



Perceiving deviance

Eli Shupe¹

Received: 16 July 2019 / Accepted: 22 November 2019
© Springer Nature B.V. 2019

Abstract

I defend the claim that we have the capacity to perceptually represent objects and events in experience as deviating from an expectation, or, for short, as deviant. The rival hypothesis is that we may ascribe the property of deviance to a stimulus at a cognitive level, but that property is not a representational content of perceptual experience. I provide empirical reasons to think that, contrary to the rival hypothesis, we do perceptually represent deviance.

Keywords Perception · Perceptual content · Rich content view · Philosophy of mind

1 Introduction

It is widely accepted that we perceptually represent basic, low-level properties like color and shape. Some philosophers have argued that these are the only properties we perceptually represent (Byrne 2009; Dretske 1995; Tye 1995). But many others argue that we also perceptually represent certain high-level properties, such as causal relations (Siegel 2009), kind properties like *being a pine tree* (van Gulick 1994; Siegel 2006), affordances like *being edible* (Nanay 2011), and even normative properties like *being wrong* and *being unjust* (Cowan 2015; Werner 2016). We can call the basic claim that high-level properties are represented in perception the Rich Content View.

The Rich Content View is an existential claim about the representational contents of perception. All we need to do in order to show that the Rich Content View is true is to show that at least one high-level property is represented in perceptual experience. A clear way forward, then, is to identify high-level properties that are good candidates for being represented in perception, and to then scrutinize whether they in fact are so represented. To this end, in this paper, I defend—on largely empirical grounds—the claim that a particular high-level property is among those that are represented in perceptual experience. My thesis, call it DEVIANCE, is this:

✉ Eli Shupe
ebshupe@gmail.com

¹ Department of Philosophy, Rutgers, New Brunswick, USA

DEVIANCE: We have the capacity to perceptually represent objects and events in experience as deviating from an expectation, or, for short, as deviant.

Here I mean ‘expectation’ in a very thin sense. An expectation needn’t be consciously occurrent—it may exist simply as a background assumption or implicit prediction. The way in which we form and maintain some of these expectations might even be entirely sub-personal.

DEVIANCE stands in contrast to the idea that we only ever *infer* that an object or event deviates from our expectations; or, alternately, that we do sometimes recognize deviance non-inferentially, but that the methods by which we do so are exclusively non-perceptual—perhaps the process is facilitated by affective states, for example.

My hope is that my project in this paper serves as a unique contribution to a well-worn debate. Not only do I argue in favor of the perceptual representation of a novel high-level property, but my arguments are supported by data from neuroscience that have not previously been brought to bear on the debate about perceptual contents. If my arguments are successful, then my methodology might generalize to defenses of the presence of other high-level perceptual contents.

In Sect. 2, I explain the DEVIANCE thesis at greater length and introduce a pair of contrast cases, *Mistake* and *Control*, the former of which involves the phenomenal experience of DEVIANCE and the latter of which does not. In Sect. 3, I show how a phenomenal contrast argument in defense of DEVIANCE using *Mistake* and *Control* might proceed. In particular, and continuing into Sect. 4, I discuss how evidence from neuroscience might come to bear on our inferences to the best explanation about the phenomenal difference between *Mistake* and *Control*. Section 5 discusses one candidate EEG study (Pieszek et al. 2014) that suggests the distinctive phenomenology of *Mistake* is best explained by DEVIANCE. Sections 6 and 7 defend a pair of claims about the responses elicited in the study; if both of these claims are true, then the study and others like it form a solid evidence base in favor of DEVIANCE.

2 Mistake and control

Experiences of unexpected objects and events commonly share certain phenomenal features. Often such experiences elicit the sensation of *surprise*. Sometimes the elicitation is weaker: an unexpected stimulus might fail to palpably shock a perceiver but still evoke feelings of unease or a more general impression of *ill-fit*. This weaker elicitation seems to be a common factor. Although not all experiences of violated expectations are accompanied by the sensation of surprise, the more minimal sensation of ill-fit does seem ubiquitous across examples (both when surprise phenomenology is present and when it is absent). Because I take it to be more fundamental in this way, it is the sensation of ill-fit that I seek to isolate and explain as the characteristic feature of experience of deviance.¹

¹ A number of ethicists have proposed fittingness as a basic normative property (Mandelbaum 1955; Chappell 2012; Audi 2013, 2015; McHugh and Way 2016). Some of these philosophers have further conceived of it as an apprehended property, and have argued that felt moral demands have their basis in the phenomenology of apprehending the fittingness or ill-fittingness of an act (Mandelbaum 1955; Audi 2013, 2015). There are commonalities between these latter philosophers and myself. Although I am con-

More specifically, I will focus on the experience of having one's *auditory* expectations violated. Since DEVIANCE being true of one modality is enough to make DEVIANCE true *simpliciter*, we can defensibly restrict our focus in this way. My exemplar case, *Mistake*, is one in which the final note of a musical sequence sounds 'wrong' to the listener. I contrast *Mistake* with *Control*, a case in which the same sequence of notes is heard but, lacking context, the listener does not experience the final note in the sequence as sounding wrong. After introducing these cases, I will explore the argument that what elicits the phenomenology of ill-fit in one case but not the other is a difference in what is perceptually represented. Here are the cases:

Control: Jenna is listening to the sound of a neighbor's piano lesson. She hears a musical triplet being played (A-B-C). After a pause, she hears another slightly different triplet (A- B-D).

Mistake: Jenna is listening to the sound of a neighbor's piano lesson. This time, before the music starts, she hears the teacher tell her neighbor to repeat the notes that she plays. Then, Jenna hears a musical triplet being played (A-B-C). After a pause, she hears another slightly different triplet (A-B-D).

Jenna perceives all of the same low-level auditory content in *Control* and *Mistake*. However, as will be familiar to the reader, in cases like *Mistake*, but not *Control*, we tend to experience the misplayed note (D) as sounding wrong. Or, in other words, we experience the sensation of ill-fit described previously. The phenomenology is the same when we listen to a singer whose voice cracks. That the note sounds wrong in one but not both of the cases above is taken to be an intuitive explanandum. DEVIANCE can be evaluated partly on the basis of whether it is a serviceable explanans, i.e. whether it is the thesis that best explains the phenomenal difference in the pair of contrast cases.

3 The phenomenal contrast argument

By far the most popular strategy for defending positive theses about the contents of perception has been the phenomenal contrast method (Siegel 2007). Once a candidate thesis like DEVIANCE is proposed, the strategy begins with the identification of two cases like *Mistake* and *Control*, i.e. two experiences that involve the representation of all of the same low-level perceptual content, but which have different overall phenomenologies. Then, the cases are used to perform an inference to the best explanation: we ask, does the addition of some particular high-level perceptual content best explain the distinctive phenomenology of *Mistake*?

To illustrate, a phenomenal contrast argument for DEVIANCE that proceeds from *Mistake* and

Footnote 1 (continued)

cerned with more basic, non-moral violations (such as misplayed notes), I, too, think that the detection of such violations is characterized by a phenomenology of ill-fit.

Control would go roughly as follows.

1. *Control* and *Mistake* differ in overall phenomenology.
2. This change in overall phenomenology is best explained as a change in the phenomenology of the auditory experience that is part of the overall experience.
3. This change in the phenomenology of the auditory experience is best explained as a change in the representational content of the auditory experience.
4. On the best explanation of (3), the change in representational content is that in *Mistake*, but not *Control*, Jenna auditorily represents that the D-note is deviant.

To expand on this, the explanation on offer is that, when you have no beliefs or other expectations about which notes an incoming triplet will contain, as in *Control*, then when you hear the triplet, auditory experience only represents the notes the triplet actually contains (i.e., the stimulus' low-level properties). But once you form a prior belief or expectation about which notes an incoming triplet will contain, as in *Mistake*, auditory experience then not only represents the notes the triplet actually contains, but also represents *whether those notes were the notes that you expected to hear* (a high-level property).

Remember, however, that on DEVIANCE's rival views one only ever infers that a note is deviant, or recognizes deviance by some other non-perceptual means. Proponents of these rival views will reject the argument above and claim that Jenna perceptually representing deviance is *not* the best explanation of the difference in her overall phenomenal state. Instead, these theorists will argue that Jenna's *auditory* experience is the same in both cases and will point to changes in one or more of Jenna's non-perceptual states as making for the overall phenomenal difference. Are any of these alternate explanations better than the explanation I provide?

This question is often adjudicated from the armchair. The best explanation is said to be the one that is the most explanatorily parsimonious, for example, or the least vulnerable to conceptual counterexamples. Instead of examining these well-worn considerations, I would like to take a somewhat different approach. Setting aside the purely theoretical advantages that one hypothesis about the contents of perception might have over another, my question will be: What role might *empirical* considerations play in deciding which explanation is best?

At the very least, I believe that empirical evidence should help to delineate the criteria of admissibility for candidate explanations. The best explanation of Jenna's phenomenal contrast will be the best explanation among those that are also consistent with our best science of the mind. This is, I think, a relatively unambitious thesis. We might press it further and argue that, when comparing two theories that are both minimally compatible with the findings of science, we should look to see if either is *better* supported by those findings, and show more favor to theories that enjoy greater support.

Many philosophers who employ the phenomenal contrast method to settle questions about the content of perceptual experience ignore or underemphasize the supplementary role that cognitive science can play in weighing candidate

explanations. Siegel (2011), for instance, discusses only how the Rich Content View (the view that we perceive high-level properties like *being a pine tree*), if it is the better than rival explanations on theoretical grounds, might subsequently guide the inquiries of cognitive science:

Aside from its bearing on philosophical issues, the Rich Content View is relevant to several other areas of research in psychology and neuroscience. [...] Where should we look for neural correlates of conscious visual experience? Whether the Rich Content View is true will influence what we will count as a neural correlate of visual experience. If the view is false, then we might expect to find these neural correlates in brain areas devoted to “early” visual processing, such as visual areas V1 and V5. If the view is true, then we should expect neural correlates to involve “later” areas, such as the fusiform face area (FFA) and the inferotemporal cortex (IT). (Siegel 2011)

But what of the other direction of influence? Might not the inquiries of cognitive science be independently brought to bear on whether we accept the Rich Content View? Scientists are already searching for the correlates of visual experience, and not always with strong prior conceptions about where to look for them. Siegel envisions that discovering the Rich Content View is true would guide the direction of cognitive research. But if the cognitive researcher independently discovered a correlate to the representation of higher-level properties, wouldn't the location of that correlate be precisely the kind of thing that could settle the question of whether such representation is perceptual or not, precisely the kind of thing, in other words, that could help philosophers decide whether the Rich Content View is true?

For instance, imagine an EEG study with two subjects, one with a species-typical brain and the other afflicted by a kind of blindsight. Say that both subjects look at an object and their EEG results show the same increased levels of activity in areas of the brain devoted to “early” visual processing (V1 and V5), but only the non-blindsighter shows increased levels of activity in parts of the brain devoted to “later” perceptual processing (FFA and IT). Furthermore, say that only the non-blindsighter reports visually *experiencing* the object; the blindsighter says that she sees nothing.

The fact that only the subject who showed “later” stage brain activity was conscious of whether they saw the object would be some evidence for the fact that the neural correlate of such experience is to be found in these “later” perceptual processing areas of the brain. And this is precisely where the Rich Content View, but not many its rivals, suggests such a correlate is to be found. So surely Siegel's observation might run in reverse. To the extent that whether the Rich Content View is true or false can tell us *where to look* for neural correlates, then where neural correlates are found might help us decide *whether the Rich Content View is true or false*. The experiment sketched above, for example, would at least lend credence to the Rich Content View, and detract from the plausibility of rival views that associate perceptual experience with only the earlier stages of perceptual processing (and thus the earlier areas of the visual system). In other words, empirical results have a role to play in the inference to the best explanation that decides on a favored theory about the contents of perceptual experience.

Unfortunately for the Rich Content View, blindsight does not work as described above. Rather, it involves damage to the V1 area of the visual system (Barbur et al. 1993). Thus, a blindsighter will *not* enjoy the same levels of early V1 activation as a typically sighted person while differing only in suppressed downstream levels of activation. Rather, blindsighters exhibit diminished V1 activation, and consequently also diminished downstream activation. So a real life experiment along the lines of the one sketched above wouldn't lend support to either the Rich Content View or its rivals. But this is just a toy example—the point was simply that other, feasible experiments might.

4 Empirical considerations

Returning to DEVIANCE, which is a species of the Rich Content View, we may ask: is DEVIANCE consistent with our best science of the mind? Moreover, are there empirical findings that are like the fantasy blindsight case above, ones which lend support to DEVIANCE over and above its rival explanations? I will argue that there are some empirical results that can contribute to our debate over DEVIANCE, namely recent experimental work on error perception in the auditory system.

First, what sort of empirical work would lend support to DEVIANCE? DEVIANCE hypothesizes that the property of *being deviant* is represented in perceptual experience. One way for DEVIANCE to be true is for the property of *being a deviant sound* to be represented in auditory experience. If deviance really is represented in audition in this way, then we should see some evidence that, when we experience sounds that upset our expectations, and which sound wrong, there is some signature brain event that occurs in centers of auditory processing, or those parts of the brain that otherwise give rise to or affect auditory experiences. Furthermore, this activity should be distinct from activity that occurs when we hear sounds that meet our expectations, and which do not sound wrong. If such activity is found, and is indeed localized to parts of the brain that determine what is experienced as heard, then that is consistent with a story on which deviance is being represented in auditory experience, like the other auditory phenomena processed in those areas. In other words, there should be some sort of unique brain state that is seen only when deviant sounds are heard, and which is localized to parts of the brain where differences in activity correlate with differences in auditory experience.

If there were *no* sensory brain event of the kind described above, then critics of DEVIANCE would likely be correct: recognizing deviance would just be a matter of entering some purely cognitive or otherwise completely non-perceptual mental state.

5 Pieszek, Schröger, and Widmann (2014)

Let us turn to one study that may be able to do the work described above, namely an EEG study conducted by Pieszek et al. (2014). I will argue that this study potentially confirms DEVIANCE, in that it appears to show (1) that we do represent whether




Fig. 1 Figure from Pieszek et al. 2014

a heard sound deviates from our expectations, and (2) that such representation is perceptual.

I believe that this study brings us an important distance towards meeting one of the challenges described above, in that it delivers a *candidate* for the neural state that makes for the phenomenal difference when we hear a note that sounds wrong. But whether or not this neural activity will satisfy our explanatory burden will depend on whether or not we consider it *sensory*. If the activity detected is cognitive rather than perceptual, then it actually lends support to the opponent of DEVIANCE, who argues that, although deviance might be represented in experience, it is not represented in *perceptual* experience. This question will be addressed following the discussion of the study.

In the study, Pieszek, Schröger, and Widmann trained 25 subjects to associate particular symbols (white-on-top, white-on-bottom) with particular notes (high-pitched and low-pitched, respectively). Once they were taught these associations, they were then shown a row of five symbols. Then, after a duration of 2 s, subjects were played a five-tone melody (while still looking at the row). Following the melody, subjects were given a two second response window within which to evaluate the congruency of the trial, i.e. whether the notes they expected to hear on the basis of the symbols were the notes they actually heard. Subjects reported their evaluation by pressing one of two buttons, one for success, when the sounds corresponded to the symbols, and one for failure, when the sounds did not correspond to the symbols. Subjects performed this task an average of two hundred times each. Half of the rows were fully congruent with the melodies, and half of the rows contained one incongruent note, making the rate of incongruent symbol-pitch pairs 10%. Figure 1 (below, from Pieszek et al. 2014) helps to illustrate the experimental design.

Subjects were typically very skilled at identifying whether or not the rows were congruent with the melodies. They had an accuracy rate of about 99% in identifying both congruent and incongruent sequences, and there was no difference in response times between the two categories (i.e., subjects did not take longer to identify incongruent sequences than congruent sequences, or vice versa). Indeed, the task is an easy one. Once the association between the visual symbols and the pitches of the notes has been learned, notes that fail to match the given symbols stand out. They are experienced as surprising, and as sounding wrong.

The experimenters analyzed EEG data from subjects as they listened to each note within the melodies. There were two distinct electrical responses, or event-related

potentials (ERPs), that were *only* detected in the brains of subjects currently listening to notes that were incongruent with the accompanying symbols: the IR (Incongruity Response) and the N2b signal. In other words, incongruent note-symbol pairs elicited both responses and congruent note-symbol pairs failed to elicit either. Thus, Pieszek, Schröger, and Widmann identified patterns of neural activity that corresponded to perceivers' experiences of hearing notes that *sounded wrong*. Let's look more closely at the signals the experimenters detected.

5.1 IR (Incongruity Response): a sensory prediction error

The IR is a modality-specific response to deviant auditory stimuli. It is localized to regions of the auditory cortex that are associated with the relatively early stages of auditory processing. The detection of the IR is consistent with a previous experimental finding of several of the authors, which detected signals resembling the IR during pitch violations (Pieszek et al. 2013). Furthermore, the IR can be thought of as an analogue to the more general 'mismatch negativity signal' (MMN), which has been extensively described in the neuroscience literature on sensory perception. The MMN signal is thought to be generated by an automatic response within the brain's sensory system to infrequent discriminable changes in auditory input (Näätänen et al. 2005).

On one of the most commonly accepted models, the predictive coding model, signals like the IR and MMN are generated by a mismatch between sensory input and automatically pre-activated auditory sensory representations (Bar 2007; Friston and Kiebel 2009). On such a model, once one has learned to associate certain symbols with certain sounds, the visual stimulus of a row of symbols actually causes one to form a kind of auditory map, which is represented in audition itself. Before each note has even played, an auditory representation of either the high note or the low note pre-activates in the early auditory cortex, depending on which symbol is seen.

This process serves to maximize the efficiency of the perceptual system, making it easier for the auditory cortex to process incoming stimuli in cases where it guesses correctly about what those stimuli will be like. In cases where the auditory cortex guesses incorrectly, it detects those mistakes at an early stage, and, instead of completing and forwarding its pre-activated representations, it newly forms and forwards correct auditory representations instead (using the feedback it acquires to hone its future predictions).²

It is important to note that the IR is almost certainly *pre-attentive* (and thus likely not conscious). This is indicated by the authors of the study when they describe the IR signal (and MMN signals) as corresponding to the detection of violations "at sensory levels of processing", which they distinguish from the "attentive detection" that occurs at "cognitive- attentive level" of processing (Pieszek et al. 2014).

That early-stage sensory signals like the IR and MMN are typically pre-attentive and pre-conscious has also been shown in other studies, such as Moreau's et al.

² This account is also in line with the functional model of symbol-to-sound match processing described in Widmann et al. (2007).

(2013) study of amusia, a form of auditory agnosia. Amusia is generally associated with deficits in processing pitch, and those afflicted often struggle to recognize and reproduce musical tones (Pearce 2005). In Moreau's, Jolicœur's, and Peretz's study, amusic subjects were played a tone that suddenly changed in pitch, a phenomenon they were unable to detect. Nonetheless, EEG results showed that the subjects' brains did register the difference *pre-attentively* "at early unconscious levels of processing", and that this pre-attentive registration was reflected by the presence of an MMN signal (Moreau et al. 2013). This is the same response that is elicited in non-amusic subjects. If both amusic and non-amusic subjects exhibit MMN responses, as Moreau, Jolicœur, and Peretz have shown, then such signals are *not sufficient* for eliciting the *experience* of deviance; and, because the IR signal is the same kind of signal as the MMN, it cannot then be a candidate correlate brain state to the experience of deviance.

What kind of neural activity does a person with amusia lack, then, when they fail to notice an unexpected change in the tone they are hearing? As it happens, Moreau et al. (2013) found that what their amusic subjects lacked was the exact same response that was elicited in subjects who successfully detected incongruent notes in the Pieszek et al. (2014) study: the N2b response.

5.2 N2b: a later prediction error

The N2b response occurs slightly after the IR, and takes place at what the experimenters call the "cognitive-attentive" level of processing (whether the signal is truly 'cognitive' in the philosopher's sense of the word will be discussed shortly) (Pieszek et al. 2014). At its most general level of description, the N2b response is a signal that reflects the conscious detection of the violation of a prediction about perceptual inputs. In other words, the *deviance* of the sound is selected for attention, draws our attention, and becomes something that we notice (Ibid, 2014; see also Näätänen and Gaillard 1983).

As mentioned previously, the Moreau et al. (2013) study of subjects with amusia also tracked this response and that work showed that neurotypical but not amusic subjects exhibited the N2b during unexpected changes in stimulus pitch.

So, to take stock: we have a perceptual representation of deviance that is not conscious (corresponding to the MMN or IR), and we have a conscious representation of deviance that may or may not be perceptual (corresponding to the N2b response).

It is important to separate the question of what we perceive from the question of what we perceptually experience. It seems safe to say that, at the very least, the results of Pieszek et al. (2014) and others suggest that we *perceive* deviance, but this is only to say that our sensory system tracks it by way of mismatch negativity signals like the MMN and IR. Whether deviance is represented in perceptual *experience*, as the DEVIANCE thesis claims, is a further question. In order to answer that question in the affirmative, I will need to argue for two distinct claims. The first claim is that MMN responses like the IR do indeed single out stimuli as *deviant*, as I have suggested, as opposed to their being possessed of some concatenation of low-level properties. The second claim is that the N2b signal corresponding to the shift

of attention toward the relevant stimulus is best understood as a perceptual rather than cognitive event. If I can successfully argue for each of these claims, then I have made a good case for DEVIANCE. I shall address both of these claims in turn.

6 Do MMNs like the IR track deviance?

It might be argued that the IR and other MMNs do not necessarily single out stimuli *as deviant*. Perhaps, one might argue, they merely single stimuli out as being possessed of some concatenation of low-level properties, namely the low-level properties *in virtue of which* we take them to be deviant.³

The first thing to note here is that the IR and the auditory MMN in general token in response to a wide range of auditory stimuli and in a variety of circumstances, where low-level perceptual features vary greatly. As a rule, the more general the circumstances under which a response is elicited, the more likely it is that a response is tracking a general, abstract property, rather than any particular surface-level feature. Furthermore, the auditory MMN may be caused by not only by deviations from regularity in pitch, as in the studies discussed above, but also by deviations in intensity or duration, and it is more pronounced and occurs more rapidly relative to *how* deviant the stimulus is from what a subject might reasonably expect (Paavilainen 2013).

Secondly, the auditory MMN does not only occur in response to passively experienced violations of regularity; interestingly, it can also be triggered by violations of action intention. In a study by Korka et al. (2019), MMNs were elicited in subjects whose button presses were intended to produce a specific tone, even in a paradigm where the button presses only produced the intended tone 50% of the time (and so the effect cannot be attributable to a failure of regularity). This experimental condition maps nicely onto certain real-world cases of experienced auditory deviance, particularly cases where the subsequent phenomenology of ill-fit is especially pronounced. The singer who botches her opening note, for example, is in an analogous situation: she is thwarted in producing her intended sound, and because of that the note she sings sounds wrong to her.

Finally, the predictive coding model of sensory processing, mentioned previously, is one of the most standard models of sensory processing, and it is *premised* on the idea that the sensory system maximizes efficiency by not constructing sensory representations from the bottom-up, but rather by making educated guesses about incoming sensory inputs and then subsequently picking out—at the earliest stages of perception—those stimuli which *deviate* from those predictions, precisely on the basis of their deviance. As the sensory system learns more and more about the regularities of its environment, it benefits more from catching out falsehoods, so to speak, than it does from freshly constructing truths. The success of the predictive coding model thus itself forms the basis for a strong argument that we should recognize deviance not only as a property that is tracked in perception, but as one that is tracked in very *early* perception.

³ Thank you to an anonymous referee for this journal for pressing me on this point.

7 Is the N2b perceptual?

If the N2b tracks a *perceptual phenomenon*, then the results of Pieszek et al. (2014) provide evidence that deviance is represented in perceptual *experience*. But perhaps the N2b does not track a perceptual phenomenon, but rather simply indexes the phenomenal experience of stimuli being cognitively classified as deviant, the way one might look at a painting and experience cognitively classifying the visual stimuli one perceives as indicative of a Rothko or a Pollock. Therefore, in order to rule out this latter cognitive explanation, and to thus make the N2b at least a strong candidate for grounding the representation of deviance in perceptual experience, we must motivate the idea that the N2b is not merely an experiential phenomenon but a perceptual phenomenon as well. Is it?

On the face of it, it may seem as if the answer is no; after all, Pieszek et al. (2014) discuss N2b as “cognitive”. However, it is important to understand that the authors’ use of “cognitive” is largely contrastive; they mean to differentiate the later N2b signal from the earlier IR, which is pre-attentive, unconscious and occurs at a much earlier stage of auditory processing. Other authors, whose motives are less contrastive, are more neutral in their descriptions and do not take pains to characterize the N2b as cognitive. So this terminological evidence alone is insufficient to establish that the N2b is not perceptual, at least in the sense relevant to philosophers. Indeed, I think we can argue that we should indeed understand the N2b as perceptual in the relevant sense.

First, the deviant stimulus to which attention is directed during an N2b response has, as I explained previously, already been classified as deviant by the IR or MMN early on in perceptual processing. Since the sound has already been classified as deviant by the perceptual system, it is reasonable to assume that it is presented as such in perceptual experience; the alternative is that the sound is stripped of that pre-assigned perceptual value and instead merely presented *simpliciter* in perceptual experience, whereupon it is *subsequently* re-categorized as deviant at the level of cognitive processing. Though not impossible, this seems unlikely.

Given the early perceptual processing of deviance, then, the argument that deviance is *not* represented in perceptual experience proceeds with much more difficulty than similar arguments concerning other high-level properties like that of *being a pine tree*. The analogous situation would be if we took *that* dispute, and then added to it the miraculous discovery of the existence of tiny, pine-tree-specific receptors in the early visual system. Were that to happen, I think it would be fair to say that Susanna Siegel and others who would forest our perceptual experience with such properties would have gained a significant edge in the debate.⁴

⁴ This argument as well as other things I have said previously might reasonably invite the question of whether deviance is actually a previously undiscussed *low-level* perceptual property, rather than a candidate high-level perceptual property. I think that to interpret deviance in this fashion would be mistaken, not to mention highly revisionary with respect to the philosophical literature’s current understanding of high- and low-level perceptual properties. However, I will save my arguments to that effect for a future paper.

The most plausible explanation for the empirical and experiential data is that the deviant sound is not merely noticed as deviant and then redundantly classified once more as such via an inference or judgment, the way one might notice a painting and subsequently identify it as a Pollock; rather, it is most probable that the sound is heard as deviant from the very beginning.

There is one final argument for claim that the N2b corresponding to the perceptual experience of deviance. This argument asks us to consider the nature of attention, to which a shift in which toward the deviant stimulus the N2b corresponds. Although the shift in attention towards deviant sounds tracked by the N2b can occur partly as a result of top-down, cognitive influences (e.g. beliefs and expectations about what sounds one is likely to hear or produce), it can be argued that the mere presence of these influences does not merit describing the attentional shift itself as a cognitive phenomenon. This is because this attentional shift is a shift in *auditory* attention, and any shift in auditory attention is a perceptual phenomenon, whatever its causes. What it is to auditorily attend to a stimulus just is for the auditory system to be preferentially devoted to processing that stimulus. As Christopher Mole writes, “The processes responsible for the allocation of attention [are] inextricable from the processes that are responsible for the perception of the things to which we attend” (Mole 2015). Thus, to the extent that the N2b tracks the direction of attention within a sensory modality, then, it may be considered a perceptual phenomenon, regardless of the contributions of top-down influences.

8 Taking stock

I hope to have made a good case that (1) ERPs like the MMN and IR track the representation of deviance in perception; and (2) later ERPs associated with conscious attention like the N2b track the representation of deviance in perceptual experience. If I am right, then studies like the one conducted by Pieszek et al. (2014) provide empirical support for DEVIANCE, the thesis that we have the capacity to perceptually represent objects and events in experience as deviating from our expectations.

Acknowledgements I would like to thank two anonymous referees at this journal, both of whom were thoughtful, prompt, and kind. I also owe debts to Sam Carter and Mia Accomando for extensive and invaluable feedback. Finally, I must express my deepest gratitude to two Susannas: Susanna Schellenberg, who provided essential feedback during this paper’s development; and Susanna Siegel, without whose work on the contents of perception this paper would never have existed.

References

- Audi, R. (2013). *Moral perception*. Princeton: Princeton University Press.
- Audi, R. (2015). *Reasons, rights, and values*. Cambridge: Cambridge University Press.
- Bar, M. (2007). The proactive brain: using analogies and associations to generate predictions. *Trends Cognitive Science*, 11(7), 280–289.
- Barbur, J. L., Watson, J. D. G., Frackowiak, R. S., & Zeki, S. (1993). Conscious visual perception without V1. *Brain*, 116(6), 1293–1302.
- Byrne, A. (2009). Experience and content. *The Philosophical Quarterly*, 59(236), 429–451.

- Chappell, R. Y. (2012). Fittingness: The sole normative primitive. *The Philosophical Quarterly*, 62(249), 684–704.
- Cowan, R. (2015). Perceptual Intuitionism. *Philosophy and Phenomenological Research*, 90(1), 164–193.
- Dretske, F. I. (1995). Meaningful perception. In M. K. Stephen & N. Daniel (Eds.), *An invitation to cognitive science: Visual cognition*. Cambridge: The MIT Press.
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society: Biological Sciences*, 364(1521), 1211–1221.
- Korka, B., Schröger, E., & Widmann, A. (2019). Action Intention-based and Stimulus Regularity-based Predictions: Same or Different? *Journal of Cognitive Neuroscience*, 31(12), 1917–1932.
- Mandelbaum, M. (1955). Societal facts. *The British Journal of Sociology*, 6, 305–317.
- McHugh, C., & Way, J. (2016). Fittingness First. *Ethics*, 126(3), 575–606.
- Mole, C. (2015). In J. Zembeikis & T. Raftopoulos (Eds.), *The cognitive penetrability of perception*. Oxford: Oxford University Press.
- Moreau, P., Jolicoeur, P., & Peretz, I. (2013). Pitch discrimination without awareness in congenital amusia: evidence from event-related potentials. *Brain and cognition*, 81(3), 337–344.
- Näätänen, R., & Gaillard, A. W. K. (1983). The orienting reflex and the N2 deflection of the event-related potential (ERP). *Advances in psychology*, 10, 119–141.
- Näätänen, R., Jacobsen, T., & Winkler, I. (2005). Memory-based or afferent processes in mismatch negativity (MMN): A review of the evidence. *Psychophysiology*, 42(1), 25–32.
- Nanay, B. (2011). Do we see apples as edible? *Pacific Philosophical Quarterly*, 92(3), 305–322.
- Paavilainen, P. (2013). The mismatch-negativity (MMN) component of the auditory event-related potential to violations of abstract regularities: a review. *International Journal of Psychophysiology*, 88(2), 109–123.
- Pearce, J. M. S. (2005). Selected observations on amusia. *European Neurology*, 54(3), 145–148.
- Pieszk, M., Schröger, E., & Widmann, A. (2014). Separate and concurrent symbolic predictions of sound features are processed differently. *Frontiers in Psychology*, 5, 1295.
- Pieszk, M., Widmann, A., Gruber, T., & Schröger, E. (2013). The human brain maintains contradictory and redundant auditory sensory predictions. *PLoS ONE*, 8(1), e53634.
- Siegel, S. (2006). Which properties are represented in perception? In T. Gendler & J. Hawthorne (Eds.), *Perceptual experience*. Oxford: Oxford University Press.
- Siegel, S. (2007). How can we discover the contents of experience? *The Southern Journal of Philosophy*, 45(S1), 127–142.
- Siegel, S. (2009). The visual experience of causation. *The Philosophical Quarterly*, 59(236), 519–540.
- Siegel, S. (2011). *The contents of visual experience*. Oxford: Oxford University Press.
- Tye, M. (1995). *Ten problems of consciousness*. Cambridge: MIT Press.
- Van Gulick, R. (1994). Deficit studies and the function of phenomenal consciousness. In G. Graham & G. L. Stephens (Eds.), *Philosophical psychopathology*. Cambridge: MIT Press.
- Werner, P. (2016). Moral Perception and the Contents of Experience. *Journal of Moral Philosophy*, 13(3), 294–317.
- Widmann, A., Gruber, T., Kujala, T., Tervaniemi, M., & Schröger, E. (2007). Binding symbols and sounds: Evidence from event-related oscillatory gamma-band activity. *Cerebral Cortex*, 17(11), 2696–2702.