The Creepy, the Bad and the Ugly: Exploring perceptions of moral character and social desirability in uncanny faces

Cite this article

Rosa, A.O.L., Villacampa, J., Corradi, G. *et al.* The creepy, the bad and the ugly: exploring perceptions of moral character and social desirability in uncanny faces. *Curr Psychol* (2021). <https://doi.org/10.1007/s12144-021-01452-w>

Abstract

We used implicit and explicit measures to study whether “real” uncanny faces (by faces of Botox users and very ugly people) will be associated with perceptions of bad moral character and social avoidance. Implicit measures showed that uncanny faces were more strongly associated with negative aesthetic evaluations (“ugliness”) than with negative moral evaluation (“moral disgust”). At the explicit level, participants preferred greater social distance from uncanny faces than from neutral faces and inferred that they shared fewer moral values with uncanny faces than with neutral faces. Contrary to our hypotheses, only Ugly faces (but not Botox faces) were perceived as more likely to commit behaviors that indicate bad moral character. However, when this analysis was restricted to “sick” immoral actions, Botox faces were perceived as more likely to be engage in these kind of behaviors than neutral faces. Although exploratory in nature, this investigation suggest that ugliness (more than creepiness) may be the crucial evaluative dimension underlying rapid moral inferences from faces.

Keywords: Uncanny valley, face perception, moral character, social perception, attractiveness

1. Introduction

Imagine that we are in the middle of a desolate place and suddenly a taxi driver offers to take us to our destination. Probably, the first thing we will do is to "scan" the taxi driver's face for information on his possible intentions ("this taxi driver doesn't seem very trustworthy"). Of course, the reliability index of this "scanner" is not guaranteed: people infer social qualities from faces regardless of their diagnostic validity (Oosterhoh & Todorov, 2008). Although there is a long tradition in research on trait inferences from faces (for a review see Olivola & Todorov, 2017), the study of how people infer moral character (i.e., the moral dimension of the personality) in contexts of social ambiguity and, in particular, the study of which cues are perceived as indicative of moral character still deserves more attention, due to its ecological relevance.

Indeed, inferring social judgments from faces is a ubiquitous practice with important social consequences. This research aims to offer new insights into the psychological mechanisms that underlie the formation of certain social evaluations ("he has the face of a bad person", "he looks like a pervert, he must be guilty"). For instance, the fact that people draw multiple social inferences from faces implies a remarkable puzzle: does the feeling that something is "not right" with a human face functions as an automatic signal that something is "not right" with the morality of that person? The potential consequences of this sociomoral heuristic may harm people with facial deformity or victims

of surgical malpractice. For instance, jurors may be influenced by the facial appearance of a defendant.

The present study addresses this research question by exploring the potential link between perception of uncanny faces (by faces of real people) and inferences of moral character and social desirability. We built on previous research showing that (i) facial attractiveness (and facial unattractiveness) bias inferences of social desirability in a congruent manner, and that (ii) unnatural facial expressions may trigger perception of deviant morality, to make two concrete predictions: (H1) uncanny faces (i.e., perceived as both ugly and creepy) will evidence stronger associations with perceptions of bad moral character and social avoidance than neutral faces, and, more specifically, (H2) faces of “Botox victims”, which are supposed to be perceived as specially “unnatural” and creepy (Pollick, 2010; Smith, 2014), will evidence the stronger link with perceptions of deviant morality. Further, we combined implicit and explicit measures in order to obtain a more detailed look at these hypotheses. Given the exploratory nature of this research, we did not, however, make any specific prediction regarding the relationship between implicit and explicit measures.

2. Moral character & first impressions

Traditionally, research on how people gain first impressions of social targets have identified warmth (i.e., a social dimension that reflect traits related to perceived intent) and competence (i.e., a cognitive dimension that capture traits that are related to perceived ability) as the two basic dimensions driving impression formation (Fiske, Cuddy & Glick, 2007). Although moral character information was previously considered as a part of the warmth dimension (Cuddy, Fiske, & Glick, 2008), other studies suggest that moral character may be a separate source of information that plays the most crucial role in driving

impressions of social targets, being the most important dependent variable in person perception research (Goodwin, Piazza, & Rozin, 2014; Goodwin, 2015).

The importance of moral character in person perception seems to be rooted in both functionalist and symbolic reasons. From a functionalist point of view, moral character evaluation is crucial to assess the quality of a potential social interaction, for instance, by determining whether a novel target is likely to be harmful or helpful to the evaluator. Indeed, research from evolutionary psychology suggest that the human mind may evolved a cheater detection mechanism to scan others for information that may signal intentions to cooperate or defect (Buchner, Bell, Mehl, & Musch, 2009; Fetchenhauer, Groothuis, & Pradel, 2010; Verplaetse, Vanneste, & Braeckman, 2007). From a symbolic perspective, some authors have proposed that moral traits are the most essential part of identity and the self (“what it means to be human”; Strohminger & Nichols, 2014).

Indeed, the necessity of making moral inferences about novel targets is so important for person perception that it may have led to the evolution of an automatic pathway of moral inference (Olivera-La Rosa, Arango-Tobón & Ingram, 2019). In recent years, moral behaviors (Uhlmann et al., 2015) and attitudes (Bocian et al., 2018), emotional displays (Szcurek, Monin, & Gross, 2012), and the perception of uncanny faces (Olivera-La Rosa, 2018) have been suggested as perceptual cues indicative of moral character. There is some evidence that disgust responds to indications of bad moral character (Giner-Sorolla & Chapman, 2017). In particular, some authors suggest that moral disgust is a reaction to a subclass of egregious moral offenses, those that reveal that an individual is lacking the normal human motives (people and behaviors that are morally “sick”). At least in Westerners, moral disgust may respond to the lower boundary of the category of humanity:

those actions that expose people moving down, who "degrade" themselves and elicit disgust in others(e.g., by exploiting the vulnerable; Rozin et al., 2008).

3. Face-ism & moral inferences

3.1. Beautiful-is-Good, Ugliness-is-Bad?

There is a long tradition in social psychology of research on “face-ism” (i.e., the tendency to stereotype people based on their facial appearance). Decades of research has shown that the face is a central source of social information, and that people infer multiple social judgments about a target person from minimal facial cues (Olivola & Todorov, 2017). Most of the time, these judgments can be described as automatic, in the sense that they occur very rapidly and may extend to early (preconscious) stages of perception (Stewart et al., 2012).

In this context, several studies have shown that facial attractiveness influences social inferences in a congruent manner: attractive people are perceived to possess more socially desirable personalities, better social skills and higher moral standards (e.g., “attractive people are friendlier than unattractive people”, Dion, Berscheid, & Walster, 1972; Eagly et al., 1991; Maestripieri, Henry, & Nickels, 2017). This Beautiful-is-Good stereotype is ubiquitous in social cognition and has been shown to bias social judgments in various domains (Eagly, Ashmore, Makhijani, & Longo, 1991). Of special interest for the present research is the finding that, most often, unattractiveness is a disadvantage more than attractiveness is an advantage in several domains of social judgment (e.g., altruism, intelligence; Griffin & Langlois, 2006). According to the authors, this Ugly-is-Bad stereotype is largely based on a negative bias: unattractive faces elicit stronger effects on social judgments than attractive faces.

Further, this sensitivity to negative valence in unattractive faces may be explained as a result of these faces as being perceived as less “face-like”, more ambiguous, and as a result, more difficult to categorize (Griffin & Langlois, 2006). In this vein, research on face perception has shown that the human visual system is especially sensitive to unusual facial features (Hassin & Trope, 2000; Nesse, 2005; Simpson, Varga, Frick, & Frigaszy, 2011). This perceptual process seems highly automatic, leading any incongruent facial features to trigger a hyper-sensitive perceptual alarm system (signaling that “something is wrong” with the target individual). For instance, there is evidence that small perceived deviations from normal human appearance produce large prediction errors in brain regions associated with human face recognition (Chattopadhyay & MacDorman, 2016). Consistent with these findings, recent research suggest that a “disfigured is bad” stereotype exists, wherein people attribute negative social attributes to individuals with facial disfigurement (Hartung et al., 2019; Jamrozik, Oraa Ali, Sarwer, & Chatterjee, 2019).

3.2. Creepiness, the uncanny valley and moral perception

The role of facial ambiguity in social perception is congruent with recent research on creepiness and the uncanny valley. With regards to the former, some authors argue that creepiness is an unpleasant emotional response that arises from some ambiguity in a potential threat and may rely on information gathered from certain facial cues of a target person (McAndrew & Koehnke, 2016; Watt, Maitland, & Gallagher, 2017). Creepiness may be better understood in the context of social interactions, being associated with violation of social norms and appraisals of untrustworthiness, which suggest that creepiness may be an adaptive response directed to increase vigilance during periods of social uncertainty (McAndrew & Koehnke, 2016). Indeed, the link between ambiguity and social danger is supported by evolutionary theory (Becker et al., 2011; Oosterhoh & Todorov,

2008) and neuroimaging evidence, which has shown a link between social anxiety and greater activation of the amygdala in response to ambiguous stimuli (Griffin & Langlois, 2006; Thomas, Drevets, Whalen, et al., 2001).

Complementing these findings, research on the uncanny valley hypothesis (Mori, 1970/2005) suggest that “unnatural” faces may trigger perceptions of deviant morality. Briefly, the hypothesis states that entities which are quite close to looking human, but don’t appear to be completely human, can produce negative feelings in an observer: the more humanlike an entity looks, the more positively it is perceived, until a point is reached at which it appears “too” humanlike, eliciting an unpleasant emotional response: the uncanny feeling (UF). Interestingly, the possibility that the UF relies on facial cues, which in turn entail moral attributions, has been suggested by some authors. For instance, some studies showed that the perception of affective deviants (i.e., targets who exhibit affective displays that deviate from normative expectations) may trigger inferences of deviant morality (Szcurek et al., 2012; Tinwell, Nabi, & Charlton, 2013), which is congruent with an evolutionary account of face evaluation: the need to make rapid guesses about another agent’s harmful intentions justifies the fast and automatic (but not necessarily 100% accurate) inference of moral judgments from faces (Oosterhoh & Todorov, 2008).

In this vein, Tinwell, Nabi, and Charlton (2013) found that the perception of psychopathic traits (i.e., aberrant facial expressions) was associated with viewers’ perception of the uncanny in characters, which suggests that some facial expressions trigger the UF. Based on these findings, a recent review suggests that the UF should be understood as an “embodied response directed to impair sociality and that it does so by triggering perceptions of psychopathy and process of dehumanization, which in turn entail moral attributions” (Olivera-La Rosa, 2018, p.44). Indeed, a recent study explored these

hypotheses by using implicit measures, finding evidence for a slight implicit association of the UF with moral disgust (relative to fear), but not evidence of an implicit association between the UF and psychopathy (Villacampa, Ingram, Corradi & Olivera-La Rosa, 2019). As a result, more empirical evidence is needed to disambiguate the relationship between uncanny faces and perceptions of deviant morality.

The purpose of this study is to fill that gap in the literature. In particular, previous research on perceptions of uncanniness and creepiness have largely focused on non-human targets (e.g., androids; for an exception see Watt et al., 2017). Therefore, whether uncanny faces of real human targets are associated with a “different” morality (which lack the same moral grounds as most people) remains an empirical question. With this aim, we used two types of real world human faces that can be related to the uncanny: faces of “very ugly” people and faces of “Botox victims” (victims of plastic surgery). Previous research suggest that “attractiveness” and “eeriness” are two dimensions involved in the experience of the uncanniness (Bartneck, Kulic, Croft, & Zoghbi, 2009; Ho & MacDorman, 2010). Indeed, the use of Botox faces is of special interest for this research, given their characteristic feeling of “unnaturalness” and their references in the uncanny literature (Pollick, 2010; Smith, 2014). To the best of our knowledge, this is the first study to assess the link between Botox faces and perceptions of deviant morality. Crucially, although in this manuscript we use the term “uncanny” in order to describe the emotional and cognitive response to some type of stimuli, in this study we did not measured the “human-likeness” dimension of our face stimuli and, therefore, we did not directly assess the “strong” version of the uncanny valley hypothesis as suggested by Mori (1970/2005).

4. Method

4.1. Participants

We recruited 174 participants (92 women; mean age = 22.4, SD = 1.41). All participants were undergraduate students at one of two universities in Colombia, who were invited to join the experiment as a part of their Psychology course credits. We recruited all of them via internal email, and they provided written consent in accordance with ethical procedures approved by the two universities. All participants had normal or corrected-to-normal vision and were over 18 years of age.

4.2. Material & Procedure

4.2.1 Implicit measures

We displayed the stimuli on a 20-inch screen (60Hz screen refresh rate) with a PC running OpenSesame v. 3.0.7 (Mathôt, Schreij, & Theeuwes, 2012) on Windows 8 (Microsoft Corporation). In order to assess the implicit associations of face stimuli, we used the Single-Target Implicit Association Test (ST-IAT; Karpinski & Steinman, 2006). In the ST-IAT, participants are asked to categorize each presented stimulus as quickly and accurately as possible. In a typical procedure, the experimenters assess the association of the target category (e.g., Botox faces) towards each pole of a bipolar attribute category (e.g., “ugliness” vs. “moral disgust”) through a series of categorization tests requiring prompt responses. The reasoning behind the ST-IAT is based on response interference or compatibility. The easiness of the task is evaluated through response latencies (i.e., reaction times; RTs): shorter latencies indicate easier stimulus/category assignment (i.e., less interference and more compatibility), suggesting a stronger implicit association (Bohner et al., 2008).

In this study, we used the ST-IAT to assess implicit associations between target faces (“Botox”, “Ugly” and “Neutral”) and two additional categories: “ugliness” (aesthetic evaluative dimension) and “moral disgust” (moral evaluative dimension). Face stimuli were

selected after a pilot study, conducted with the aim of determining which faces received higher ratings on the ugliness and creepiness dimensions. A crucial aspect of this research was to adopt an ecological approach to the role of uncanny faces in social perception. Therefore, we wanted to ensure that our face stimuli would be perceived as uncanny in the real world. On this basis, we decided to focus on male faces given their greater propensity to elicit the creepiness response (McAndrew & Koehnke, 2016; Watt et al., 2017).

To further enhance ecological validity, 60 additional participants rated an initial pool of 21 faces depicting neutral emotional expressions (7 ugly faces, 7 Botox faces and 7 neutral faces). Fourteen faces were selected from the Internet (Botox and ugly faces;) and 7 neutral faces were selected from the Chicago Face Database (Ma, Correll, & Wittenbrink, 2015) by using two 7-point Likert scales: the aesthetic scale (from 1: “very ugly” to 7: “very handsome”) and the creepiness scale (from 1: “very creepy” to 7: “not creepy at all”). Therefore, each face was rated in two independent scales. Our selection criteria for the “ugly” category was to include the three faces with lower ratings in the aesthetic dimension (i.e., more ugly). For the “Botox” category, we selected the three faces with lower ratings in the creepiness dimension (i.e., more creepy). Neutral faces were selected from The Chicago Face Database (Ma, et al., 2015) under the criteria that they showed middle ratings in the aesthetic dimension and high ratings in the creepiness dimension (i.e., “not creepy at all”). However, it is worthy to mention that, although the main criteria for selecting “ugly” and “Botox” faces were their low ratings in the aesthetic and creepiness scales (respectively), both type of faces also evidenced low ratings in both dimensions (see Appendix A).

“Ugliness” and “moral disgust” categories consisted of three ugliness related words (e.g., feo [ugly]) and three “moral disgust” words (e.g., pervertido [perverted]), respectively. Words stimuli were selected after a pilot pretest in which 19 additional

participants rated an initial pool of 13 words (5 ugliness-related words and 8 moral disgust words) according to their degree of association with their respective category. Based on this pretest, we chose three unambiguously identifiable stimuli (item) for each category (Appendix B).

Following the same procedure as Bluemke and Frieze (2008), each stimulus was presented at least twice, adding up to 21 trials per combined block. Face stimuli, coupled and uncoupled aesthetic /moral stimuli occurred in a ratio of 6:6:9 trials (Table 1). The categorization task started with 12 trials for the training block, prior to the first combined block. This training block considered only two categories (ugliness and moral disgust, but not target faces) and the scores obtained were not explored in the analysis of the results. The order for the three versions of the ST-IAT (Botox, Ugly and Neutral) was counterbalanced between participants, and the order of the item/category assignment trials was randomized within participants (Table 1).

Table 1. Category assignment and stimuli proportions across ST-IAT blocks for an example participant.

4.2.2. Explicit measures

Once participants completed the implicit task, they were asked to respond to three explicit measures related to social perception. The order for the three explicit measures was counterbalanced between participants. Before the task, we emphasized that we were asking participants for their first reactions and that it was important to respond quickly.

Participants responded to a variation of the social distance scale (Bogardus 1933; adapted by Szczurek et al., 2012) which we had modified to be appropriate for the Colombian context (e.g., “I would be completely comfortable if the target was a member of my neighborhood”); higher punctuation indicate more social distance. In addition, participants

completed the shared moral values scale (Szczurek et al., 2012). In this scale, participants indicated how much they felt the face target shared their moral values (“I feel the target individual shares most of the same moral values as me”), using a 6-point scale ranging from 1 (strongly disagree) to 6 (strongly agree). Perception of moral character was measured with predictions of the target face’s behaviors (for a similar approach see Bocian et al; 2018). In this task, participants indicated which of the nine target face was more likely to commit the immoral behavior depicted in the screen. For this purpose, we selected five immoral behavior associated with moral disgust, that is, moral assessments that "degrade" or "dehumanize" the perpetrator (“Enjoy child pornography”) and one “ordinary” immoral behavior (“Decide to take a wallet that he found lost at a restaurant table”), for a total of six immoral behavior (6 trials). (Appendix C).

5. Results

5.1. Implicit results

Based on past ST-IAT research (Bluemke & Frieze, 2008), we omitted participants who committed 30% or more errors (i.e., incorrect responses in the item/category assignment) in at least one of the three ST-IATs. Therefore, our final analysis was based on a sample size of 136 participants (74 women). Likewise, we recoded the trial latencies (RTs) that were below 300 ms (0.07% of the total trials of the task) or above 3,000ms (3.14%) to the respective values (i.e., Rts below 300ms or above 3000ms were recoded to 300ms and 3000ms, respectively) and replaced the 3.61% trials that were errors by the block mean of correct latencies plus 600ms (for a detailed description of the rationale behind this procedure, see Richetin et al., 2015).

ST-IAT effects were calculated on the basis of the attribute trials only, by computing the widely used D-measure algorithm (i.e., a computed score from attribute trials indicating

an implicit association for one of the two attribute categories resulting from quicker response times if the target stimuli had to be sorted into the respective category; Greenwald et al., 2003). To explore ST-IAT effects we calculated three separate D-measures. This analysis revealed that both ugly faces and Botox faces were more strongly associated with ugliness (aesthetic evaluation) than with bad moral character (moral evaluation) ($D=0.193$ and $D=0.122$, respectively). Nevertheless, based on the conventional level of strength adopted by previous research (Blanton et al., 2015)¹, this can only be considered a slight implicit association. No other D-measure reached the conventional level of strength (see Appendix D).

5.2. Explicit results.

As predicted, t-test of pairwise comparisons using Bonferroni-adjusted alpha levels revealed that both ugly faces ($M=5.96$) and Botox faces ($M=5.92$) were more “socially uncomfortable” (i.e., elicited greater desire of social distance) than neutral faces ($M=4.07$), $Z= 9.61$, $p < .001$ and $Z=9.15$, $p < .001$; respectively. With regard to moral perception, results from the “shared moral values” scale showed that both ugly faces and Botox faces were rated as more morally “different” (i.e., less perceived shared moral values) than neutral faces, $Z= 6.59$, $p < .001$, and $Z=6.39$, $p < .001$, respectively). In both tasks, we found no significant differences between Botox faces and ugly faces.

In an unpredicted finding, the analysis of the “faces and moral character” with χ^2 test of association revealed that ugly faces were more likely to be rated as having bad moral character than Botox faces and neutral Faces, $\chi^2 (2, N = 816) = 245$, $p < .001$ (Ugly faces

1 In the words of Blanton and colleagues (2014, p.3): “Although no explicit rationale was ever published, cut-values centered around the D values of 0.15 (“slight preference”), 0.35 (“moderate preference”), and 0.64 (“strong preference”) were adopted. These same cut-points are used in the other IAT measures used to assess psychological states (Anthony Greenwald, personal communication, August 2002).”

= 54.41 %, Botox faces = 23.41 % and Neutral faces = 22.18%). Therefore, we found no significant differences in moral character ratings between Botox faces and neutral faces. However, it is worthy to mention that when the analysis was restricted to the moral disgust items (items 1-5, see Appendix C) significant differences were found between Botox faces and neutral faces ($\chi^2(1, N = 317) = 9.11, p = .003$, Botox faces = 56.15 % and Neutral faces = 43.85%).

6. Discussion & Conclusion

This investigation was designed to extend previous research on face perception of uncanny targets and moral inferences. In particular, we combined explicit and implicit measures to explore the potential link between perception of uncanny faces and inferences of moral character and social desirability. Based on previous research showing that: (i) aesthetic features bias inferences of social desirability and that, (ii) the UF may be associated with perceptions of deviant moral character, we predicted that uncanny faces will evidence stronger associations with appraisals of bad moral character and will elicit a greater desire of social avoidance than neutral faces. Further, we hypothesized that Botox faces, which are supposed to be perceived as more “unnatural” and creepy, will evidence the stronger link with perceptions of deviant morality.

Our results supported our first hypothesis. The analysis of the explicit tasks revealed that participants preferred greater social distance from uncanny faces than from neutral faces and inferred that they shared fewer moral values with uncanny faces than with neutral faces. Interestingly, participants evaluated ugly faces (but not Botox faces) to be more likely to be involved in actions that evidenced bad moral character, suggesting that ugliness, more than creepiness, may be the crucial dimension driving unfavorable sociomoral evaluations from faces.

Complementing this finding, we found that uncanny faces were strongly associated with negative aesthetic evaluations (“ugliness”) than with negative moral evaluations (“moral disgust”). This was not the case for neutral faces. This result suggests that, at least implicitly, “ugliness” seems to be the dominant evaluative dimension when it comes to quickly categorize novel face stimuli into aesthetic and moral categories. Indeed, although we expected that Botox faces (given its lower ratings in the creepiness dimension) would be strongly associated with bad moral character than ugly faces, it is worthy to underline that both type of face stimuli evidenced low ratings in ugliness and creepiness², which difficult the interpretation of these findings in terms of their inner dimensions.

Overall, these findings are consistent with previous research on the Beauty-is-Good stereotype, showing that attractiveness (and ugliness) increase (and decrease) the perception of socially desirable traits (Dion, Berscheid, & Walster, 1972; Eagly et al., 1991). This claim is consistent with research showing that evaluation of aesthetic and moral attributes is supported by partly overlapping neuro-cognitive mechanisms (Tsukiura & Cabeza, 2011; Zaidel & Nadal, 2011), suggesting that physical and moral trait attributes may be coded along a single evaluative dimension by a shared brain circuit (Ferrari et al., 2017). More specifically, our results support Griffin and Langlois’s (2006) research on the Ugly-is-Bad stereotype and the effects of negative bias on social judgment from faces (see also Hartung et al., 2019; Jamrozik et al., 2019) . Interestingly, the authors suggest that ugly faces and other social stimuli may be perceived in a threatening manner because of their ambiguity.

2 In particular, only Face #2 (Ugly category) showed middle values in the creepiness dimension (Appendix A).

In this vein, these findings provide general support for our theoretical prediction regarding the role of the uncanniness in moral cognition, which also assumes that stimulus ambiguity plays a central role in negative social evaluations (McAndrew & Koehnke, 2016; Tinwell, Nabi, & Charlton, 2013; Olivera-La Rosa, 2018; Watt, Maitland, & Gallagher, 2017). In fact, previous research on the uncanny valley has suggested that those stimuli that are located at the category boundary between “artificial” and “human” are perceived as ambiguous and difficult to process, generating the UF and social anxiety (Burleigh, Schoenherr, & Lacroix, 2013; Yamada, Kawabe, & Ihaya, 2013; for a review see Wang, Lilienfeld, & Rochat, 2015).

Further, given that human faces are the primary focus of scrutiny in social interactions (Izard, 1994), and that evolution would have erred on the side of caution to minimize false-positive errors (at the expense of increasing false-negative errors; Nesse, 2005), difficulties in reading intentions in ambiguous faces may signal potential social danger in the context of unfamiliar social interactions (Becker et al., 2011). Congruently, research on affective learning demonstrate that people learn the affective value of other people very quickly and under minimal learning conditions, being able to associate the affective value of a target’s moral behavior with their face in a few seconds (Bliss-Moreau, Barrett, Wright, 2008; Ferrari, Oh, Labbree & Todorov, 2020; Todorov & Olson, 2008). In this vein, research on the cheater detection mechanism also suggests that we have evolved an intuitive ability to quickly and easily draw inferences on whether someone is a cheater or a cooperator (Buchner, Bell, Mehl, & Musch, 2009; Fetchenhauer, Groothuis, & Pradel, 2010; Verplaetse, Vanneste, & Braeckman, 2007).

Due to the exploratory nature of this study, there are some important limitations. First, the slight implicit associations between uncanny faces and ugliness may be taken

with caution as it must be remembered that this was only a small effect. It may be argued that the distinction between the “ugly” and “creepy” category is a subtle one, and that asking participants to quickly categorize stimuli into these categories may be resulted particularly difficult (relative to the categorization of the stimuli into the typical “good” vs. “bad” IAT procedure). An alternative explanation to this finding is that “moral disgust” items were largely sexual in nature (e.g., “perverted”), which may influenced the categorization task. Indeed, this finding is consistent with previous implicit attempts to capture the UF, which indicate that IAT-like procedures may not be valid measures for uncanny stimuli (as suggested by Wang and Rochat, 2017, based on Zlotowski et al.’s, 2015, results; see also Villacampa et al., 2019). The results of the present study support this suggestion.

Second, the finding that ugly faces (but not Botox faces) were perceived as more likely to commit actions that revealed bad moral character was somewhat surprising. We may hypothesize that neutral faces were perceived as displaying no affective expression (“flat affect”) towards immoral vignettes, which has been previously associated with perception of social deviants (Szczurek et al., 2012). It is worth mentioning that participants’ explicit ratings of creepiness showed that all neutral faces were rated as “not creepy”, which suggests that, at least at an explicit level, our neutral stimuli were not perceived as uncanny. However, it may be that participants’ social inferences about the neutral faces were ultimately influenced by these faces being displayed alongside immoral behaviors. Future studies should control for displays of “flat affect” in the context of social judgments.

Interestingly, when predictions of target’s behaviors were based only on statements related with moral disgust (i.e., moral behaviors that "dehumanize" the perpetrator), Botox

faces were perceived as more likely to commit these behaviors (relative to neutral faces). At a first glance, this finding would suggest that the link between creepiness and perceptions of bad moral character may be more salient in the case of “sick” immoral actions, that is, those actions that reveal that an individual morality is lacking the same moral grounds as most people. This finding may be interpreted by considering the cultural context of the participants, most being from Medellin (Colombia), a city in which delinquency rates are high and thereby may facilitate a “hyper-sensitive” perceptual association with delinquents. Crucially, it may be argued that participants perceived our set of ugly faces as depicting people of low socioeconomic status, who might be more likely to be involved in criminal activity. On the other hand, “Botox people” may be associated with high socioeconomic status, which may explain why these faces were not associated with “ordinary” moral violations (e.g., stealing), but only with a more “sick” type of behavior (“e.g., pedophilia). However, given that we only include one “ordinary” moral violation in this task, this hypothesis is highly speculative and requires further empirical study.

In line with this finding, our results suggest that ugliness may be more important than creepiness when making moral inferences about novel faces. However, this later conclusion should be taken with caution because of the methodological difficulty in selecting stimuli: both “ugly” faces and “Botox” faces were rated as very ugly and also as creepy. Indeed, we adopted an “ecological” approach to the study of moral inferences from uncanny faces, which required to use “real” uncanny face stimuli that were not available in any existing face database. This was an interesting methodological challenge, given that our pretest revealed that ugliness and creepiness seem to be overlapping evaluations. Therefore, one interpretation of our findings could be that they rely on a general affect-based associative mechanism (negative valence) and not necessarily on trait inferences from

uncanny faces. Although we cannot dismiss this possibility, it does not seem likely because we found that only ugly faces (but not Botox faces) were significantly associated with immoral behaviors.

Similarly, we found some difficulties with the final selection of the ugliness-related Spanish words, which eventually limited stimulus selection to only three stimuli per category. However, previous research showed that IAT's with only two items representing each attribute and category showed almost the same effects than those using eight items per attribute and category (Nosek, Greenwald & Banaji, 2005). Future research should address these limitations, for instance, by using negative valence faces that are not related with the uncanny (e.g, unattractive faces that are not creepy or faces depicting negative emotions) and that also allow to use a larger set of stimuli for each category. Furthermore, we believe that future research should explore whether facial ambiguity functions as a perceptual signal of bad moral character, for instance, by studying the potential link between androgynous faces (by real people) and moral perception. We also call for future research to extend our study by investigating the role of uncanny female faces in social perception.

In sum, our study showed that uncanny faces were perceived as less socially desirable and more different (in terms of their moral evaluations) than neutral faces. Botox faces and ugly faces elicited a similar pattern of social inferences, except when these faces were related with immoral behaviours. However, this perceptual pathway seemed to rely more on aesthetic evaluations (ugliness) than on more “social” evaluations (creepiness), possibly indicating that the Ugly-is-Bad stereotype plays a crucial role in making rapid moral inferences from novel faces.

Conflict of Interest. All the authors declare no conflict of interest.

Ethical Approval. All procedures performed in the studies involving human participants were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed Consent. Informed consent was obtained from all individual participants included in the study.

Data Availability Statement: The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- Bartneck, C., Kulic, D., Croft, E., & Zoghbi, S. (2009). Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics*, 1(1), 71–81.
- Becker, D. V., Mortensen, C. R., Ackerman, J. M., Shapiro, J. R., Anderson, U. S., Sasaki, T., ... & Kenrick, D. T. (2011). Signal detection on the battlefield: Priming self-protection vs. revenge-mindedness differentially modulates the detection of enemies and allies. *PloS one*, 6(9), e23929.
- Blanton, H., Jaccard, J., & Burrows, C. N. (2015). Implications of the implicit association test D-transformation for psychological assessment. *Assessment*, 22(4), 429–440. DOI: 10.1177/1073191114551382
- Bliss-Moreau, E., Barrett, L. F., & Wright, C. I. (2008). Individual differences in learning the affective value of others under minimal conditions. *Emotion*, 8(4), 479–493.

- Bluemke, M., & Frieze, M. (2008). Reliability and validity of the Single-Target IAT (ST-IAT): Assessing automatic affect towards multiple attitude objects. *European Journal of Social Psychology*, 38, 977-997. doi:10.1002/ejsp.487
- Bocian, K., Baryła, W., Kulesza, W., Schnall, S., & Wojciszke, B. (2018). The mere liking effect: Attitudinal influences on attributions of moral character. *Journal of Experimental Social Psychology*, 79(2018), 9–20. <http://doi.org/10.1016/j.jesp.2018.06.007>
- Bogardus, E. S. (1933). A social distance scale. *Sociology and Social Research*, 17, 265–271.
- Bohner, G., Siebler, F., González, R., Haye, A., & Schmidt, E. A. (2008). Situational flexibility of in-group-related attitudes: A single category IAT study of people with dual national identity. *Group Processes & Intergroup Relations*, 11, 301- 317.
doi:10.1177/1368430208090644
- Buchner, A., Bell, R., Mehl, B., & Musch, J. (2009). No enhanced recognition memory, but better source memory for faces of cheaters. *Evolution and Human Behavior*, 30(3), 212 – 224. <https://doi.org/10.1016/j.evolhumbehav.2009.01.004>
- Burleigh, T. J., Schoenherr, J. R., & Lacroix, G. L. (2013). Does the uncanny valley exist? An empirical test of the relationship between eeriness and the human likeness of digitally created faces. *Computers in Human Behavior*, 29, 759–771.
- Chattopadhyay, D., & MacDorman, K. F. (2016). Familiar faces rendered strange: Why inconsistent realism drives characters into the uncanny valley. *Journal of Vision*, 16(11):7, 1-25.

Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2008). Warmth and competence as universal dimensions of social perception: The stereotype content model and the BIAS map.

Advances in Experimental Social Psychology, 40, 61–149. doi:10.1016/S0065-2601(07)00002-0

Dion, K., Berscheid, E., & Walster, E. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, 24(3), 285-290.

Eagly, A. H., Ashmore, R. D., Makhijani, M. G., & Longo, L. C. (1991). What is beautiful is good, but...: A meta-analytic review of research on the physical attractiveness stereotype. *Psychological bulletin*, 110(1), 109-128.

Ferrari, C., Oh, D.W., Labbree, B., & Todorov, A., (2020). Learning the affective value of people: More than affect-based mechanisms. *Acta Psychologica*, 203, 103011

Fetchenhauer, D., Groothuis, T., & Pradel, J. (2010). Not only states but traits—Humans can identify permanent altruistic dispositions in 20's. *Evolution and Human Behavior*, 31(2), 80 – 86. <https://doi.org/10.1016/j.evolhumbehav.2009.06.009>

Fiske, S. T., Cuddy, A. J. C., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in Cognitive Sciences*, 11(2), 77-83.

<http://dx.doi.org/10.1016/j.tics.2006.11.005>

Giner-Sorolla, R., & Chapman, H. A. (2017). Beyond purity: Moral disgust toward bad character. *Psychological Science*, 28(1), 80–91. <https://doi.org/10.1177/095679761667319>

Goodwin, G. P. (2015). Moral Character in Person Perception. *Current Directions in Psychological Science*, 24(1), 38–44. <https://doi.org/10.1177/0963721414550709>

Goodwin, G. P., Piazza, J., & Rozin, P. (2014). Moral character predominates in person perception and evaluation. *Journal of Personality and Social Psychology*, 106(1), 148-168.

<http://dx.doi.org/10.1037/a0034726>

Greenwald, A. G., Nosek, B. A., & Banaji, M. R. (2003). Understanding and using the implicit association test: I. An improved scoring algorithm. *Journal of Personality and Social Psychology*, 85, 197-216. doi:10.1037/0022-3514.85.2.197

Griffin, A.M., & Langlois, J.H. (2006). Stereotype directionality and attractiveness stereotyping: Is beauty good or is ugly bad? *Social Cognition*, 24, 212–246

Hartung, F., Jamrozik, A., Rosen, M.E. et al. (2019) Behavioural and Neural Responses to Facial Disfigurement. *Scientific Reports*, 9, 8021 <https://doi.org/10.1038/s41598-019-44408-8>

Hassin, R., & Trope, Y. (2000). Facing faces: Studies on the cognitive aspects of physiognomy. *Journal of Personality and Social Psychology*, 78, 837–852.

Ho, C.-C., & MacDorman, K. F. (2010). Revisiting the uncanny valley theory: Developing and validating an alternative to the Godspeed indices. *Computers in Human Behavior*, 26(6), 1508–1518. <https://doi.org/10.1016/j.chb.2010.05.015>

Izard, C.E. (1994). Innate and universal facial expressions: Evidence from developmental cross-cultural research. *Psychological Bulletin*, 115, 288–299.

Jamrozik, A., Oraa Ali, M., Sarwer, D. B., & Chatterjee, A. (2019). More than skin deep: Judgments of individuals with facial disfigurement. *Psychology of Aesthetics, Creativity, and the Arts*, 13(1), 117–129. <https://doi.org/10.1037/aca0000147>

Karpinski, A., & Steinman, R. B. (2006). The single category implicit association test as a measure of implicit social cognition. *Journal of Personality and Social Psychology*, *91*(1), 16-32. doi:10.1037/0022-3514.91.1.16

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*, 1122–1135.

<https://doi.org/10.3758/s13428-014-0532-5>

Maestripieri, D., Henry, A., & Nickels, N. (2017). Explaining financial and prosocial biases in favor of attractive people: Interdisciplinary perspectives from economics, social psychology, and evolutionary psychology. *Behavioral and Brain Sciences*, *40*, E19.

doi:10.1017/S0140525X16000340

Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. *Behavior Research Methods*, *44*, 314-324.

doi:10.3758/s13428-011-0168-7

McAndrew, F. T., & Koehnke, S. S. (2016). On the nature of creepiness. *New Ideas in Psychology*, *43*, 10-15.

Mori, M. (1970/2005). The uncanny valley. (K. F. MacDorman, & T. Minato, Trans.). *Energy*, *7*, 33–35.

Nesse, R. M. (2005). Natural selection and the regulation of defenses: A signal detection analysis of the smoke detector principle. *Evolution and Human Behavior*, *26*(1), 88-105.

Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2005). Understanding and Using the Implicit Association Test: II. Method Variables and Construct Validity. *Personality and Social Psychology Bulletin*, *31*(2), 166–180. <https://doi.org/10.1177/0146167204271418>

Olivera-La Rosa, A. (2018). Wrong outside, wrong inside: A social functionalist approach to the uncanny feeling. *New Ideas in Psychology*, 50, 38-47.

Olivera-La Rosa, A., Arango-Tobón, O.E., & Ingram, G.P. (2019). Swiping right: face perception in the age of Tinder. *Heliyon*, 5(12):e02949. Doi: 10.1016/j.heliyon.2019.e02949

Olivola & Todorov (2017). The biasing effects of appearances go beyond physical attractiveness and mating motives. *Behavioral and Brain Sciences*, 40,:e38

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, 105(32), 11087–11092.

Pollick, F. E. (2010). In search of the uncanny valley. *Lect. Note Inst. Comput. Sci. Telecomm.* 40, 69–78.doi:10.1007/978-3-642-12630-7_8

Rozin, P., Haidt, J., & McCauley, C. R. (2008). Disgust. In M. Lewis, J. M. Haviland-Jones & L. F. Barrett (Eds.), *Handbook of emotions*, 3rd ed. (pp. 757-776). New York: Guilford Press.

Simpson, E. A., Varga, K., Frick, J. E., & Frigaszy, D. (2011). Infants experience perceptual narrowing for nonprimate faces. *Infancy*, 16(3), 318-328.

Smith, P. (2014). Of ‘near pollution’ and non-linear cultural effects: Reflections on Masahiro Mori and the Uncanny Valley. *American Journal of Cultural Sociology*, 2(3), 329-347.

Stewart, L. H., Ajina, S., Getov, S., Bahrami, B., Todorov, A., & Rees, G. (2012). Unconscious evaluation of faces on social dimensions. *Journal of Experimental Psychology: General*, 141(4), 715-727.<http://dx.doi.org/10.1037/a0027950>

- Strohminger, N., & Nichols, S. (2014). The essential moral self. *Cognition*, 131(1), 159-171. <http://dx.doi.org/10.1016/j.cognition.2013.12.005>
- Szczurek, L., Monin, B., & Gross, J. J. (2012). The Stranger Effect: The Rejection of Affective Deviants. *Psychological Science*, 23(10), 1105–1111. <https://doi.org/10.1177/0956797612445314>
- Thomas, K. M., Drevets, W. C., Whalen, P. J., Eccard, C. H., Dahl, R. E., Ryan, N. D., & Casey, B. J. (2001). Amygdala response to facial expressions in children and adults. *Biological Psychiatry*, 49(4), 309-316.
- Tinwell, A., Nabi, D. A., & Charlton, J. P. (2013). Perception of psychopathy and the Uncanny Valley in virtual characters. *Computers in Human Behavior*, 29(4), 1617-1625.
- Todorov, A., & Olson, I. R. (2008). Robust learning of affective trait associations with faces when the hippocampus is damaged, but not when the amygdala and temporal pole are damaged. *Social cognitive and affective neuroscience*, 3(3), 195-203.
<https://doi.org/10.1093/scan/nsn013>
- Tsukiura, T., & Cabeza, R. (2011). Shared brain activity for aesthetic and moral judgments: Implications for the beauty-is-good stereotype. *Social Cognitive and Affective Neuroscience*, 6(1), 138–148.
- Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A Person-Centered Approach to Moral Judgment. *Perspectives on Psychological Science*, 10(1), 72–81. <https://doi.org/10.1177/1745691614556679>

- Verplaetse, J., Vanneste, S., & Braeckman, J. (2007). You can judge a book by its cover: The sequel. A kernel of truth in predictive cheating detection. *Evolution and Human Behavior*, 28(4), 260–271. <https://doi.org/10.1016/j.evolhumbehav.2007.04.006>
- Wang, S., Lilienfeld, S. O., & Rochat, P. (2015). The uncanny valley: Existence and explanations. *Review of General Psychology*, 19(4), 393–407.
- Wang, S. & Rochat, P. (2017). Human perception of animacy in light of the uncanny valley phenomenon. *Perception*, 46 (12), 1386–1411.
- Watt, M. C., Maitland, R. A., & Gallagher, C. E. (2017). A case of the “heeby jeebies”: An examination of intuitive judgements of “creepiness”. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement*, 49(1), 58–69.
- Villacampa, J., Ingram, G.P., Corradi, G., & Olivera-La Rosa, A. (2019). Applying an implicit approach to research on the uncanny feeling. *Journal of Articles in Support of the Null Hypothesis*, 16 (1), 11–22.
- Yamada, Y., Kawabe, T., & Ihaya, K. (2013). Categorization difficulty is associated with negative evaluation in the “uncanny valley” phenomenon. *Japanese Psychological Research*, 55, 20–32.
- Zaidel, D. W., & Nadal, M. (2011). Brain intersections of aesthetics and morals: perspectives from biology, neuroscience, and evolution. *Perspectives in Biology and Medicine*, 54(3), 367–380.
- Złotowski, J. A., Sumioka, H., Nishio, S., Glas, D. F., Bartneck, C., & Ishiguro, H. (2015). Persistence of the uncanny valley: the influence of repeated interactions and a robot’s

attitude on its perception. *Frontiers in Psychology*, 6, 883.

<http://doi.org/10.3389/fpsyg.2015.00883>