Moral Beauty, Inside and Out

Ryan P. Doran

Published online: 05 Aug 2020.

Submit your article to this journal ⤤

View related articles ⤤

View Crossmark data ⤤

Australasian
Association of
Philosophy

Routledge
Taylor & Francis Group

Check for updates

Moral Beauty, Inside and Out

Ryan P. Doran

University of Cambridge

ABSTRACT
In this article, robust evidence is provided showing that an individual's moral character can contribute to the aesthetic quality of their appearance, as well as being beautiful or ugly itself. It is argued that this evidence supports two main conclusions. First, moral beauty and ugliness reside on the inside, and beauty and ugliness are not perception-dependent as a result; and, second, aesthetic perception is affected by moral information, and thus moral beauty and ugliness are on the outside as well.

1. Introduction

The claim that moral beauty or ugliness exist can be formulated generally in the following way:

> *General-MBU*. If a person is morally good then, to this extent, they are beautiful; or, conversely, if a person is morally bad then, to this extent, they are ugly.

In section 2, I distinguish between three more specific forms of General-MBU, including the form that has recently been defended in an indirect fashion by reference to the use of language [Gaut 2007] and to certain data from the field of social psychology [Paris 2018]. In section 3, I argue that these indirect arguments may be resisted by sceptics of moral beauty (ugliness): the use of language and the psychological data cited may be variously deflated in terms of non-literal intent or the operation of an error-generating mechanism. In sections 4–6, I provide evidence that is robust against these deflationary challenges and shows, first, that an individual's moral character can contribute to the perceived aesthetic quality of their appearance, and, second, that moral character is in itself beautiful (ugly). As such, this paper suggests that, while beauty and ugliness are not perception-dependent, moral beauty and ugliness exist partly because moral character tends to modify our perception of appearances.[1]

2. Three Location-Specific Conceptions of General-MBU

By further specifying the location of beauty (ugliness), it can be seen that there are at least three important ways in which General-MBU could be

---

[1] Hereafter, 'moral beauty or ugliness' will be referred to as 'MBU'.

true.[2] One location-specific conception, and the one defended most recently, holds the following (adapted from Paris [2018: 643] and Gaut [2007: 120]):

> *Trait-MBU*. If a trait is a moral virtue then, to this extent, it is beautiful; or, conversely, if a trait is a moral vice then, to this extent, it is ugly.

Since someone's character is the set of their traits, this suggests a second conception:

> *Character-MBU*. If a person's character is virtuous then, to this extent, it is beautiful; or, conversely, if a person's character is vicious then, to this extent, it is ugly.[3]

Since people's characters and their traits are dispositional, and as such have no perceivable properties in themselves, this conception entails that beauty or ugliness are not perception-dependent.

This corollary suggests a third location-specific conception, which holds that perceivable effects of moral virtue (vice) can be beautiful (ugly):

> *Effect-MBU*. If a person is known to have a virtuous trait then, to this extent, their appearance can be beautiful; or, conversely, if a person is known to have a moral vice then, to this extent, their appearance can be ugly.

Unlike its dispositional sisters, the truth of this conception would not have any consequences for whether beauty or ugliness are perception-dependent. However, if true, it may suggest that aesthetic perception is cognitively penetrable to an interesting extent.

3. The Indirect Arguments for MBU

Two philosophers have recently argued indirectly for the truth of one of the dispositional conceptions—Trait-MBU.

Gaut [2007] deploys a dialectical strategy: Trait-MBU is *prima facie* true, and so should be accepted in the absence of convincing reasons to reject it. One of the principal ways in which Gaut pursues this strategy is by appealing, from the armchair, to language usage. At least some people utter locutions expressing MBU, and these should be taken to be intended literally, according to Gaut, on the ground that the two defeaters of literal usage—namely, being evidently false or true but irrelevant when taken literally—do not apply to such locutions. Unlike locutions such as 'I'm not over the moon' (obviously true but irrelevant) and 'my boss is a pig' (obviously false), Gaut suggests that locutions ostensibly expressing MBU are neither obviously true but irrelevant nor obviously false, as indicated by the fact that there is disagreement about the existence of moral beauty and ugliness.

---

[2] This taxonomy does not exhaust the logical space that can be occupied in this debate. There are other possible locations of beauty (ugliness): e.g. beauty (ugliness) could be located in someone's moral *behaviour*. Conceptions of MBU can also vary in terms of whether a narrow or broad sense of beauty (ugliness) is intended (e.g. Beardsley [1981: 505–7]). Moreover, conceptions might be quantified in different ways: it may be the case that all or only some moral traits can contribute to an individual's beauty (ugliness). Indeed, it has even been thought by some, such as Shaftesbury [1711] and Gaut [2007], that non-morally-good (-bad) traits can contribute to an individual's beauty (ugliness), although these do not strictly fall within the purview of MBU. There may also be logical links between the sense of beauty (ugliness) intended and the question of which traits contribute to an individual's beauty (ugliness). At times, Kant's narrow conception of beauty leads him to carve the beauty of traits along lines that are orthogonal to the moral-non-moral distinction: he posits that *warm* traits such as agreeableness and tenderheartedness are beautiful [1764: 22–6]. Finally, it is at least logically possible to hold the corresponding immoralist positions. To illustrate what one location-specific conception might be: if a trait is vicious (virtuous), then to this extent it is beautiful (ugly).

[3] Gaut might also hold Character-MBU: see, e.g., [2007: 115].

In contrast, Paris [2018] principally cites evidence that people who have positive characteristics are judged to be more physically attractive and that people appear to reflectively endorse this fact (see, for example, Nisbett and Wilson [1977]). Paris argues that the best explanation of these findings is Trait-MBU.

The great advantage of these indirect strategies is that they afford a way of securing the existence of MBU without having a sufficient account of beauty or ugliness, such accounts having proved elusive. Indeed, to argue for MBU from a putative account of beauty or ugliness is to risk making MBU dependent on the success of that account: if the account sinks, it may drag the argument for MBU under with it.

Setting aside those metaphysical assumptions—such as perception-dependence— that are at stake in the debate about certain conceptions of MBU, these strategies largely proceed with assumptions that both supporters and critics of MBU are likely to endorse. Those on both sides of the debate will agree that, no matter what beauty and ugliness are (for example whether they're mind-independent or mind-dependent properties, and of what specific kind), they can at least be reliably detected through the experience of the folk as reflected in their use of language and responses in relevant empirical studies, provided that certain conditions apply.

We do not tend to think that beauty and ugliness can only be accurately detected by God, for instance; and, while it is certainly thought that expertise might be relevant to the detection of artistic beauty and ugliness, it is far from clear that it is relevant to any great extent in the detection of human beauty and ugliness. Indeed, even among the few who do think that expertise is required for the detection of human beauty, it has been thought to be needed for the specific detection of *moral* beauty.[4] As a result, if anything, the indirect strategies may in fact be thought to favour scepticism about MBU.

So, arguably, the indirect strategy of appealing to the folk concepts of beauty and ugliness constitutes a robust and attractive way of settling whether MBU exist, provided that certain conditions hold. The most important such conditions have already been suggested by Gaut and Paris. The evidence provided by the folk's use of language and responses in relevant empirical studies only reveal the semantic content of the folk concepts, and come to bear on the metaphysical question of whether MBU exist, provided that: the concepts of beauty or ugliness are targeted, and that the responses are not the result of errors or pragmatic factors.

Notwithstanding the advantages of indirect strategies generally, sceptics about MBU might worry that Gaut's and Paris's forms of these strategies do not satisfy these conditions.

First, the linguistic and empirical evidence offered by Gaut and Paris might be deflated in terms of pragmatic factors such as non-literal intentions, despite the considerations that Gaut carefully proposes to guard against this possibility.

Against Gaut, sceptics might press that whether the content of locutions expressing MBU seem evidently false or true but irrelevant depends crucially on whether one holds that MBU exist, which is precisely what is at stake in appealing to such considerations. The description of an obviously ugly individual as a 'beautiful person', for example, may be taken to metaphorically express the individual's goodness, since that person is obviously not *physically* beautiful. But, equally, it may be taken to be literally true and to describe *moral beauty*.

---

[4] In the Symposium [c.370 BC: 2010b], Socrates claims that the appreciation of physical beauty leads to the cultivation of a sensitivity to higher beauties, including moral beauty.

Sceptics might also press that Gaut's restrictions on when the content of locutions are intended or understood non-literally, are best understood as the conditions that *would hold* for non-literal meaning if the parties involved were *maximally oriented towards truth-preservation*.

Suppose, *ex hypothesi*, that people do indeed aim to maximise truth-preservation in communicative exchanges, and that locutions expressing MBU are intended non-literally. In such a world, it's true that locutions such as 'she's a beautiful person' would only be used to describe moral goodness in a very limited number of contexts, such as where the formal character of the individual concerned is obviously incongruent with their moral character (for example, a hideous appearance and a virtuous character)—if indeed these locutions would be used at all in such a world.

The problem is that we are not *as a matter of fact* always maximally oriented towards truth-preservation, especially when it comes to lexical items—such as 'beautiful'—that can serve expressive functions. In exclaiming that a morally good person is beautiful, we may be concerned exclusively with vividly expressing our approval or admiration of them. Or we may simply wish to express the idea that the person has a stable positive disposition (as beauty is just such a disposition) in a florid manner. And since the issue at stake is that of how people, *as a matter of fact*, intend locutions expressing MBU, and not how they are likely to intend such locutions if they are truth-maximisers, sceptics may reasonably worry that Gaut's criteria cannot fairly be used to establish that locutions expressing MBU are intended literally.

For the same reason, participants in the studies that Paris cites may not intend their responses literally. Indeed, the problem may be worse in some such cases. For example, in presenting student participants with a video of a teacher who belittled students, and asking them to rate the teacher's attractiveness/*irritatingness*, Nisbett and Wilson [1977] seem likely to have unwittingly encouraged answering in a way that relieves pragmatic pressures.

Turning specifically to Paris's [2018] inference-to-the-best-explanation strategy, sceptics might worry that it is currently seriously constrained, and in fact actually threatened, by the available evidence.

The main studies that Paris marshals, such as those by Nisbett and Wilson [1977], can at best be thought to support only Effect-MBU rather than Trait-MBU. After all, these studies concern *physical* attractiveness, and Effect-MBU and Trait-MBU are logically independent of one another. Importantly, as a consequence, these studies do not provide support for the idea that beauty and ugliness are not perception-dependent. Setting aside this problem, sceptics might also have reasonable doubts about whether these studies even support Effect-MBU, for two further reasons.

First, it may be thought that the results of these studies are simply the result of what psychologists call 'misattribution'—a form of response-error—and so do not provide support for Effect-MBU. A longstanding program of research in social psychology has demonstrated that evaluative judgments generally, including judgments of attractiveness, are made on the basis of affective phenomena. When people are asked to make evaluative judgments, they often do so by asking themselves, often implicitly, 'how do I feel?' Making judgments in this way can lead to errors, partly because we tend to assume that the object of our attention at a given moment is the cause of our affective states at that moment. For example, Dutton and Aron [1974] found that those who had walked across a fearsome bridge judged an experimenter of the opposite sex to be more attractive than did others who had crossed a bridge that was not

fearsome. Participants who crossed the fearsome bridge wrongly inferred that the feelings of arousal were caused by the experimenter. In the case of the studies cited by Paris, participants may have found it difficult to quarantine the feelings elicited by the target's character in making judgments as to the target's physical attractiveness.[5]

Indeed, such misattribution is particularly likely to occur in these studies, for two reasons. First, while an individual's traits are invisible, they are *practically co-extensive* with something that isn't—namely, the individual's appearance. If an individual's traits lie anywhere, they are located somewhere just underneath the individual's appearance, which itself often *expresses* these traits. As a result, an individual's appearance is often likely to be the more salient object of attention and therefore will readily be assumed to be the cause of any positive or negative feelings arising from these traits. Second, the more similar an incidental affective response is to the response that is relevant for making judgments of the property concerned, the more readily misattribution will occur (see, for example, Pham [1998]). Since judgments of physical attractiveness are of a relatively thin, response-dependent property, they will be particularly susceptible to being made on the basis of any merely positive or negative affective response felt at the time.

Second, sceptics might reject the idea that the studies cited support moral *beauty per se*, on the ground that the concept of attractiveness is distinct from that of beauty. One reason for this is that attractiveness can involve elements—such as a romantic interest—that may be thought to be at best *irrelevant*, and at worst actually *anathema*, to beauty. While some have suggested that human beauty can involve a romantic interest (for example, Nehamas [2007] and Gaut [2007]), the dominant view since the eighteenth century is that beauty involves disinterestedness of some kind (for example, Kant [1890] and Beardsley [1981]). Indeed, such a view has even been held by some of those who believe that some form of MBU exist (for example, Shaftesbury [1711]).

Kant [1890] posits a particularly strong conception of disinterestedness, and of its relationship to beauty. For him, beauty is that which gives rise to delight apart from *any* interest: roughly, to the extent that we are appreciating something's beauty we wish merely to contemplate it, rather than to act in any manner towards it; and, relatedly, we are concerned with its appearance, rather than its existence.

Alongside Kant, many of the philosophers who hold weaker conceptions of disinterestedness and of its relation to beauty also believe that the desire to become romantically involved with an individual is one of the interests that cannot be involved in beauty (see, for example, Burke [1757: 162–3]; and Shaftesbury [1711: 319], although cf. [ibid.: 63]). As a result, for these philosophers, *to the extent* that judgments of physical attractiveness are made on the basis of such an interest, they are not judgments of beauty.

Indeed, there are good empirical reasons to think that physical attractiveness can involve such interests, as evinced by the 'closing time' effect, where the availability of a potential partner influences their physical attractiveness [Gladue and Delaney 1990]. As a result of the foregoing, even those who would otherwise be sympathetic to Effect-MBU might worry that the main evidence that Paris cites might not be consistent with prominent notions of disinterestedness.[6]

---

[5] Consistent with affective misattribution, Nisbett and Wilson [1977: 256] themselves interpret their findings as reflecting 'a fundamental inability to resist the affective influence of global evaluation on evaluation of specific attributes' and found that, when asked, participants denied that the target's character had influenced their rating.

[6] The desire to express moral approval (disapproval) may be another such interest.

Moreover, sceptics might also point out that it is instructive to look more broadly at how this issue is relevant to other conceptions of MBU. They might press that, in so far as beauty is being understood as mere attractiveness, General-MBU and the dispositional conceptions of MBU look at best to be trivially true and at worst to be actually false. If the issue at stake boils down to whether it is true that we are *attracted* to morally good people, characters or traits thereof, then it is difficult to see how anyone reasonable could disagree with these conceptions of MBU. For who could possibly object to the claim that we like, or feel ourselves gravitate towards, moral goodness? This, sceptics might urge, should give pause to supporters of these conceptions of MBU.

Furthermore, sceptics might advise supporters of these conceptions of MBU that it would be imprudent for them to insist on such a conception of beauty. There is evidence of an exception to the rule that moral goodness attracts, an exception that would force us to jettison these conceptions of MBU[7] and to embrace some form of immoralism about beauty. The evidence shows that some of us are attracted to morally bad characteristics at times—specifically, the so-called 'dark-triad' of psychopathy, narcissism, and Machiavellianism (for example, Carter et al. [2013]). In the face of this, the sceptic might point out that it is fortunate for supporters of MBU that attractiveness is not beauty, since sexiness can contribute towards attractiveness, and immorality can contribute towards sexiness.

Finally, it might be thought that, until there is evidence of a plausible mechanism by which moral character affects physical beauty or ugliness, we cannot be confident that *some* deflationary explanation isn't true. In connection with this, there seem to be two plausible candidate mechanisms. One possibility is that moral character merely amplifies the formal character of someone's appearance by, for example, causing us to selectively attend to existing ugliness- or beauty-making features. Another, more radical, possibility is that moral character can *penetrate* the perception of beauty or ugliness. In this case, it may be difficult to imagine *how* this could happen. After all, arguably, a change in the perception of an individual's physical beauty (ugliness) would require a change in the perception of at least one of the properties that can contribute to the realisation base of physical beauty (ugliness), including aesthetic properties such as balance (imbalance) and delicateness (coarseness) and, in turn, non-aesthetic properties such as shape, colour, size, proportion, and arrangement. And why should knowledge of someone's moral character have any effect on the perception of *these* properties?

In this article, I wish to provide support for the indirect strategies pursued by Gaut and Paris, by providing evidence to address the sceptical worries raised above. Evidence is required for MBU that is not able to be deflated in terms of pragmatic factors such as non-literal intent, or in terms of the operation of an error, such as affective misattribution or an interest that cannot be involved in beauty. Moreover, in the case of Effect-MBU specifically, evidence of a plausible mechanism is needed.

4. General-MBU

Given that General-MBU follows from, and itself implies, at least one of the location-specific conceptions, an empirical corollary of General-MBU was tested in two studies.

---

[7] At least when they are formulated conjunctively, as Gaut and Paris do.

4.1 Study 1

Materials and Method

Photographs of two moderately beautiful and two moderately ugly female models from the Radboud Faces Database [Langner et al. 2010] were used, along with two counterbalanced vignettes describing the model as morally good or bad. The model was described as, for example, caring (callous) and honest (dishonest). Participants were randomly presented with one photograph and vignette, and were asked to indicate the extent to which they agreed that the model is beautiful. As in all of the studies reported here, participants were then asked to complete manipulation and attention checks (which are especially important in light of the sceptical challenges outlined above), and to provide demographic information such as gender and sexual orientation to determine whether romantic attraction was possible. Participants in all of the studies reported here were recruited from across the United States on MTurk. Ethical approval for the studies reported here was granted by the Department of Philosophy's Ethics Committee at the University of Sheffield.

Results

279 participants took part in the study. Of these, 73 withdrew or failed the manipulation and attention checks. These participants were excluded from the analysis, leaving 206 participants (66% female, mean age: 38). A three-way ANOVA—formal character (ugly and beautiful) x moral character (good and bad) x possibility of romantic attraction (possible and not possible)—was conducted, using heteroskedasticity-consistent standard errors, as a Levene's test indicated that homogeneity of variance could not be assumed and that the numbers of participants in each group were not equal (for a summary, see, for example, Field and Wilcox [2017])[8]. The results reveal a main effect of moral character only ($F(1, 198) = 226.24$, $p<.001$, Cohen's $d = 1.89$): participants in the morally good group agreed considerably more with the judgment of beauty (Mean = 2.50, SD = 1.54) than did participants in the morally bad group (Mean = -1.95, SD = 2.21, see Figure 1(a)). There were no significant interactions, and no main effects of formal character or the possibility of romantic attraction.

Discussion

These data show that moral character affects judgments of beauty and ugliness.[9] Moral character did not influence the judgments made by participants for whom romantic attraction was possible either more, or less, than it affected the judgments made by participants for whom romantic attraction was not possible. As a result, we can conclude that it seems unlikely that the effect of moral character on judgments of beauty and ugliness can be deflated as being the result of a romantic interest in the model. Moreover, given that judgments of attractiveness have been found to be susceptible to romantic interests (see section 3) and that these data suggest that judgments of beauty and ugliness are not similarly sensitive, this study provides tentative evidence that the concepts BEAUTY (UGLINESS) and ATTRACTIVENESS (UNATTRACTIVENESS) may be empirically distinct, and that beauty and ugliness may be disinterested in some sense.

---

[8] For the same reason, the ANOVAs reported in studies 1, 2, and 4 also utilised heteroskedasticity-consistent standard errors.

[9] Future studies might wish to confirm that the effects obtained in studies 1–4 are bidirectional—that viciousness makes for ugliness and virtuousness makes for beauty—by including a morally neutral control; and that disagreement at least sometimes indicates judgments of ugliness, rather than merely confident judgments of plainness.
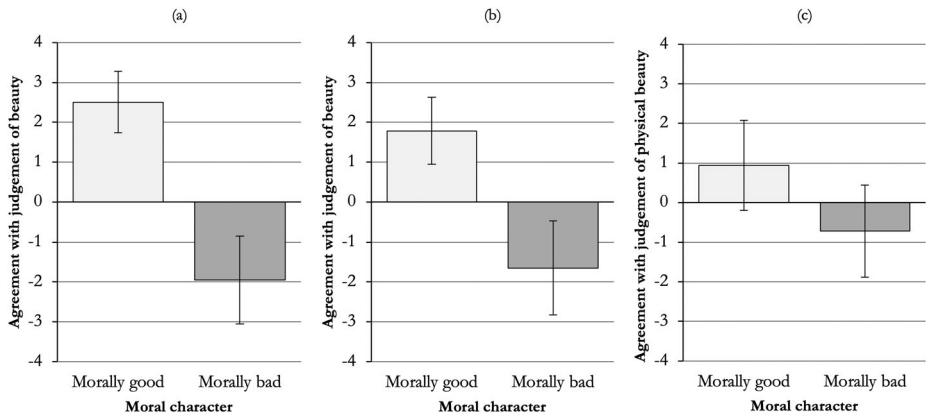
Figure 1. (a) Study 1 (judgment of beauty, no deflationary control); (b) Study 2 (judgment of beauty, deflationary control); (c) Study 3 (judgment of physical beauty, deflationary control). Scales run from -4 (Strongly disagree) to 4 (Strongly agree), and error bars indicate ± 0.5 standard deviations.

4.2 Study 2

As I have argued in section 3, sceptics might worry that these findings do not reflect anything about the folk concepts of beauty (ugliness) or come to bear on the nature of beauty (ugliness). Rather, they may argue that these results can be deflated in two ways. These results may merely indicate the influence of pragmatic factors, such as a participant's desire to express their approval (disapproval) of the model's moral character.[10] These results may also reflect the operation of an error-generating mechanism, such as affective misattribution. In making judgments of beauty (ugliness), participants may reflect on how they feel, and mistake the feelings that are caused by the model's moral character for those that are caused by the formal character of the model's appearance.

Fortunately, these deflationary possibilities can be mitigated. The influence of pragmatic factors such as the desire to approve (disapprove) can be alleviated by giving participants a clear way of expressing the relevant desire. In an analogous context, Knobe [2003] mitigated the effect of participants' desire to praise or blame on judgments of intentional action, by giving participants the opportunity to express this desire before making the relevant judgment. Similarly, and notwithstanding the special aggravating factors noted in section 3, affective misattribution can normally be mitigated by making the true source of participants' affect salient to them. Schwarz and Clore [1983], for example, eliminated the effect of positive affect caused by sunny weather on participants' judgments of life satisfaction by asking them what the weather was like before they judged how satisfied they were with their lives.

Method
A variation of study 1 was conducted. To both give participants the opportunity to express their desire to approve (disapprove), and to make salient to them the true

---

[10] Given that only the manipulation of moral character—and not of formal character—affected judgments of agreement, a sceptic might be especially worried that participants are using their judgment of agreement (disagreement) to express approval (disapproval) rather than to report aesthetic reality.

cause of any affect that they experience, a question asking participants to express their approval (disapproval) was presented before participants were asked to indicate the extent to which they agreed that the model is beautiful.

Results

336 participants took part in the study. Of these, 91 participants withdrew or failed the manipulation and attention checks, and three participants did not provide enough information to determine whether romantic attraction was possible. These participants were excluded from the analysis, leaving 242 participants (70% female, mean age: 38).

A one-way ANOVA revealed that participants in the morally good group approved of the model's actions a great deal more than did participants in the morally bad group ($F(1, 240) = 3693.8$, $p<.001$, Cohen's $d = 7.78$).[11] A three-way ANOVA—moral character x formal character x possibility of romantic attraction—revealed main effects of moral character ($F(1, 234) = 145.75$, $p<.001$, Cohen's $d = 1.67$) and formal character ($F(1, 234) = 13.80$, $p<.001$, Cohen's $d = .31$): participants in the morally good group agreed much more with the judgment of beauty (Mean = 1.78, SD = 1.68) than did participants in the morally bad group (Mean = -1.65, SD = 2.37, see Figure 1(b)); and participants in the formally good group agreed slightly more than did participants in the formally bad group.

Discussion

Since there was a small reduction in the size of the effect of moral character on agreement with the judgment of beauty (compare graphs (a)–(b)), the sceptic's worry that evidence ostensibly supporting General-MBU might in fact reflect pragmatic factors or affective misattribution appears to be partially vindicated. Participants do indeed seem, *in part*, to be either using their judgment of beauty (ugliness) to express their approbation (disapprobation), or misattributing affect arising from the model's character to the formal character of the model's appearance. Crucially though, since the effect of moral character on agreement remains strikingly large, the evidence for General-MBU can only be partially deflated by appeal to such factors.

In addition, as there was again no main effect of the possibility of romantic attraction, or interaction with moral character, it does not seem that romantic attraction is responsible for the effect of moral character.

5. Effect-MBU

As noted in section 2, there are at least three prominent ways in which General-MBU can be true—Effect-MBU, Trait-MBU, or Character-MBU. So, does the beauty (ugliness) of moral goodness (badness) reside on the outside, the inside, or both? Let me begin with Effect-MBU.

---

[11] Analyses of this measure in later studies are not reported. In all cases, the manipulation had a significant and extremely large effect on ratings of approval in the expected direction.

5.1 Study 3

To investigate whether the effects observed in studies 1 and 2 reflect a change in judgments of the models' physical appearance, a variation of study 2 was conducted. This study also sought to eliminate one further deflationary explanation. As noted in section 3, it may be that morally good (bad) people are described as 'beautiful' ('ugly') non-literally not only to express approval (disapproval), but also to express the thought that the individual has a stable positive (negative) disposition.

Method

A variation of study 2 was conducted. Participants were given the opportunity to express the thought that the model has a stable positive (negative) disposition in addition to their approval (disapproval) on a page before viewing the model and indicating the extent to which they agreed that the model is 'physically beautiful'.

Results

279 participants took part in the study. Of these, 59 participants withdrew or failed the manipulation and attention checks, and the possibility of romantic attraction could not be determined for one participant. These participants were excluded from the analysis, leaving 219 participants (76% female, mean age: 38). A one-way ANOVA revealed that participants in the morally good group were much more inclined to think that the model was a good person than did participants in the morally bad group ($F(1, 217) = 2403.19$, $p<.001$, Cohen's $d = 6.62$). A three-way ANOVA—moral condition x formal condition x possibility of romantic attraction—revealed a main effect of moral character ($F(1, 211) = 17.55$, $p<.001$, Cohen's $d = .72$, see Figure 1(c)) and formal character ($F(1, 211) = 10.32$, $p<.01$, Cohen's $d = .55$) qualified by an interaction between the two ($F(1, 211) = 3.91$, $p<.05$). Pairwise comparisons indicated an ordinal interaction, with the effect of moral character on agreement with the judgment of physical beauty being larger for formally ugly models (Mean Immoral = -1.49, SD = 2.09; Mean Moral = .69, SD = 1.88, $p<.001$, Cohen's $d = 1.10$) than for formally beautiful models (Mean Immoral = .30, SD = 2.24, Mean Moral = 1.13, SD = 2.55, $p = .054$, Cohen's $d = .35$).[12]

Discussion

These data suggest that beliefs about a person's moral character affect the perception of the aesthetic quality of that person's physical appearance. As was the case in studies 1 and 2, there was no main effect of the possibility of romantic attraction, or significant interactions involving this factor. As such, this study provides additional support for the idea that participants' judgments cannot be deflated as reflecting a romantic interest.

As the effect is present even when participants have expressed their approval (disapproval) and the idea that the model has a stable good (bad) disposition, and thereby also had made salient the source of any affect that they may have felt, it does not seem likely that the effect can be deflated in terms of these pragmatic factors or misattribution.

---

[12] It could be that participants who received the formally bad models after coming to know that they were morally bad found it easier to bring the information about their character to mind when deciding whether they were physically beautiful. However, since this interaction has a p value of .049 and it was not replicated in any of the other studies, the veracity of this result is doubtful.

Finally, while still being striking, the effect is considerably smaller than when participants are asked to agree with a judgment of beauty (compare graph (c) with (a) and (b)). As such, these data imply that MBU might lie both inside and out: if all of the changes in judgments of beauty and ugliness in study 2 were mediated by changes in perceptions of physical beauty and ugliness, then one would expect the effects observed in studies 2 and 3 to be roughly the same size. This implies that at least one of the dispositional conceptions is also true.[13] I provide further support for this in section 6.

5.2 Study 4

Study 3 shows that moral character affects judgments of physical beauty (ugliness), and provides defeasible evidence that this is not due to affective misattribution. But, as I note in section 3, there are good reasons to think that affective misattribution may be difficult to expunge if it is occurring, and it may reasonably be thought that we cannot be confident of Effect-MBT until we have evidence of a plausible mechanism by which it could be true. Faced with these challenges, it is helpful to focus on the grounds on which judgments of physical human beauty (ugliness) are made.

To focus on beauty, it is plausible that the beauty of female faces tends to lie in the properties of balance, delicateness, and perhaps softness (at least in the West). Symmetry tends to make for beauty in all faces, and 'femininity' tends to make for beauty in female faces at least—where more feminine-looking faces are those that look more neotenous, are smaller, and have less prominent brows and jaws, smaller noses, larger eyes, and fuller lips (as summarised by Perrett [2010]).

How can this help to rebuff the sceptical threat posed by the possibility of affective misattribution? Unlike judgments of relatively thin properties, such as physical attractiveness, merely negative or positive affective responses (and even those that satisfy constraints such as having a disinterested aetiology) are insufficient for true predications of these thick properties. Having a delicate appearance, for example, is a matter of *looking* a particular kind of way, and not a matter of simply tending to make people feel good or bad. As a result, the mere positive or negative responses that the models' characters are likely to give rise to are much less likely to be misattributed and used to make judgments of those aesthetic properties that realise the models' formal beauty. The same is true of judgments of the aesthetic properties that tend to ground ugliness.

The relationship between thick aesthetic properties and physical beauty and ugliness also suggests a mechanism by which Effect-MBU could be true, and so may assuage the sceptical worry that some deflationary explanation of evidence of Effect-MBU is probably warranted. Beliefs about an individual's moral character may be able to amplify, or even themselves cause, the perception of certain properties that will tend to contribute to their formal beauty or ugliness. How is this?

Someone who has the moral vice of being insensitive to the suffering of others may be considered a hard individual. Likewise, someone whose facial features are well-defined and prominent might be considered to look hard. An individual who acts

---

[13] Since two measures were changed, the reduction in effect size may be due *solely* to the satisfaction of a second pragmatic factor—namely, the desire to express the thought that the individual has a stable positive (negative) disposition. Given that the effect size was only slightly smaller than the one obtained in study 4, where this additional factor was not mitigated, this doesn't seem likely.

without favour or prejudice may justly be described as 'fair', 'even-handed', or 'balanced'. Similarly, someone whose facial features relate to one another evenly—in size, shape and arrangement—may be considered to look balanced.

Not only do properties such as hardness and balance admit of moral and aesthetic analogues; these analogues can also form at least part of the realisation base of beauty (ugliness) and moral goodness (badness). For example, as I have just noted, having features that are balanced with respect to one another, in spatial arrangement and size, will tend to make for beauty, and being fair-minded will also tend to make for being morally virtuous.

Given this symmetry, it seems plausible that a change in beliefs about the quality of an individual's character might occasion a change in the perception of beauty or ugliness via a change in the perception of the corresponding aesthetic quality. So, does this seem to be true?

Method

A variation of study 3 was conducted. The individuals described in the counterbalanced vignettes used in studies 1–3 differ in a number of moral characteristics that admit of aesthetic analogues. In the morally bad variant, for example, the individual (i) exploits a vulnerable person, and so will be thought to be callous (where hard is the aesthetic analogue); (ii) deceives, and so will be thought to be oily (where oily is also the aesthetic analogue); (iii) steals, and so will be thought to be unfair-minded (where imbalanced is the aesthetic analogue); and (iv) indulges herself rather than acting for others, and so will be thought to be coarse (where coarse is also the aesthetic analogue). *Mutatis mutandis*, the individual in the counterbalanced vignette will be thought to have the opposite characteristics.[14]

After reading about the model's actions and judging the extent to which they approved (disapproved) of the model on one page, participants were then presented with a photograph of the model and asked to indicate the extent to which they found the model's appearance to be delicate, hard, balanced, and oily, each on a 7-point scale. In a conceptual replication of study 3, participants were then asked to indicate the extent to which they agreed that the model's 'appearance is beautiful'.

Results

364 participants took part in the study. Of these, 104 participants withdrew or failed the attention and manipulation checks, and the possibility of romantic attraction could not be determined for two participants. These participants were excluded from the analysis, leaving 258 participants (67% female, mean age: 36).

A three-way ANOVA—moral character x formal character x possibility of romantic interest—revealed main effects of formal character ($F_{(1, 250)}$=43.06, $p < .001$, Cohen's d=.76) and moral character ($F_{(1, 250)} = 40.57$, $p < .001$, Cohen's d=.81)—with participants in the morally good group agreeing more with the judgment of physical beauty (Mean = 4.71, SD = 1.44) than did participants in the morally bad group (Mean = 3.49, SD = 1.58).

---

[14] The only exception is fairness: the individual in the good variant gives gifts rather than steals, and those who give gifts are not necessarily fairer. However, participants may assume this, given that she is otherwise described as a paragon of virtue.

From a parallel multiple mediation analysis conducted using ordinary least squares path analysis [Hayes 2013], both moral and formal character indirectly influenced agreement with the judgment of physical beauty through their effect on the appearance of balance and delicacy, respectively. As can be seen in Figure 2, participants presented with a model with good moral character tended to find the model's appearance to be more delicate ($a = .76$) and balanced ($a = 1.12$) and less oily ($a = -.87$) and hard ($a = -.57$) than did participants who were presented with a model with bad moral character. Similarly, participants presented with a model with good formal character found the model's appearance to be more delicate ($a = 1.04$) and balanced ($a = .80$) and less hard ($a = -.78$), although not less oily. Each a is an estimate of the difference made to a given aesthetic property by just the difference in moral or formal character. So, for example, the good moral character vignette can be thought to have made the model look .76 points more delicate than did the bad moral character vignette on a 7-point scale. Participants who found the model's appearance to be more delicate and balanced tended to agree more strongly that the model's appearance is beautiful ($b = .34$ and .30 respectively). By contrast, participants who found the model's appearance to be less hard and oily did not tend to agree more strongly that the model's appearance is beautiful. Each b estimates the difference made to participants' agreement with the judgment of physical beauty by a one-point increase in the aesthetic property concerned alone.
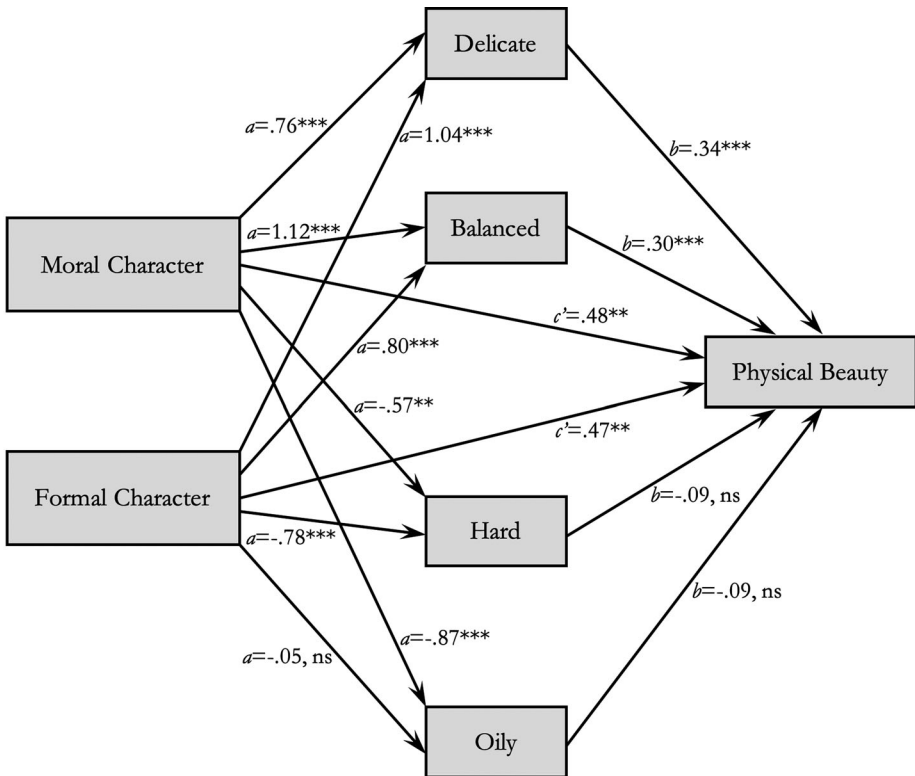

Figure 2. A parallel multiple mediation model showing the effect of formal and moral character on the aesthetic properties of faces. ***= $p<.001$, **= $p<.01$, and ns = not significant.

Bias-corrected bootstrap confidence intervals for the effect of moral and formal character on perceptions of physical beauty indirectly via the appearance of delicacy ($ab_{Moral} = .26$, $ab_{Formal} = .35$) and balance ($ab_{Moral} = .34$, $ab_{Formal} = .24$) indicate that these indirect effects were significant at the .05 level. By contrast, bias-corrected bootstrap confidence intervals for the effect of moral and formal character on perceptions of physical beauty indirectly via the appearance of oiliness ($ab_{Moral} = .08$, $ab_{Formal} = .004$) and hardness ($ab_{Moral} = .05$, $ab_{Formal} = .07$) indicate that there were no indirect effects here. In addition to the indirect effects, there was evidence that both good moral and formal character increased perceptions of physical beauty independently of their effects on the perceived delicacy, balance, hardness, and oiliness of the models' appearances ($c'_{Moral} = .48$, $c'_{Formal} = .47$).

Discussion

The results of this study support a number of claims. First, these data provide a conceptual replication of study 3 (where participants were asked to agree with a judgment of 'physical beauty'), and, as such, provide further support for Effect-MBU. Second, these data suggest that beliefs about moral character generally—and perhaps about specific moral qualities such as being caring (being callous)—affect the perception of aesthetic properties that admit of moral analogues, such as balance, oiliness, delicateness, and hardness. Third, these data suggest that the change in agreement with the judgment of physical beauty is mediated at least in part by changes in the perception of properties that reliably realise beauty in female faces at least—namely, balance and delicateness. As such, these data provide evidence of a plausible mechanism by which Effect-MBU could be true, and suggest that deflationary explanations such as affective misattribution are unlikely to be warranted in a wholesale manner.[15]

Even though changes in moral character affected the perception of the oiliness and hardness of the models' appearances, these changes did not in turn affect the perception of physical beauty. In the case of oiliness, this is probably because the change in the appearance of oiliness was not sufficient to affect the aesthetic character of the models' faces: only a surfeit or complete absence of oiliness tend to contribute to the ugliness of a face, as when a face appears greasy or dry, respectively. In the case of hardness, this is probably because having a morally good or bad character made the models' faces appear soft or hard, respectively, and both appearing hard and soft can make for beauty: female faces can be beautiful by being both soft and dainty (for example, many of Waterhouse's depictions of women), or by being hard and statuesque (for example, many of Rosetti's depictions of women).[16]

---

[15] These data also suggest that the effect cannot be deflated in terms of valence-congruent conceptual priming. The formal character of the models' appearances did not affect ratings of oiliness ($a = -.05$, ns), which is usually a bad quality, as would be predicted if the formal character of the model simply primed the concept OILINESS in a valence-congruent manner. A better explanation for this is that the participants are sensitive to its descriptive content—namely, the *look* of oiliness.

[16] In so far as the kinds of properties that make for beauty (ugliness) in a broad sense can vary, depending on the kind of object of which a given case is an instance, conceptions of Effect-MBU and General-MBU that concern beauty (ugliness) in a broad sense may not always be true. If good moral character makes appearances more delicate, then, in so far as delicateness can be an ugliness-making property of men in a broad sense, moral character could make men uglier in a broad sense. By contrast, in so far as the conceptions are formulated in terms of beauty in a narrow sense, they may be more universally true. Since beauty in a certain narrow sense is typically found in, for example, things that are smooth, small, delicate, mild-coloured, and have gradual variation, such as greyhounds, myrtle trees, rolling hills and meandering rivers (Burke [1757]; see also Kant [1764: 14–18] for similar

Although it is far beyond the scope of these data to say much more about the specific mechanisms by which moral characteristics influence their aesthetic analogues, likely candidates include attentional amplification and cognitive penetration.[17] In the latter case, the mechanism may involve conceptual metaphors and imagery. To focus on just one of the pairs of analogues here, it seems likely that the moral quality of being callous and the aesthetic quality of looking hard are understood in terms of hardness just because physical interactions with people with these properties have felt hard or would feel so: the bodies and faces of callous people do not soften to show care in the face of objects that warrant these changes, and the bodies and faces of hard looking people are likely to be literally hard to the touch. More masculine-looking people, for example, tend to look harder in virtue of having greater bone- and muscle-growth. In a similar manner to how visual representations penetrate auditory representations in the McGurk effect [McGurk and MacDonald 1976], the effect of beliefs about someone's callousness may penetrate visual perceptions of hardness via shared tactile representations of hardness.

Importantly, this study suggests that even those who (howsoever mistakenly) think that beauty and ugliness are perception-dependent should accept the existence of some form of MBU.[18]

6. Character-MBU

Studies 1–4 provide evidence that supports General-MBU and Effect-MBU which cannot be fully deflated in terms of the most likely and threatening pragmatic factors and errors. The large difference in the effect sizes between studies 1–2 and 3–4 suggests that one of the dispositional conceptions might also be true. I here focus on just Character-MBU, as this conception is more amenable to being investigated empirically.

The evidence in favour of Effect-MBU presents special challenges in devising experiments to provide support for Character-MBU. On the one hand, if participants are asked to judge the beauty (ugliness) of a person, then any effect of moral character may be attributable to a change in the person's appearance. On the other hand, if participants are asked to judge whether a person's character is beautiful (ugly) directly, then it is implied that the possibility that entities such as characters can be beautiful (ugly) is assumed by the experimenter, and, as a result, a demand is placed on participants to answer in accordance with Character-MBU.[19] For these reasons, a new design was

---

examples), good moral character may always make things more beautiful in this sense by, for example, making them appear more delicate.

[17] Future studies might test the amplification hypothesis by seeing whether beliefs about moral goodness affect the aesthetic qualities of faces that entirely lack beauty-making features (and the same is true, *mutatis mutandis*, for beliefs about moral badness).

[18] I leave it open whether this is best thought of as showing that appearances can *truly* become more beautiful via moral character or merely *seem* so. This will partly depend on whether contextual properties are relevant in determining the beauty (ugliness) of appearances. But I note that even if it is thought that these data don't support Effect-MBU, on the ground that someone's character is not a part of their appearance, then, interestingly, they may nonetheless support General-MBU as both character and appearance are parts of a person.

[19] Indeed, some of the other evidence that is cited by Paris [2018] is similarly vulnerable to deflation in terms of demand effects. The items used by Diessner et al. [2008] to measure sensitivity to MBU presuppose its existence, and, since presuppositions but not entailments survive negation [Levinson 1983: 191], it is not clear that participants can successfully reject MBU.

employed that sought to target the conceptual coherence of Character-MBU and whether the concepts of beauty and ugliness contain a perceptibility constraint.

6.1 Study 5

Materials and Method

Participants were presented with one of two vignettes. In one, participants were told that two male friends—Chris and Adam—meet a female individual—Caroline—who strikes them as having an ugly appearance, repulses them, and is noticed to have features—such as an unsymmetrical and ill-proportioned face—that ground their judgments of her ugliness. Chris is reported to subsequently discover that Caroline is a good person. He relays this to Adam, concluding in direct speech that, while Caroline looks ugly, she is a beautiful person. In response, Adam expresses some puzzlement, and suggests that Chris does not really mean that she is a *beautiful* person, but merely that she is a very *good* person.

Participants were asked to indicate the extent to which they approved of Caroline's actions, agreed that Caroline is a very good person, and whether it was possible that Chris *really* found Caroline to be a *beautiful* person as a result of his discovery. The second variation was the same, only with the moral and aesthetic valences switched (that is, beautiful appearance, bad character).

There are several features to emphasise about this design. First, Chris is confirmed to do a number of things—to judge Caroline's appearance to have a determinate aesthetic quality (for example ugliness), to undergo one of the responses that is arguably reliably caused and warranted by those qualities (for example, repulsion), and to notice the presence of lower-order perceptual properties that ground these judgments and responses (for example, asymmetry). Moreover, Chris is confirmed to continue to find Caroline's appearance to have the same aesthetic quality after discovering what her moral character is like. As a result, no matter what participants may think is involved in truly finding someone's appearance to be beautiful or ugly, and even if participants believe that Chris's discovery of Caroline's moral character may *pro tanto* affect his perception of the beauty or ugliness of her appearance, we can be confident that participants know that he finds her appearance to be incongruent with her moral character. As a result, should they report that it is possible that Chris could find her to be a beautiful or ugly person, they can justly be taken to be referring to how he could find her in virtue of her *moral character*.

Second, a number of aspects of the vignettes attempt to ascertain a participant's concept of beauty or ugliness without begging the question in favour of, or against, Character-MBU. It may plausibly be thought that the truth-makers in this matter are specific facts about Chris's psychological state—for example, the presence of a certain kind of perceptual or affective response. For this reason, participants were not told anything about the psychological state that his utterance expresses. Moreover, it may naturally be assumed that Chris's expression should be taken at face value in the absence of reasons to be sceptical, not least because it may be thought that he has privileged access to whether he finds something beautiful or ugly. For this reason, a strong dissenting voice was included to prevent such an assumption being made lightly. As a result, the design employed constitutes a fair test of the claim that the folk will rationally endorse the conceptual coherence of Character-MBU.
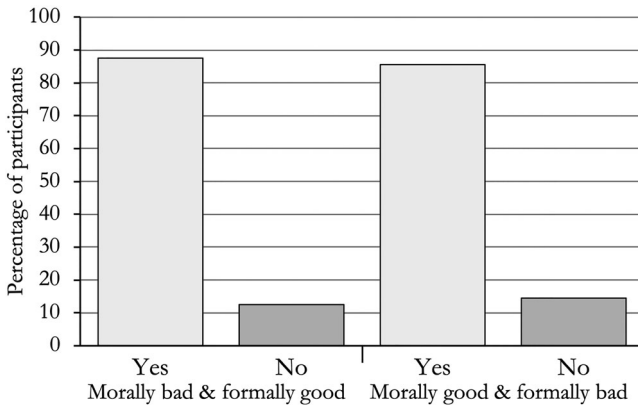
Figure 3. Percentage of participants that thought that it was possible that Chris really found Caroline to be a beautiful or ugly person as a result of his discovery.

Results

242 participants took part in the study. Of these, 78 participants withdrew or failed the attention and manipulation checks. These participants were excluded from the analysis, leaving 164 participants (70% female, mean age: 38).

Goodness-of-fit chi-square tests indicated that participants were much more likely to judge that it is possible that Chris really finds Caroline to be a beautiful person ($\chi^2(1) = 38.37$, $p<.001$) or ugly person ($\chi^2(1) = 49.50$, $p<.001$) as a result of his discovery—with 85.5% of participants agreeing in the moral beauty condition, and 87.5% agreeing in the moral ugliness condition (see Figure 3). A test-of-independence chi-square test indicated that participants in the morally good condition were as likely to judge that it is possible that Chris really finds Caroline to be a beautiful person as participants in the morally bad condition were to judge that it is possible that Chris really finds Caroline to be an ugly person ($\chi^2(1) = .14$, $p = .71$, ns).

Discussion

These data suggest that the majority of the folk endorse the conceptual coherence of Character-MBU, and find the idea of moral ugliness and moral beauty equally coherent, even when they are given the opportunity to express their approval (disapproval), and their belief that moral goodness (badness) is a stable part of the target individual.

Since the participants in this study seem to be competent users of 'beautiful' and 'ugly', and the vast majority claim that Character-MBU is conceptually coherent in a literal manner, the data from this study suggest that beauty and ugliness are not perception-dependent.

7. Conclusion

I have distinguished between three location-specific conceptions of MBU, and have argued that the existing indirect arguments marshalled in favour of Trait-MBU are vulnerable to certain deflationary objections: they may variously be explained away in

terms of pragmatic factors such as non-literal intent and sources of error such as affective misattribution or the influence of an interest that cannot be involved in beauty. I have provided evidence that supports three forms—General-MBU, Effect-MBU, and Character-MBU—which is robust against these worries, while at the same time *partially* vindicating them. As such, these studies allow us to be more confident that beauty and ugliness are not perception-dependent, and, moreover, that the perception of aesthetic qualities can be affected by beliefs about moral character. Moral beauty and ugliness, then, seem to reside on the inside *and* the outside.[20]

Disclosure Statement

Funding

ORCID

Ryan P. Doran 🆔 http://orcid.org/0000-0003-3986-4693

References

Beardsley, M. 1981. *Aesthetics: Problems in the Philosophy of Criticism*, Indianapolis: Hackett.

Burke, E. 1757 (1788). *A Philosophical Enquiry into the Origin of Our Ideas of the Sublime and the Beautiful*, London: J. Dodsley.

Carter, G.L., A.C. Campbell, and S. Muncer 2013. The Dark Triad Personality: Attractiveness to women, *Personality and Individual Differences* 56: 57–61.

Diessner, R., R.C. Solom, N.K. Frost, L. Parsons, and J. Davidson 2008. Engagement with Beauty: Appreciating Natural, Artistic, and Moral Beauty, *The Journal of Psychology: Interdisciplinary and Applied* 142/3: 303–32.

Dutton, D.G. and A.P. Aron 1974. Some Evidence for Heightened Sexual Attraction under Conditions of Anxiety, *Journal of Personality and Social Psychology* 30/4: 510–17.

Field, A.P. and R.R. Wilcox 2017. Robust Statistical Methods: A Primer for Clinical Psychology and Experimental Psychopathology Researchers. *Behaviour Research and Therapy* 98: 19–38.

Gaut, B. 2007. *Art, Emotion and Ethics*, Oxford: Clarendon Press.

Gladue, B.A. and J.J. Delaney 1990. Gender Differences in Perception of Attractiveness of Men and Women in Bars, *Personality and Social Psychology Bulletin* 16/2: 378–91.

Hayes, A.F. 2013. *Introduction to Mediation, Moderation and Conditional Process Analysis: A Regression-Based Approach*, New York: The Guildford Press.

Kant, I. 1764 (2011). *Observations on the Feeling of the Beautiful and Sublime*, ed. P. Frierson and P. Guyer, Cambridge: Cambridge University Press.

Kant, I. 1790 (2000). *Critique of the Power of Judgement*, trans. P. Guyer and E. Matthews, Cambridge: Cambridge University Press.

Knobe, J. 2003. Intentional Action in Folk Psychology: An Experimental Investigation, *Philosophical Psychology* 16/2: 309–24.

Langner, O., R. Dotsch, G. Bijlstra, D.H.J. Wigboldus, S.T. Hawk, and A.D. van Knippenberg 2010. Presentation and Validation of the Radboud Faces Database, *Cognition and Emotion* 24/8: 1377–88.

Levinson, S.C. 1983. *Pragmatics*, Cambridge: Cambridge University Press.

Nehamas, A. 2007. *Only a Promise of Happiness: The Place of Beauty in a World of Art*, Princeton: Princeton University Press.

Nisbett, R.E. and T.D. Wilson 1977. The Halo Effect: Evidence for Unconscious Alteration of Judgments, *Journal of Personality and Social Psychology* 35/4: 250–6.

McGurk, H. and J. MacDonald 1976. Hearing Lips and Seeing Voices, *Nature* 264/5588: 746–8.

Plato. c.370 BC (2010). Symposium, in *Greek and Roman Aesthetics*, ed. and trans. O. Bychkov and A. Sheppard, Cambridge: Cambridge University Press: 22–5.

Paris, P. 2018. The Empirical Case for Moral Beauty, *Australasian Journal of Philosophy* 96/4: 642–58.

Perrett, D. 2010. *In Your Face: The New Science of Human Attraction*, Basingstoke, UK: Palgrave Macmillan.

Pham, M.T. 1998. Representativeness, Relevance, and the Use of Feelings in Decision Making, *Journal of Consumer Research* 25/2: 144–59.

Schwarz, N. and G.L. Clore 1983. Mood, Misattribution, and Judgments of Well-Being: Informative and Directive Functions of Affective States, *Journal of Personality and Social Psychology* 45/3: 513–23.

Shaftesbury [Cooper, Anthony Ashley, Third Earl of Shaftesbury] 1711 (1999). *Characteristics of Men, Manners, Opinions, Times*, ed. L.E. Klein, Cambridge: Cambridge University Press.