M., Simchon, A., Saribay, S. A., Owsley, N. C., Jang, C., . . . Coles, N. A. (2021). To which world regions does the valencedominance model of social perception apply? Nature Human Behaviour, 5(1), 159–169.

Kidder, C. K., White, K. R., Hinojos, M. R., Sandoval, M., & Crites, S. L., Jr. (2018). Sequential stereotype priming: A meta-analysis. *Personality and Social Psychology Review*, *22*(3), 199–227.

Kleiman, S., & Rule, N. O. (2013). Detecting suicidality from facial appearance. *Social Psychological and Personality Science*, *4*(4), 453–460.

Koch, A., Imhoff, R., Dotsch, R., Unkelbach, C., & Alves, H. (2016). The ABC of stereotypes about groups: Agency/socioeconomic success, conservative–progressive beliefs, and communion. *Journal of Personality and Social Psychology*, *110*(5), 675–709.

Kriegeskorte, N., Mur, M., & Bandettini, P. (2008). Representational similarity analysis–connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience, 2*, Article 4. https://doi.org/10.3389/neuro.06.004.2008

Kunda, Z., & Thagard, P. (1996). Forming impressions from stereotypes, traits, and behaviors: A parallel-constraint-satisfaction theory. *Psychological Review*, *103*(2), 284–308.

Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*(4), 1122–1135.

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

Macrae, C. N., & Martin, D. (2007). A boy primed Sue: Feature-based processing and person construal. *European Journal of Social Psychology*, *37*(5), 793–805.

Oldmeadow, J. A., Sutherland, C. A., & Young, A. W. (2013). Facial stereotype visualization through image averaging. *Social Psychological and Personality Science*, *4*(5), 615–623.

Olivola, C. Y., Tingley, D., & Todorov, A. (2018). Republican voters prefer candidates who have conservative-looking faces: New evidence from exit polls. *Political Psychology*, *39*(5), 1157–1171.

Olivola, C. Y., & Todorov, A. (2010). Fooled by first impressions? Reexamining the diagnostic value of appearance-based inferences. *Journal of Experimental Social Psychology*, *46*(2), 315–324.

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, *105*, 11087–11092.

Re, D. E., & Rule, N. O. (2016). Predicting firm success from the facial appearance of Chief Executive Officers of non-profit organizations. *Perception*, *45*(10), 1137–1150.

Rule, N. O., & Ambady, N. (2010). Democrats and Republicans can be differentiated from their faces. *PLOS ONE*, *5*(1), Article e8733.

Rule, N. O., Ambady, N., Adams, R. B., Jr., & Macrae, C. N. (2008). Accuracy and awareness in the perception and categorization of male sexual orientation. *Journal of Personality and Social Psychology*, *95*, 1019–1028.

Rule, N. O., Garrett, J. V., & Ambady, N. (2010). On the perception of religious group membership from faces. *PLOS ONE*, *5*(12), Article e14241.

Rule, N. O., Macrae, C. N., & Ambady, N. (2009). Ambiguous group membership is extracted automatically from faces. *Psychological Science*, *20*, 441–443.

Rule, N. O., & Sutherland, S. L. (2017). Social categorization from faces: Evidence from obvious and ambiguous groups. *Current Directions in Psychological Science*, *26*(3), 231–236.

Stolier, R. M., Hehman, E., & Freeman, J. B. (2020). Trait knowledge shapes a common structure across social cognition. *Nature Human Behaviour*, *4*(4), 361–371.

Stolier, R. M., Hehman, E., Keller, M. D., Walker, M., & Freeman, J. B. (2018). The conceptual structure of face impressions. *Proceedings of the National Academy of Sciences*, *115*(37), 9210–9215. https://doi.org/10.1073/pnas.1807222115

Talamas, S. N., Mavor, K. I., Axelsson, J., Sundelin, T., & Perrett, D. I. (2016). Eyelid-openness and mouth curvature influence perceived intelligence beyond attractiveness. *Journal of Experimental Psychology: General*, *145*(5), 603–620.

Todorov, A., Olivola, C. Y., Dotsch, R., & Mende-Siedlecki, P. (2015). Social attributions from faces: Determinants, consequences, accuracy, and functional significance. *Annual Review of Psychology*, *66*, 519–545.

Tskhay, K. O., & Rule, N. O. (2013). Accuracy in categorizing perceptually ambiguous groups: A review and meta-analysis. *Personality and Social Psychology Review*, *17*(1), 72–86.

Wilson, J. P., & Rule, N. O. (2015). Facial trustworthiness predicts extreme criminal-sentencing outcomes. *Psychological Science*, *26*(8), 1325–1331.

Xie, S. Y., Flake, J. K., Stolier, R. M., Freeman, J. B., & Hehman, E. (2021). Facial impressions are predicted by the structure of group stereotypes. *Psychological Science*.

Author Biographies

Maryam Bin Meshar is a doctoral student at New York University. Her research focuses on social cognition and judgment.

Ryan M. Stolier is a user experience researcher at Facebook. His research focuses on human cognition and behavior.

Jonathan B. Freeman is an Associate Professor of Psychology and Neural Science at New York University. His research examines the cognitive and neural mechanisms underlying person perception.

Handling Editor: Lowell Gaertner