

Gunning for affective realism: Emotion, perception and police shooting errors

Raamy Majeed

To cite this article: Raamy Majeed (10 Mar 2023): Gunning for affective realism: Emotion, perception and police shooting errors, *Philosophical Psychology*, DOI: [10.1080/09515089.2023.2188890](https://doi.org/10.1080/09515089.2023.2188890)

To link to this article: <https://doi.org/10.1080/09515089.2023.2188890>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 10 Mar 2023.



Submit your article to this journal [↗](#)



Article views: 358



View related articles [↗](#)



View Crossmark data [↗](#)

Gunning for affective realism: Emotion, perception and police shooting errors

Raamy Majeed 

Lecturer in Philosophy, Department of Philosophy, University of Manchester, Manchester, UK

ABSTRACT

Affective realism, roughly the hypothesis that you “perceive what you feel”, has recently been put forward as a novel, empirically-backed explanation of police shooting errors. The affective states involved in policing in high-pressure situations result in police officers literally seeing guns even when none are present. The aim of this paper is to (i) unpack the implications of this explanation for assessing police culpability and (ii) determine whether we should take these implications at face value. I argue that while affective realism stands to diminish, if not eliminate, the moral and legal responsibilities of officers who have made shooting errors, the empirical data itself does not directly support such a radical rethink of police culpability.

ARTICLE HISTORY

Received 26 September 2022
Accepted 3 March 2023

KEYWORDS

Affective realism; emotion;
racism; shooter bias;
predictive processing

1. Introduction

Police officers in the U.S shoot and kill roughly a thousand people each year, with a large number of these cases being attributed to shooting errors.¹ This has come under the spotlight in recent years as part of a wider conversation about the police use of force against African Americans. With regards to the shootings themselves, no one particular race or ethnicity is immune, but they disproportionately affect people racialized as black.² The standard explanation for these tragic statistics is that police officers are prone to making errors in judgment. Having to make rapid decisions in potentially life-threatening circumstances, officers often mistake a neutral object (e.g., a wallet, a cellphone) or action (e.g., someone reaching into their pockets) as a threatening one, and act on the basis of such errors. Various factors might contribute to this, including officer fatigue, situational features, e.g., the neighborhood in question, unconscious bias and of course explicit racism.³

CONTACT Raamy Majeed  raamy.majeed@manchester.ac.uk  Lecturer in Philosophy, Department of Philosophy, University of Manchester, Manchester, UK

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

This way of explaining police shooting errors has been turned on its head by a more radical explanation, one that is said to be supported by the latest work done at the intersection of cognitive psychology and neuroscience. At the very heart of this is the hypothesis of affective realism: roughly the view that your feelings can influence what you see; not what you think you see but what you literally see.⁴ Affective realism gains its practical stripes on account of the novel way it explains police shooting errors: “In a dangerous, high-pressure situation such as a possible crime scene, it’s conceivable that some police shooters actually see a weapon when none is present” (Barrett & Wormwood, 2015, p. 3). They might mistakenly see, say, a wallet as a gun.⁵

To be clear, proponents of affective realism are not saying that this is the primary cause of police shooting errors or that race isn’t a factor. With regards to the latter, they are careful to point out that affective realism might actually help explain how race features in such errors. If the suspect is a black male, negative racial stereotypes, e.g., about black men being dangerous, might make it more likely that an officer will see the object in the suspect’s hand as a gun.

The way affective realism is viewed within the context of police shooting errors, then, is cautious, but it still has the potential to prove morally and legally explosive. For example, according to standard explanations of police shooting errors, there is a case to be made that police officers are, to a significant degree, culpable for their actions, as they are acting on the basis of errors in judgment that are avoidable. By contrast, the affective realist explanation undercuts this argument, as officers, at least in some cases, can be seen to be acting on the basis of perceptual errors which by their very nature seem unavoidable.⁶ In this vein, affective realism, where relevant, stands to drastically diminish police culpability, if not do away with it completely.

Before we get ahead of ourselves, however, I think the question of whether affective realism is actually a plausible, empirically-backed explanation of police shooting errors is in need of closer scrutiny. In this paper, I argue that there are, as a matter of fact, two distinct views assumed under the broad banner of “affective realism”, which lends itself to two different ways it can be employed to explain police shooting errors. Moreover, I argue that the empirical data doesn’t directly support either explanation. In fact, I show that the case is weakest for the explanation actually on offer by proponents of affective realism; the one that involves a radical rethink of police culpability. If I am right, unchecked endorsements of affective realism as an explanation of police shooting errors are both misleading and potentially harmful.

The paper is structured as follows. First, I provide a brief overview of affective realism and explore more closely the moral and legal ramifications that ensue from the affective realist explanation of police shooting errors

(Sect. 2). I then provide an exposition of the empirical data employed to justify this explanation. I argue that the experimental findings themselves don't actually support the explanation. Rather, the explanation gains its plausibility on account of the predictive processing framework used to interpret these findings; a framework which assumes that what we see is a construction of the brain's best guess at incoming sensory information (Sect. 3). Finally, I reexamine the application of affective realism to police shooting errors in light of this discussion (Sect. 4).

2. Affective realism and police shooting errors

Affect, in a general sense, is the feeling you experience at any given moment. More precisely, affect is a phenomenon characterized along the dimensions of valence and arousal. A state has a valence if it can be evaluated as being positive or negative, e.g., whether it feels pleasant or unpleasant. Moreover, a state can exhibit arousal insofar as it comes with a degree of intensity, e.g., whether it makes you feel calm or agitated. Emotions and moods are affective states, but affect is supposed to be a feature of all, if not most, of our conscious experiences, and is typically thought to stem from interoception: the way your brain represents sensations from your own body.⁷

While affect involves internal sensations, perception is typically characterized as our sensory experience of the external world. Seeing, hearing, tasting, smelling and touching are all species of perception. However, unlike "affect" the term "perception" is often used in an ambiguous sense, and can mean the raw sensations we have of the outside world or the ways we interpret these sensations. Affective realism, roughly the view that "you perceive what you feel", is supposed to concern the former. For example, Barrett and Wormwood define affective realism as "the tendency of your feelings to influence what you see – not what you think you see, but the actual content of your perceptual experience" (Barrett & Wormwood, 2015, p. 1).

The case for this will be explored in [Section 3](#), but for now note that, if true, it is significant for a number of different reasons. First and foremost, it heralds a new way of understanding perception. Cross-modal perceptual effects are pervasive and well-understood. We know that smell affects taste, that vision affects sound, and so on.⁸ However, whether affective states, e.g., feelings, emotions and moods, can influence perceptual experience in a similar fashion is less clear. The proponents of affective realism are keen to stress that affective states can change the content of your perceptual experiences in a similar fashion. For example, it was found that participants see neutral faces as either smiling or scowling when paired with (unconscious) affective information (E. H. Siegel et al., 2018).

Affective realism, then, should be of interest to both the philosophy and psychology of perception. However, the significance of affective realism isn't just academic. As Siegel *et al.* observe, "affective realism stands to fundamentally alter our understanding of how perception influences decision making in real-world scenarios where errors can have costly, potentially deadly, consequences" (E. H. Siegel *et al.*, 2018, p. 502). This is clearest in the way affective realism is applied within the context of police shooting errors.

Police officers are often operating in high-pressure, potentially life-threatening, situations where their feelings are bound to run high. When faced with making split-second decisions in such affect-laden circumstances, police officers might see a neutral object (e.g., a wallet, a cellphone) as a threatening one (e.g., a gun), and make a decision to shoot on this basis. The important point is that such a decision is made not just based on a false belief or an error judgment, but on the basis of a visual experience which is itself in error. This gives us a completely new way of understanding what might be going wrong when police make real-life shooting errors. Moreover, such an understanding calls for nothing short of a radical reassessment of how we think about police culpability.

There are two dimensions to this, one moral and another legal. Consider the legal dimension. According to police training guidelines in most U.S. states, the use of deadly force is only justified when it is "reasonable", i.e., when an officer reasonably believes that they or someone else is under imminent threat or serious bodily injury. Within this setting, one question that proves important is, Will the way jurors apply this standard vary depending whether they think the officer thought they saw a gun or whether they actually saw one? I venture that jurors will be more likely to deliver a favorable verdict to the defendant if they think affective realism is at play; if they believe the officer literally saw a gun.⁹

What's more, this I gather mirrors how we think we ought to regard culpability in such cases. This is the moral dimension. According to standard explanations of police shooting errors, there is a case to be made that police officers are, to a significant degree, morally responsible for their actions, as they are acting on the basis of errors in judgment that are avoidable. In general, moral considerations follow an "ought-implies-can" principle, where an agent has a moral obligation to perform some action only if they can perform it. Misjudgments, and by extension the shooting errors to which they give rise, are avoidable, so there seems to be a moral obligation for not committing shooting errors. The affective realist explanation undercuts this argument, as officers, at least in some cases, can be seen to be acting on the basis of perceptual errors which by their very nature seem unavoidable. If these perceptual errors are unavoidable, it looks like the shooting errors caused by them aren't avoidable either; or at the very

least, they are significantly less avoidable. Affective realist explanations, subsequently, have the potential to eliminate, if not significantly minimize, the moral responsibility of police officers, even in cases where their errors prove fatal.¹⁰

Proponents of affective realism aren't blind to these concerns. Barrett (2017), for instance, is careful to emphasize that she is not denying that police officers are to blame for their errors. She justifies this by appealing to the legal notion of foreseeability: the extent to which you can foresee a probable harm if you act in a negligent way. Within law, foreseeability, if applicable, dictates that you are culpable for a harmful action – whether you intend harm or not. According to Barrett, foreseeability is at play when it comes to police shooting errors, including when racial bias is a factor:

If your brain predicts that an African American youth in front of you is holding a weapon, and you perceive a gun where none is present, you have some degree of culpability even in the face of affective realism, because it is your responsibility to change your concepts. If you educate yourself and inoculate yourself against such stereotypes, expanding your conceptual system with the goal to change your predictions, you still might mistakenly see a gun where none is present, and a tragedy still might occur. But your culpability is diminished somewhat, because you've acted responsibly to change what you can (Barrett 2017: 250).

On this reasoning, if you are aware of data which suggests that negative racial stereotypes can impact how you feel and act, then you have a duty to try to rid yourself of such stereotypes. If you happen to act on them, then you might still be culpable, but this is a lot less worse than had you not tried to shield yourself from such stereotypes to begin with.

I think Barrett is right to assume that culpability comes in degrees. Moreover, I think she is also right that the degree to which you are culpable for a shooting error can vary depending on whether you are aware of affective realism and its potential harm, and whether you have taken any preventive measures to mitigate them. Still, I think the foreseeability argument, and others like it, don't waylay all of the concerns that stem from the affective realist explanation of police shooting errors. If police officers are making mistakes because they literally see a gun, regardless of foreseeability concerns, they are still less culpable for their error than had they not literally seen a gun. This still amounts to a radical revision of how we ought to think about police culpability.

So how should we proceed? Should we take these moral and legal implications at face value? I think answers to such questions depend on more fundamental questions that haven't been fully explored. In particular, How plausible, really, is the affective realist explanation of police shooting

errors? Do empirical data, as a matter of course, back the claim that police officers literally see guns even when none are present?

3. The case for affective realism in police shooting errors

Proponents of affective realism defend their explanation of police shooting errors by appealing to a wide range of experimental findings. In what follows, I attempt to impose some structure on these findings by classifying the relevant experiments into two categories: valence-detection tasks in affect-laden faces and object-recognition tasks concerning affect-laden objects. I argue that these categories foreshadow two distinct affective realist hypotheses, a misattribution hypothesis and a misidentification hypothesis, which is often obscured in theoretical discussions of the phenomenon. Furthermore, I argue that once we make this distinction, we see that the empirical findings themselves simply do not support either hypothesis. The support for them, rather, stems from the predictive processing framework used to interpret these findings.

3.1 *Valence-detection tasks and the misattribution hypothesis*

The first series of experiments that lend credibility to affective realism concern various types of valence-detection tasks in affect-laden faces. In these experiments, neutral faces are typically paired with affective stimuli, and participants are then asked to identify the valence of the faces themselves. For example, in a series of studies, Bliss-Moreau et al. (2008) paired images of neutral faces with sentences describing positive, negative or neutral behaviors, and participants were then asked to judge the valence of the faces in question. In particular, they were asked to categorize whether the faces were positive, negative or neutral. They found that participants were more likely to categorize neutral faces as positive when paired with positive sentences and as negative when paired with negative ones.

A similar set of studies were carried out by Anderson et al. (2012) but with two main differences. First, they paired neutral faces with unconscious affective stimuli, in this case other faces that were either smiling, scowling or neutral. Second, participants were asked to make not only affective judgments but personality ones as well. More specifically, they were asked to judge whether the faces were more pleasant or unpleasant, but also whether they were more or less truth-worthy, attractive, likeable, competent and so on. They found that when primed with the unseen smiling faces, participants were not only more likely to judge the seen faces as pleasant but also more trustworthy, likeable, competent etc. Likewise, when primed with unseen scowling faces, participants were not only more likely to judge the seen faces as unpleasant but also as less trustworthy, likeable, component

and so on. On the whole, these studies are said to show that “affective consequences of unseen faces are misattributed to consciously seen structurally neutral faces” (pg. 2).

So far as I can tell, the first instance of the term “affective realism” occurs in the paper by Anderson et al. (2012). Furthermore, they also offer a possible explanation for this phenomenon:

We can speculate that affective misattribution causes a kind of “affective realism” that is a normal consequence of how the brain processes visual sensations in the context of sensations from the body. At times, somatic and visceral sensations from the body are in the foreground of consciousness so that they are experienced as our own internal reactions to the world: we like or dislike a food, a painting, or a person. When such sensations are in the background because they are not in the focus of attention, they are experienced as a property of the object itself, leading perceivers to experience those objects as affectively significant. (Anderson et al., 2012, p. 13)

The important point for us is not whether this explanation is correct but what it tells us about the phenomenon of affective realism itself. As evident, Anderson *et al.* take affective realism to be a form misattribution, one where the affective features of our interoceptive experiences are misattributed to our visual experiences. Call this *the Misattribution Hypothesis*.

It is worth pointing out that the empirical findings actually lend themselves two kinds of misattribution claims: misattribution influences our perceptual judgments, and misattribution influences the content of our perceptual experiences. This ambiguity is nicely brought out by Bliss-Moreau *et al.* when they point out that “It is not clear, however, if minimal affective learning serves to bias perception per se (i.e., how the perceiver literally sees a target face) or higher order person perception (i.e., how the perceiver categorizes a target face), or both” (Bliss-Moreau et al., 2008, p. 12). So which misattribution claim is relevant to affective realism? As I read them, proponents of affective realism are concerned with misattributions that affect our perceptual experiences. This is what makes affective realism novel.

Within the scope of the experiments that concern affect-laden faces, we find the clearest defense of this idea by E. H. Siegel et al. (2018). As in the series of studies conducted by Anderson *et al.*, participants in these trials were primed with unseen faces that were either smiling, scowling or neutral. However, unlike in the previous studies, participants were then asked to perform a face matching task. Here there were first shown a target neutral face. They were then shown five faces (ranging from slightly scowling to slightly smiling) and were told to identify the image that best matched the target face. It was found that when primed with unseen smiling faces, participants were more likely to identify a smiling face. *Mutatis mutandis* for the scowling faces. What is telling is the way Siegel *et al.* interpret their findings. “We demonstrated that affective realism extends beyond broad

social judgments to the visual perception of neutral faces: Individuals perceive structurally neutral faces as more smiling or scowling when paired with unconscious, affective information” (pg. 497).

Our primary concern is to determine whether affective realism is an empirically viable explanation of police shooting errors. Proponents of the misattribution hypothesis recognize this potential. As Siegel *et al.* observe, “the affective-realism hypothesis may help to explain why police officers perceive targets as more or less threatening depending on the interoceptive information they receive”. (pg. 503). They don’t elaborate on the matter but an explanation could go something like this. When police officers are feeling stressed, fearful or anxious, — all affective states with a negative valence — they mistakenly attribute these negative feelings to the objects they see around them, in which case they might literally experience these objects themselves as having a negative valence. If something like this is right, it is possible that police officers might, *ipso facto*, also experience these objects as threatening, and decide to shoot on the basis of such experiences.

This is fairly speculative, but there are two things worth bearing in mind. First, as we have seen, the empirical findings themselves seem to underdetermine whether the observed affect-based effects concern perceptual judgments or perceptual experiences themselves. Whether any given set of empirical data can ever demonstrate changes to the content of perceptual experiences is bound to be contentious.¹¹ This is a familiar point and I shall not belabor it here except note that we see this play out within the relevant studies as well. The data captures behavioral phenomena, e.g., judgments, categorizations and the like, and it remains contentious whether any such behaviors can track changes in the contents of perceptual experiences themselves.

Proponents of affective realism are sensitive to some of these concerns. Wormwood *et al.* (2019), in particular, aim to provide a “direct test” for the misperception interpretation:

In the present studies, we sought to empirically demonstrate that affective realism is a special case of affective misattribution by manipulating the timing offset between the presentation of an affective stimulus and a to-be-judged (target) stimulus, allowing us to carefully tease apart whether incidental affect can influence the experience of the target stimulus itself and not just the perceivers’ post hoc judgment of the target (Wormwood *et al.* 2019: 789).

As in the previous studies, participants were primed with unconscious affective faces and were then asked to make personality judgments of the target neutral faces that were presented consciously. The main difference is that the present studies also manipulated when the target faces were presented; each target face was presented either before, concurrent with or after a suppressed affective face. It was found that while seen neutral faces were

evaluated more positively when paired with suppressed positive faces and more negatively when paired with suppressed negative faces, these effects were only significant when the seen faces were presented concurrent with the suppressed ones. Wormwood *et al.* take this as conforming a change in perceptual experience itself instead of just a change in post-perceptual judgment on the assumption that simple cases of misattribution *qua* misjudgment would not predict a timing effect. The assumption is not explicitly justified but it is not without merit. If information from the suppressed faces were merely used to make affective judgments, there would be no reason why the evaluations of the seen faces should be stronger when concurrent with the suppressed ones. By contrast, if information from suppressed faces were also integrated into the contents of perceptual experience, we would expect the evaluations of the seen faces to be stronger when presented at the same times as the suppressed ones.

I think this is a step in the right direction. Here we find an experimental manipulation that aims to strengthen a particular phenomenological interpretation of the experimental findings. Unfortunately, it won't, by itself, address all of the concerns we have about whether we can read-off conclusions about perceptual experience from behavioral data. For instance, the findings by Wormwood *et al.* (2019) are still compatible with a number of rival hypotheses. One such hypothesis is that the findings better track an effect of memory. When the suppressed faces were not presented concurrently with the seen ones, participants had to rely on their memory of the former to evaluate the later. By contrast, when both sets of faces were presented concurrently with each other, participants were able to rely not just on their memory but information presently available to them in their environment as well. This might explain why the evaluations were more significant in the concurrent cases than the non-concurrent ones. My point is not that such rival hypotheses can provide better explanations of the experimental data but that such hypotheses aren't ruled out by the addition of a simple time manipulation task. Subsequently, for now we can note that although there are empirical grounds for supposing the misattribution hypothesis, these remain far from conclusive.

Putting aside the issue of what can actually be gleaned from the empirical data, the second issue concerns how we are really supposed to apply these (alleged) findings to cases where police officers have mistakenly shot unarmed civilians. The clearest application of affective realism to such cases is outlined by Barrett and Wormwood (2015), where they argue that police officers might mistakenly see harmless objects, like wallets and cellphones, as threatening objects, e.g., as guns. However, this interpretation of police shooting errors isn't supported by the misattribution hypothesis. All that it indicates, if anything, is that officers might experience a neutral object as threatening. This is different from officers perceiving a neutral

object as a threatening object. The difference is subtle but significant. As we saw earlier, a radical revision to how we think about police culpability stems from supposing that police officers literally see a gun when none is present. Even on a charitable reading of the affect-laden face studies, this isn't something supported by the empirical data.

To elaborate, the misattribution hypothesis supposes that we misattribute features of our interoceptive experiences to features of our perceptual experiences. This puts constraints on the sorts of things we can attribute to the contents of perceptual experiences. Objects such as guns aren't features of our interoceptive experiences. Subsequently, they can't be misattributed to our perceptual experiences. Feelings can be, but things like guns simply can't. Does this mean that the misattribution hypothesis is irrelevant to shooting errors? Not at all. It may very well be that officers make shooting errors because they see objects as affectively significant as a consequence of misattribution. The point, however, is that if affective realism occurs because of misattribution, its proponents provide the wrong sort of explanation of police shooting errors; one that if not false is incredibly misleading.

3.2 Object-recognition tasks and the misidentification hypothesis

The next set of experiments fare a lot better in terms of supporting the explanation of police shooting errors actually on offer by proponents of affective realism. These experiments involve object-recognition tasks concerning affect-laden objects. For example, Baumann and DeSteno (2010) ran a series of studies which aimed to investigate whether the way emotions affect our judgments of threat extends to object-recognition. In the first part of the study, participants were emotionally aroused by asking them to recall events that made them happy or angry under the pretext of completing a memory task. The effectiveness of this procedure was assessed by asking them to complete a questionnaire that included various feeling descriptors. In the second part, participants were shown images of individuals holding guns or neutral objects. Here each target image was shown for 750 ms, and participants were told to press a button to indicate whether the individual in question was holding a gun or a neutral object. They had to do this within 750 ms as well. Baumann and DeSteno found that "anger increases the probability that neutral objects will be misidentified as ones related to violence" (pg. 1).

It is important to give this study some context. Baumann and DeSteno are working alongside the shooter bias paradigm. Shooter bias concerns the tendency to shoot on the basis of racial cues. Shooter bias trials take various forms, but they all concern computer-simulations where subjects are asked to "shoot" armed suspects and to avoid shooting unarmed ones.¹² In the

original shooter trials, Correll et al. (2014) ran a series of studies which demonstrated various racially biased decision-making patterns. In general, they took these studies to show that “if a target was African American, participants generally required less certainty that he was, in fact, holding a gun before they decided to shoot him” (pg. 1325). Baumann and DeStano aimed to investigate the influence of emotion on decisions to shoot by modifying the shooter trials in two respects. First, the possible effects of race were controlled for by making all the target individuals white. Second, instead of asking participants to “shoot” armed suspects, they were asked to identify whether the target individuals were holding a gun or a neutral object. Their results show that situational influences on threat-detection needn’t just be racial. They can be emotional as well. In particular, they found that anger led to identification errors irrespective of the target’s race.

What concerns us here is how these studies enrich our understanding of affective realist explanations of real-life police shooting errors. To that end, the first thing to note is that they give us a different way of understanding affective realism. As we saw earlier, within the context of affect-laden face studies, affective realism is discussed as a misattribution hypothesis: affective features of our introspective experiences are misattributed to our visual experiences. Baumann and DeStano don’t discuss affective realism, but their work is typically taken to indicate an altogether different interpretation of affective realism, one where affective influence on perceptual experience can be seen to be a form of misidentification. Call this *the Misidentification Hypothesis*.

I think there is a strong case to be made that it is the misidentification hypothesis that proponents of affective realism have in mind when they employ their view to explain police shooting errors. For example, Barrett and Wormwood, the first to make explicit the possible links between affective realism and police shooting errors, describe their explanation as an “explanation for police shootings that involve the misidentification of weapons” (pg. 3). However, as before, whether the empirical data actually supports the form of misidentification assumed by proponents of affective realism can be called into question. The crux is, affective realism is supposed to concern changes to the contents of perceptual experiences themselves, which means the relevant form of misidentification must be a form of misperception. This, once again, doesn’t seem to be supported by the empirical data itself.

Baumann and DeStano, for instance, demonstrate that anger increases the probability that subjects will “misidentify” neutral objects as guns. Crucially, they see this as a form of misjudgment not misperception. This is evident in the way they explain their findings: “angry participants set a much lower threshold for saying that a target is holding a gun; they require much less information before they are willing to claim a target individual is

threatening” (pg. 5). Nowhere do they claim that angry participants actually see neutral objects as guns. Rather, what they claim is that angry participants have a lower threshold for when they are willing to say that they see a neutral object as a gun. What is even more telling is that Baumann and DeStano themselves view the misperception hypothesis as an avenue for future research. In particular, they claim, “in addition to anticipating encountering more guns, they might also perceive neutral objects as actually looking more like guns. An exploration of when and if perceptual distortion might occur is an important next step to understanding emotion’s impact on threat detection” (pg. 15).

As evident, Baumann and DeStano take their findings to demonstrate a form of misjudgment, not misperception, but are they right? There are aspects to their results that give more weight to the misjudgment hypothesis than the misperception one. For instance, they found that participants made fewer misidentification errors when they were given a bit more time before they had to decide whether the objects were guns or neutral objects. If anger affects how we literally see objects around us, one would expect this effect to persist outside the 750 ms mark. By contrast, if what is at play is a form of misjudgment – perhaps stemming from an ambiguous percept or an inability to consciously register an unambiguous percept as neutral – we would expect subjects to correct for this when given more time. This won’t rule out the misperception hypothesis completely, however, I think it is a point in favor of the misjudgment hypothesis.

Even if I am wrong about this, the fact remains that the available data doesn’t entail the misperception hypothesis. Here it is instructive to look at S. Siegel’s (2020) work in the philosophy of perception. Like the proponents of affective realism, Siegel is interested in changes to the contents of perceptual experience; in her case, she is interested in changes brought on by cognition as opposed to affect. However, unlike proponents of affective realism, she is sensitive to the difficulty in empirically establishing such changes. For example, she offers a diverse range of rival hypotheses which might explain why subjects misclassify neutral objects as weapons in shooter trials.¹³ In the interests of brevity, I’ll focus on just four of these hypotheses.

The first one involves disbelief: the neutral objects look to the participants just as they are, but they disbelieve their perceptual experience and misclassify them as guns. The second involves attention: the neutral objects look to the participants “somewhat like a gun” because their affective states draw their attention to features of the objects that are “congruent with being a gun” (pg. 103).¹⁴ The third concerns an introspective error: the neutral objects look to the participants exactly like they are, but they make an introspective error in which they take themselves to be experiencing guns. Finally, the fourth hypothesis involves making a hasty judgment: the neutral objects look “somewhat” like what they actually are and “somewhat” like

a gun. However, having to make quick decisions before perceiving all the relevant detail, participants misclassify the objects as guns because of their affective states. The point for us is that all four hypotheses involve misjudgment, not misperception, and are compatible with Baumann and DeStano's findings.

I appreciate that I have only discussed one set of studies concerning object-recognition. The surprising fact of the matter is that proponents of affective realism actually don't cite many experiments along these lines. The only other notable exception is the one carried out by Wormwood *et al.* (2019), which aimed to determine the role of anger on threat perception. This study, like the previous one, involved a modified version of the shooter trials. In this instance, the main difference was the effects of anger were measured during the week of the Boston Marathon bombings and then one and five months later. It was found that participants were less sensitive to the difference between threatening and non-threatening objects during the week of the bombings but not afterward. What is notable is that although these results were said to confirm an effect on "threat perception", they remain ambiguous between misjudgment and misperception for the same reasons as the ones mentioned above.

At present, the lesson, then, is this. If affective realism occurs because of a type of misidentification, its proponents do provide the right sort of explanation of police shooting errors. The problem, however, is that the misidentification hypothesis itself simply isn't supported by the available empirical findings, and *ergo* neither is the explanation of police shooting errors which employs this hypothesis. The empirical data, it turns out, underdetermines whether police officers literally see guns when none are present or merely judge that they do.¹⁵

3.3 Predictive processing

Thus far, we have seen that affective realism is associated with two hypotheses: a misattribution hypothesis where affective features of our interoceptive experiences are misattributed to our visual experiences, and a misidentification hypothesis where such features play a role in object-recognition. Proponents of affective realism employ the misidentification hypothesis in their explanation of police shooting errors, but as we have seen, neither this nor the misattribution hypothesis is directly supported by the experimental findings cited in favor of affective realism. This raises the question, Why do proponents of affective realism take the empirical data to support their view? Why, in particular, do they read the relevant data as confirming a form of misidentification that involves misperception and not just misjudgment?

My diagnosis is that proponents of affective realism have a tendency to interpret the relevant data within a predictive processing framework, which licenses a misperception interpretation of these findings. Predictive processing (PP), very roughly, is “the theory that the brain is a sophisticated hypothesis-testing machine, which is constantly involved in minimizing the error of its predictions of the sensory input it receives from the world” (Hohwy, 2013, p. 1).¹⁶ According to this theory, the main task of the brain is to minimize prediction errors, which involves minimizing the mismatch between predictions of sensory inputs generated internally and actual sensory inputs received externally. Most proponents of predictive processing model this process using a hierarchy of Bayesian modeling. The rough idea is that the brain makes predictions in accordance with the principles of Bayesian inference, which is a statistical method for eliminating uncertainty by combining existing information with new evidence. Moreover, this process is hierarchical because it occurs at multiple levels, with predictions generated at higher levels influencing what happens at the levels below them.

Proponents of PP aim to explain perception, cognition and action via this framework. For instance, on a PP account of perception, perceptual processing involves making predictions about incoming sensory data in a way that minimizes prediction errors. Moreover, this is achieved by our perceptual systems approximating Bayesian principles in a hierarchical setting. Proponents of PP take this to imply that what we see, hear, touch, taste, smell and feel, in effect, are constructs of the brain’s best guess at incoming sensory information.¹⁷ To be clear, they don’t deny that the brain takes new information from the environment into account, but according to them, it does so only when such information is in conflict with what is predicted. Once this happens, a prediction error is generated and the brain updates its predictions, and then constructs what we see, hear, taste, feel and so on based on these new predictions.

The PP framework is relevant for affective phenomena in two respects. First, PP is employed as part of the story of how emotions form. For example, consider Barrett’s (2017) constructionist theory of emotion. As in the case of perception, this theory models interoception using predictive processing: the brain predicts incoming sensory information from our body, and our affective states are constructed by this process. These affective states are a key ingredient in how we construct emotion, but they don’t suffice to bring about emotions by themselves. Emotions are produced when affective states are combined with various complex high-level processing, in particular, when they are categorized using our emotion concepts. This process of categorization itself is once again modeled based on predictive processing. There is plenty more to say about the theory of constructed emotion,

but for now, what is relevant is that it employs PP in its explanation of how emotions form.

Proponents of affective realism, e.g., Barrett and her various collaborators, tend to subscribe to the theory of constructed emotion, but it is important to recognize that affective realism isn't entailed by this theory. To reiterate, the theory of constructed emotion is a theory about how emotions form, not how emotions affect perception. Affective realism, as we have seen, concerns the latter. However, like the theory of constructed of emotion, proponents of affective realism employ PP in their explanation of the phenomenon in question. This is the second way PP is relevant for affective phenomena. Here the key idea is that "affective predictions" play a role in object-recognition:

We suggest that the brain's prediction about the meaning of visual sensations of the present includes some representation of the affective impact of those (or similar) sensations in the past. An effective prediction, in effect allows the brain to anticipate and prepare to act on those sensations in the future. Furthermore, affective predictions are made quickly and efficiently, only milliseconds after visual sensations register on the retina. From this perspective, sensations from the body are a dimension of knowledge—they help us to identify what the object is when we encounter it, based, in part, on past reactions (Barrett & Bar, 2009, 1326).

As I understand them, Barrett and Bar take affect to signal an object's value, which in turn influences object-recognition. As we saw earlier, affect is constructed from the brain's best guess at incoming sensory information from our body, whereas percepts are constructed from the brain's best guess at incoming sensory information from our environment. What's new is the hypothesis that the former influences the latter. To elaborate, Bayesian models are hierarchical. When it comes to perceptual processing, we see that the predictions the brain makes about interoception influence the predictions it makes about perception. This, in a nutshell, is the PP story about how affect shapes perception.

What is significant about the application of PP for our purposes is that affective influences on perception are part of the story we tell about how the brain constructs our percepts. In this way, affect shapes not just our perceptual judgments but our perceptual experiences as well. Moreover, this in turn gives us a unique perspective from which to view the empirical findings. If our prior expectations shape what we literally see, then any affect-based predictions might also influence the contents of our perceptual experiences. Significantly, this makes it plausible that participants primed with affective stimuli misidentify objects because they actually misperceive them. Consider the affect-laden object-recognition tasks mentioned earlier. Any of these experiments can now be plausibly interpreted as ones that involve misidentification by way of misperception.

For example, as we saw earlier, Baumann and DeStano demonstrate that anger increases the probability that subjects will “misidentify” neutral objects as guns. I have argued that these findings underdetermine whether participants misperceive or misjudge neutral objects as guns, whereas Baumann and DeStano themselves opt for a misjudgment interpretation. By contrast, assuming a PP framework, Fridman et al. (2019) take these findings to demonstrate affect-based influences on perception experience itself:

[I]ndividuals induced to experience an instance of anger were more likely to exhibit biased perceptual decision making in a gun detection task, such that they were more likely to make misidentification errors “seeing” unarmed individuals as armed than vice versa (Baumann & DeSteno, 2010). Critically, this biased perception was causally explained by anger’s influence on predictions: angry participants expected to encounter more armed suspects, and controlling for these expectancies mitigated the impact of anger on threat perception (Fridman et al., 2019, p. 6).

Building on Barrett and Bar’s claims about affective predictions influencing object-recognition, Fridman *et al.* interpret the empirical data as confirming that affective predictions can result in participants “seeing” neutral objects as guns. Granted, the use of scare-quotes is confusing. However, elsewhere, their commitment to affective realism is a lot clearer. For instance, they note that “affective experience critically shapes what we *expect to* and *actually do* see, hear, and smell” (pg. 5, original italics).

More generally, PP explanations of affective realism are ubiquitous, and proponents of affective realism tend to provide accounts of police shooting errors within such a framework. This is evident in the way Barrett and Wormwood explain police shooting errors:

The brain is a predictive organ. A majority of your brain activity consists of predictions about the world — thousands of them at a time — based on your past experience . . . These neural “guesses” largely shape what you see, hear and otherwise perceive . . . In every moment, your brain consults its vast stores of knowledge and asks, “The last time I was in a similar situation, what sensations did I encounter and how did I act?” . . . If you are in a part of town with a high crime rate, your brain may well predict a weapon (Barrett & Wormwood, 2015, pp. 1–2).

The take home message here is quite striking. Proponents of affective realism present their explanation of police shooting errors as one that is backed by the latest empirical findings. Nevertheless, the data they cite actually don’t, in and of themselves, support their explanation. Instead, this explanation is best seen as stemming from a series of experimental data interpreted from the point of view of PP.

To be fair, advocates of PP take this theory itself to be supported by the latest science. Assessing this claim is beyond the scope of this paper. For now, it suffices to make two points. First, we can make a distinction between PP being a fruitful way of modeling the subpersonal processes involved in

perceptual processing and PP being a way to understand the nature of perceptual experience itself. Its advocates make both claims, but by and large, the implications of PP for perceptual experience has proved controversial.¹⁸ Second, even setting these issues aside, most discussions of PP tend to focus on the way cognition shapes the content of perceptual experience. Its proponents, of course, claim that affective states can also shape this content. Miller and Clark (2018), for instance, draw on affective realism by way of arguing that “affect and content must be co-computed” (pg. 2571). However, as we see, this isn’t something actually supported by the available data. If anything, proponents like Miller and Clark get the order mixed-up. They employ affective realism to motivate a PP interpretation of various affective phenomena without seeming to recognize that affective realism itself can only be justified by assuming PP!

There is plenty more to say about the connection between PP and affective realism, but to summarize the points so far, the misidentification hypothesis, which proponents of affective realism assume to explain police shooting errors, isn’t directly supported by the empirical data. Rather, we arrive at this hypothesis when we interpret such findings through the lens of PP. I don’t take this to be a knock-down objection to affective realist explanations of police shooting errors. (PP, after all, might prove to be the correct view of perception). However, it shows us that such explanations are, as a matter of fact, much more controversial and theory-laden than they initially seem.¹⁹

4. Implications

The motivation for this work has been practical. Proponents of affective realism provide an explanation of police shooting errors that seeks to radically revise our understanding of police culpability, with significant moral and legal ramifications. What’s more, they present their explanation as one that is backed by the latest science. I think proponents of affective realism should be lauded for pushing further the academic debate on this controversial matter, but I also think the way they have done so is misleading in the following respects.

First, most pressingly, it is simply not true that the affective realist explanation of police shooting errors is backed by the latest empirical data. Such data can only be seen to back this when viewed through the lens of predicting processing. To be clear, predicting processing as a research paradigm is on the rise. However, whether it entails the sorts of claims about perceptual experience its proponents suppose remains a matter of debate. Subsequently, the claim that police officers might literally see neutral objects as guns in high-pressure situations isn’t purely an empirical

one. This needs to be conveyed to the wider public, as well as policymakers, lawyers, and police officers themselves.

Second, there are actually two possible ways affective realism might be relevant for police shooting errors. If the misattribution hypothesis is correct, affective features of our interoceptive experiences can be misattributed to our visual experiences. This suggests that police officers might misattribute the negative feelings they undergo to what they see in their environment and thereby judge neutral objects as threatening. By contrast, if the misidentification hypothesis is correct, affective features of our interoceptive experiences can make us misperceive objects in our environment. This tells us that police officers, when they feel negative, might misperceive neutral objects, e.g., wallets, as threatening ones, such as guns. Both hypotheses provide us with an explanation of police shooting errors, but affective realist explanations of the phenomenon typically assume the misidentification hypothesis. This means proponents of affective realism have missed a trick in their attempts to employ the arsenal of affective realism to explain police shooting errors. Affective states, like feelings, emotions and moods, might help explain some instances of police shooting errors, but not necessarily because such states cause officers to literally see things like wallets and cellphones as guns.

Third, while neither hypothesis is supported by the empirical data by themselves, we also saw that the bulk of the empirical data for affective realism concerns the misattribution hypothesis. This suggests that if proponents of affective realism want to explain police shooting errors in an empirically robust manner, their best bet is to focus on the explanations made available by this hypothesis instead of the one they have focused on so far. In other words, instead of promoting the idea that police officers make errors because they see harmless objects as weapons, they should instead focus on exploring the ways affect might influence threat detection.

To elaborate, there is a subtle but important distinction to be drawn between the claim that affect can make us experience neutral objects as threatening and the claim that affect can make us experience neutral objects as threatening objects. Proponents of affective realism blur this distinction when they cite empirical findings relevant to both claims by way of supporting an explanation of police shooting errors that actually assumes only the second claim. A careful examination of the empirical literature finds both claims to be questionable. Moreover, of the two claims, we see that the findings, if anything, lend more credibility to the first claim; the one that doesn't imply that police officers mistakenly shoot unarmed suspects because they literally see harmless objects as weapons.

So what does all this hold for the moral and legal ramifications of affective realist explanations of police shooting errors? Should we take them at face value? A radical rewrite of how we approach police culpability stems from

the misidentification hypothesis: in high-pressure situations, police officers literally see guns when none are present, which in turn significantly diminishes their moral and legal responsibility. These surprising results of the affective realist research paradigm, however, are not empirically born out. For the time being at least, I we should, therefore, hold off taking them at face value. This, of course, is not to deny that affect might play a role in police shooting errors or that this in turn might make us recalibrate our assumptions about police culpability. These remain open questions. What becomes clear, however, is that those affected by police shooting errors, the victims and police officers themselves, are owed a more careful discussion of these issues than that which we have been able to provide thus far.

Notes

1. While there is no standardized official government data on the matter (Fyfe, 2002; Shane, 2016), data collected by *The Washington Post* since 2020 show the numbers to be relatively stable and near the thousand-mark. The precise number of cases attributed to shooting errors is even harder to discern as police departments don't typically make incident reports publicly available. However, shooting errors are frequently used as a defense for the cases that are reported, as in the deaths of Tamir Rice in 2012, Philando Castile in 2016, and Stephon Clark in 2018, just to mention a few examples (B.B.C, 2020).
2. Ibid. Also see Peebles (2020) and Haddad (2021).
3. e.g., see Alpert (2012), Novy (2012), and Fachner and Carter (2015). See Barrett and Wormwood (2015) and Fridman *et al.* (2019) for a discussion.
4. e.g., see Anderson *et al.* (2012), Barrett and Wormwood (2015), Barrett (2017) and Wormwood *et al.* (2019).
5. Also see Barrett (2017), Wormwood *et al.*, (2019), E. H. Siegel *et al.* (2018) and Fridman *et al.* (2019).
6. As I shall discuss in Sect. 2, proponents of affective realism question this claim. For a general discussion of the malleability of perception, see Zeimbekis and Raftopoulos (2015) and S. Siegel (2017).
7. For a detailed exposition of affect, see Russell and Barrett (1999), Russell (2003), Barrett and Bliss-Moreau (2009) and Barrett (2017).
8. e.g., see Calvert *et al.* (2004) and Stokes *et al.* (2015).
9. As Siegel observes, "On the prevailing view, both in philosophy and in law, what it's reasonable for people to believe depends in part on how it's reasonable to respond to the way things look to them" (S. Siegel, 2020, p. 12).
10. Personally, I don't find such arguments persuasive, but they are commonplace in the implicit bias literature. See Washington and Kelly (2016) for a discussion.
11. We see this played out in various debates in the philosophy of perception, e.g., Firestone and Scholl (2015) and Zeimbekis and Raftopoulos (2015), as well as the philosophy of mind more broadly, e.g., Nagel (1974) and Chalmers (1995).
12. For a review, see Correll *et al.* (2014), Kahn and McMahon (2015) and Mekawi and Bresin (2015).

13. Siegel's example is the shooter trials done by Payne (2001), but these hypotheses work equally well in the context of the modified shooter trials undertaken by Baumann and DeStano.
14. In general, affective realists don't explore potential attentional biases on threat perception, e.g., see Phelps et al. (2006) and Preciado et al. (2017).
15. As a referee helpfully explains, the situation might actually be worse for affective realism. Some empirical data suggests that threat actually enhances perceptual sensitivity, e.g., see de Voogd et al. (2022), as well as the attentional studies mentioned earlier. If such findings are representative, they help make an even stronger case against the misidentification hypothesis. That is, not only is the hypothesis not supported by the data affective realists themselves cite – which is what I try to show here —, some empirical data appear to directly contradict it. Exploring this stronger line of criticism is something I shall leave for a further date.
16. Also see Friston (2005), Clark (2016), Hohwy (2020) and Mendonça et al. (2020).
17. e.g., see Frith (2007), Clark (2016, 2018) and Wilkinson (2014).
18. See Jenkin and Siegel (2015), Newen et al. (2017), Macpherson (2017) and Drayson (2018).
19. It is also worth bearing in mind that the truth of PP doesn't ensure that the misidentification hypothesis is true. It's proponents would still need to address the stronger line of criticism not explored here, i.e., that there is evidence, e.g., de Voogd et al. (2022), to suggest that threat actually enhances perceptual sensitivity.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Raamy Majeed  <http://orcid.org/0000-0001-8864-1643>

References

- Alpert, G. P. (2012). Working toward the truth in officer-involved shootings. *FBI: Law Enforcement Bulletin* (J. Rivera & L. Lott, Eds.), 1–7. <https://leb.fbi.gov/articles/featured-articles/working-toward-the-truth-in-officer-involved-shootings>
- Anderson, E., Siegel, E., White, D., & Barrett, L. F. (2012). Out of sight but not out of mind: Unseen affective faces influence evaluations and social impressions. *Emotion*, 12(6), 1210–1221. <https://doi.org/10.1037/a0027514>
- Barrett, L. F. (2017). *How emotions are made*. Pan Books Ltd.
- Barrett, L. F., & Bar, M. (2009). See it with feeling: Affective predictions during object perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364 (1521), 1325–1334. <https://doi.org/10.1098/rstb.2008.0312>
- Barrett, L. F., & Bliss-Moreau, E. (2009). Affect as a psychological primitive. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 41, pp. 167–218). Academic Press.
- Barrett, L. F., & Wormwood, J. (2015, April 17). Opinion | When a gun is not a gun. *The New York Times*. <https://nyti.ms/1GUPAZN>

- Baumann, J., & DeSteno, D. (2010). Emotion guided threat detection: Expecting guns where there are none. *Journal of Personality and Social Psychology*, 99(4), 595–610. <https://doi.org/10.1037/a0020665>
- BBC. (2020, June 3). Timeline of US police killings. *BBC News*. <https://www.bbc.com/news/world-us-canada-52905408>
- 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. <https://doi.org/10.1037/1528-3542.8.4.479>
- Calvert, G., Spence, C., & Stein, B. E. (2004). *The handbook of multisensory processes*. MIT Press.
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Clark, A. (2016). *Surfing uncertainty : Prediction, action, and the embodied mind*. Oxford University Press.
- Correll, J., Hudson, S. M., Guillermo, S., & Ma, D. S. (2014). The police officer's dilemma: A decade of research on racial bias in the decision to shoot. *Social and Personality Psychology Compass*, 8(5), 201–213. <https://doi.org/10.1111/spc3.12099>
- de Voogd, L. D., Hagenberg, E., Zhou, Y. J., de Lange, F. P., & Roelofs, K. (2022). Acute threat enhances perceptual sensitivity without affecting the decision criterion. *Scientific Reports*, 12(1), 9071. <https://doi.org/10.1038/s41598-022-11664-0>
- Drayson, Z. (2018). Direct perception and the predictive mind. *Philosophical Studies*, 175(12), 3145–3164. <https://doi.org/10.1007/s11098-017-0999-x>
- Fachner, G., & Carter, S. (2015). *An assessment of deadly force in the Philadelphia police department*. Collaborative Reform Initiative, Office of Community Oriented Policing Services. <https://www.phillypolice.com/assets/directives/cops-w0753-pub.pdf>
- Firestone, C., & Scholl, B. J. (2015). Cognition does not affect perception: Evaluating the evidence for “top-down” effects. *The Behavioral and Brain Sciences*, 39. <https://doi.org/10.1017/s0140525x15000965>
- Fridman, J., Barrett, L. F., Wormwood, J. B., & Quigley, K. S. (2019). Applying the theory of constructed emotion to police decision making. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.01946>
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Frith, C. D. (2007). *Making up the mind : How the brain creates our mental world*. Blackwell Pub.
- Fyfe, J. J. (2002). Too many missing cases: Holes in our knowledge about police use of force. *Justice Research and Policy*, 4(1–2), 87–102. <https://doi.org/10.3818/jrp.4.1.2002.87>
- Haddad, M. (2021). *How many people have been killed by US police since George Floyd?* <https://www.aljazeera.com/news/2021/5/25/how-many-people-have-police-killed-since-george-floyd>
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
- Hohwy, J. (2020). New directions in predictive processing. *Mind & Language*, 35(2), 209–223. <https://doi.org/10.1111/mila.12281>
- Jenkin, Z., & Siegel, S. (2015). Cognitive penetrability: Modularity, epistemology, and ethics. *Review of Philosophy and Psychology*, 6(4), 531–545. <https://doi.org/10.1007/s13164-015-0252-5>
- Kahn, K. B., & McMahon, J. M. (2015). Shooting deaths of unarmed racial minorities: Understanding the role of racial stereotypes on decisions to shoot. *Translational Issues in Psychological Science*, 1(4), 310–320. <https://doi.org/10.1037/tps0000047>

- Macpherson, F. (2017). The relationship between cognitive penetration and predictive coding. *Consciousness and Cognition*, 47, 6–16. <https://doi.org/10.1016/j.concog.2016.04.001>
- Mekawi, Y., & Bresin, K. (2015). Is the evidence from racial bias shooting task studies a smoking gun? Results from a meta-analysis. *Journal of Experimental Social Psychology*, 61, 120–130. <https://doi.org/10.1016/j.jesp.2015.08.002>
- Mendonça, D., Curado, M., & Gouveia, S. (2020). *The philosophy and science of predictive processing*. Bloomsbury Publishing Plc Bloomsbury Academic.
- Miller, M., & Clark, A. (2018). Happily entangled: Prediction, emotion, and the embodied mind. *Synthese*, 195(6), 2559–2575. <https://doi.org/10.1007/s11229-017-1399-7>
- Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, 83(4), 435–450. <https://doi.org/10.2307/2183914>
- Newen, A., Marchi, F., & Brössel, P. (2017). Cognitive penetration and predictive coding. Pushing the debate forward with the recent achievements of cognitive science. *Consciousness and Cognition*, 47, 1–112. <https://doi.org/10.1016/j.concog.2016.12.001>
- Novy, M. (2012). Cognitive distortions during law enforcement shooting. *Activitas Nervosa Superior*, 54(1–2), 60–66. <https://doi.org/10.1007/bf03379584>
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, 81(2), 181–192. <https://doi.org/10.1037/0022-3514.81.2.181>
- Peebles, L. (2020). What the data say about police brutality and racial bias — and which reforms might work. *Nature*, 583(7814), 22–24. <https://doi.org/10.1038/d41586-020-01846-z>
- Phelps, E. A., Ling, S., & Carrasco, M. Emotion facilitates perception and potentiates the perceptual benefits of attention. (2006). *Psychological Science*, 17(4), 292–299. PMID: 16623685; PMCID: PMC1555625. <https://doi.org/10.1111/j.1467-9280.2006.01701.x>
- Preciado, D., Munneke, J., & Theeuwes, J. Was that a threat? Attentional biases by signals of threat. (2017). *Emotion*, 17(3), 478–486. Epub 2016 Nov 7. PMID: 27819442. <https://doi.org/10.1037/emo0000246>
- Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, 110(1), 145–172. <https://doi.org/10.1037/0033-295x.110.1.145>
- Russell, J. A., & Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, 76(5), 805–819. <https://doi.org/10.1037/0022-3514.76.5.805>
- Shane, J. M. (2016). Improving police use of force: A policy essay on national data collection. *Criminal Justice Policy Review*, 29(2), 128–148. <https://doi.org/10.1177/0887403416662504>
- Siegel, S. (2017). *The rationality of perception*. Oxford University Press.
- Siegel, S. (2020). Bias and Perception. In E. Beeghly & A. Madva (Eds.), *An introduction to implicit bias: Knowledge, justice, and the social mind*. Routledge. PhilPapers; Routledge. <https://philpapers.org/rec/SIEPAI-2>
- Siegel, E. H., Wormwood, J. B., Quigley, K. S., & Barrett, L. F. (2018). Seeing what you feel: affect drives visual perception of structurally neutral faces. *Psychological Science*, 29(4), 496–503. <https://doi.org/10.1177/0956797617741718>
- Stokes, D., Matthen, M., & Biggs, S. (2015). *Perception and its modalities*. Oxford University Press.
- Washington, N., & Kelly, D. (2016). Who's responsible for this? Moral responsibility, externalism, and the knowledge about implicit bias. In M. Brownstein & J. Saul (Eds.), *Implicit bias and philosophy Volume 2: Moral responsibility, structural injustice, and ethics* (pp. 10–36). Oxford University Press.

- The Washington Post. (2020, January 22). Every fatal police shooting since 2015. *The Washington Post*. <https://www.washingtonpost.com/graphics/investigations/police-shootings-database/>
- Wilkinson, S. (2014). Accounting for the phenomenology and varieties of auditory verbal hallucination within a predictive processing framework. *Consciousness and Cognition*, 30, 142–155. <https://doi.org/10.1016/j.concog.2014.09.002>
- Wormwood, J. B., Siegel, E. H., Kopec, J., Quigley, K. S., & Barrett, L. F. (2019). You are what I feel: A test of the affective realism hypothesis. *Emotion*, 19(5), 788–798. <https://doi.org/10.1037/emo0000484>
- Zeimbekis, J., & Raftopoulos, A. (2015). *The cognitive penetrability of perception: New philosophical perspectives*. Oxford University Press.