

References

- Alston, W. 1989. *Epistemic Justification*. Ithaca, NY: Cornell University Press.
- Annis, D.B. 1978. A contextualist theory of epistemic justification. *American Philosophical Quarterly* 15: 213–9.
- Barlow, H. 1974. Inductive inference, coding, perception, and language. *Perception* 3: 123–34.
- Chisholm, R. 1989. *Theory of Knowledge*, 3rd edn. Englewood Cliffs, NJ: Prentice-Hall.
- Cohen, S. 1988. How to be a fallibilist. *Philosophical Perspectives* 2: 91–123.
- Cohen, S. 1999. Contextualism, scepticism, and the structure of reasons. *Philosophical Perspectives* 33: 57–89.
- DeRose, K. 1992. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research* 52: 913–29.
- Dretske, F. 1981. The pragmatic dimension of knowledge. *Philosophical Studies* 40: 363–78.
- Goldman, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy* 73: 771–91.
- Gregory, R. 1997. *Eye and Brain: The Psychology of Seeing*, 5th edn. Princeton, NJ: Princeton University Press.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. New York and Oxford: Oxford University Press.
- Nagel, J. 2008. Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy* 86: 279–94.
- Pryor, J. 2000. The skeptic and the dogmatist. *Noûs* 34: 517–49.
- Rock, I. 1983. *Logic of Perception*. Cambridge: MIT Press.
- Siegel, S. 2017. *The Rationality of Perception*. New York and Oxford: Oxford University Press.
- Stanley, J. 2005. *Knowledge and Practical Interests*. New York and Oxford: Oxford University Press.
- Stanley, J. 2007. Replies to Gilbert Harman, Ram Neta, and Stephen Schiffer. *Philosophy and Phenomenological Research* 75: 196–210.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Replies to Begby, Ghijsen and Samoilova

SUSANNA SIEGEL

I'm grateful to Endre Begby, Harmen Ghijsen, and Katia Samoilova for engaging with *The Rationality of Perception* and for writing such interesting and productive commentaries. Taken together, the three commentaries cover a diverse range of topics.

Begby focuses on my claims in Chapter 10 about testimony, and challenges my epistemological analysis of the relationship between individuals' minds and their cultural milieu. Ghijsen concentrates on my theories of epistemic charge and inference, and takes issue with my account of how inference can modulate epistemic charge. Samoilova connects my theory of epistemic charge with a central topic in the theory of knowledge: whether contextual factors can affect how much justification is needed for knowledge. I reply to them in alphabetical order.

1. Reply to Begby

1.1 Blameworthy vs. ill-founded

Whit forms his racialized beliefs about black people passively, absorbing his outlook from his cultural milieu. I argue that the process by which Whit's beliefs are formed and maintained redound badly on him epistemically. He is less rational than he could be, if he didn't have these beliefs.

Begby assumes that if an individual's belief redounds badly on her, then she is culpable and blameworthy for having that belief. In characterizing my view, he takes this assumption for granted (I have added italics to the relevant words):

[T]he ill-foundedness doesn't simply detract from the justification, it renders the resultant belief irrational and *epistemically blameworthy*.

It may be that such beliefs aren't just lacking in some epistemically relevant dimension, e.g. justification. Rather, they are actually irrational. And to the extent that they are irrational, they reflect poorly on the epistemic subject: the subject is *epistemically culpable* for acquiring and maintaining the belief in question.

[T]o the question of whether he is being *irrational, and therefore blameworthy*, in believing as he does, I stand firm: so far as Siegel's example gives us any reason to believe, he might be doing exactly what he should be doing (epistemically speaking) in forming the belief that he does.

I reject Begby's assumption. Having a belief could redound poorly on your rational standing, even if you are not culpable or blameworthy for having that belief. If you solve a logic puzzle, your rational standing is greater than it would be if you had all the information you needed to solve the puzzle but reasoned poorly with that information. In the scenario where you reason poorly, you might be doing the best you can with the resources you have. But there is no point in blaming you, and it seems wrong to say that you're blameworthy. If we analyse poor rational standing in terms of blame, we won't be able to straightforwardly characterize your reasoning as poor.

In the kinds of cases I discuss in *The Rationality of Perception*, the concepts of blame and culpability have no clear application. Blame implies that

the person doing the blaming can rightfully ask the blamed person to do something differently, such as change their beliefs, or apologize or compensate anyone they may have harmed. Culpability is the flip side of blame. These normative notions are made for appraising and regulating social relationships. In this respect, they belong with other epistemic normative notions designed to help shape social relationships, such as the notions that determine when it is reasonable to trust another person or social entity or the notions of immanent critique and external critique that figure in the theory of social criticism.

In contrast, the purpose of the normative notions I work with is to describe the normative facts of the situation, rather than to devise ways to improve them. Throughout *The Rationality of Perception*, I assume that there are facts about what rational standing different possible subjects would have, given their overall outlook and how it was formed. The main point of my analysis of Whit's case is that among the many marks left on an individual's mind by their social context, we find an impact on the status of their beliefs as well-founded or ill-founded. Just as we can study the effects of variously sized social contexts on cognition, I argue, we can also study the effects of those contexts on a person's rational standing.

Begby rejects my conclusion on the grounds that Whit 'may have done everything right' and is 'performing ideally well, epistemically speaking, with limited resources'. As per Begby's title, Whit is thinking straight in a warped environment.

Once one disconnects redounding poorly on one's rational standing from being culpable, one could in principle agree with Begby that Whit blamelessly forms and maintains his belief and is not culpable, while still holding that Whit's belief redounds poorly on him. Begby says his 'present concern is not about justification, but about the question of epistemic culpability'. Someone who shared Begby's concern might then agree with my main point that Whit's belief redounds poorly on him.

But Begby's comments suggest that he disagrees with my main point, because he thinks the epistemic badness in Whit's case is located wholly outside Whit's mind. According to Begby's assumption linking poor rational standing to blameworthiness, since Whit is not culpable for his racialized belief, that belief does not redound poorly on him either. I disagree.

Whit's friends and associates have the same racialized beliefs that Whit has. From their point of view, the ease with which Whit maintains his outlook, the utter lack of dissonance it causes in him and his unreflective comfort in his outlook all belong to the trappings of reassurance by which their beliefs are maintained. The workings of Whit's mind are part of the workings of the cultural milieu, and they impact the minds of the people around him. Whit's mind is not simply a landing pad on which social forces leave their mark. Similarly, the racialized presumption that Whit absorbs is entrenched in his milieu in part because other people's minds already operate the way his mind

operates. In these ways, individual minds are not epistemically separable from the bad-making features of Whit's social environment.

Since Whit and his friends are in exactly the same epistemic situation when it comes to racialized beliefs, Begby's position entails that they are all thinking straight. The 'warped environment' then has to be analysed independently of the operation of their minds. This consequence robs us of the resources we need to analyse just what it is that's warped about the environment. Part of what's warped is that the outlook shared by Whit and his friends feels normal to them. It's part of their social habitus.

To study the epistemic impact of social contexts on individuals' minds, we need normative notions that can link the individual's mind with their social milieu. The normative notions of ill-foundedness and well-foundedness can draw this link, even though these notions have traditionally been applied analytic epistemology only to the minds of individuals. My analysis of Whit's case takes an epistemic notion designed for analysing belief-forming processes within an individual and makes the case that it has a wider application to the interface between individuals' minds and their social context.

How does the notion of ill-foundedness extend to the interface between the individual's mind and her social milieu? When an individual's belief is ill-founded, it is formed or maintained epistemically badly. Ill-foundedness of a belief is distinct from failing to be supported by evidence – another dimension of justification. A person could have strong evidence for p , even when their belief that p is ill-founded. This could happen if the means by which the belief is formed or maintained does not take account of the evidence. For instance, you might have strong evidence that I'm angry at you, and then your belief that I'm angry at you would be evidentially supported. But for all that, your belief could be ill-founded, because the reason you have it (and keep on having it) is that it's Wednesday, and every Wednesday you think that everyone is mad at you. Here, the factor that explains why you have the belief is independent of the evidence you have for it. The fact that ill-foundedness of a subject's belief is distinct from evidential support is one reason why we need a more fine-grained set of notions than the term 'justification' allows.

Similarly, Whit's racialized belief could be ill-founded, regardless of whether the evidence he has for it is good (because nothing in his social milieu pulls against it and some things seem to favour it), or bad (because it is spotty or misleading or impacted by moral factors) or some of each.

The point of the notion of ill-foundedness is to bring into focus the normative significance of the ways we have of forming and maintaining beliefs. At the level of cultural milieu, we can find a potential disconnect between evidence that counts in favour of a proposition and the factors that explain why that proposition is entrenched in a culture. It could have been part of a cultural myth that everything material has a weight, even before Euler and Newton made the discoveries that gave rise to our current concept of weight (on which nothing material is too small to be measured for weight). In that

scenario, the culturally entrenched presumption could start out ill-founded and end up well-founded.

By locating the epistemic badness entirely outside Whit, Begby's verdict has no resources to describe the epistemic impact on Whit of his social context. It's the difficulty of distinguishing social from individual factors that makes the notion of blame and culpability ill-suited as tools of normative analysis. In contrast, the notion of ill-foundedness does better at connecting an individual to her social context. What's needed to complete the picture is a better understanding of how social practices enable the absorption of ill-founded presumptions.¹

1.2 *Assertion in the USA*

In making the case that ill-foundedness can be transmitted from the mind of the world to an individual's mind, I compare the epistemology of testimony between individuals, on the one hand, with the epistemology of testimony between an individual and their social milieu, on the other. If we ask what social mechanisms enable individuals to absorb presumptions in the seamless way that Whit does, we'd have to look to social practices as well as to discursive 'messages' in advertising, narratives and other discursively structured cultural products (a point emphasized by Haslanger (forthcoming)).

But for the purpose of asking whether ill-foundedness is transmitted from a social milieu to an individual, I find it useful to compare the epistemic features of the transmission with the epistemology of testimony, so that we can locate similarities and differences.

I grant for the sake of argument that ill-foundedness is not transmitted in inter-personal testimony. If your mother tells you the water is unsafe to drink and you believe her, but her belief that the water is dangerous is ill-founded, then according to the view I grant for the sake of argument, you could end up with a well-founded belief. My point in granting this position is to show that the considerations that favour it do not extend to the relationship between an individual and her cultural milieu. Those considerations include the facts that your mother is concerned for you and that she is making an assertion with the purpose of conveying information for your benefit. In contrast, neither of these factors applies to the interface between the mind of the world and an individual.

Begby worries that this disanalogy would apply only to a society in which no one ever explicitly asserted the thing about black men that Whit comes to believe, but the outlook was conveyed in other ways. If it were common practice to assert the thing that Whit believes, then that would be a case of interpersonal testimony. And I granted for the sake of argument that Whit would be justified in believing what he's told, for instance, by his mother or aunts or uncles. On that variant of Whit's case, Begby suggests, by my own

1 I thank Sally Haslanger for helping me see this last point.

lights, Whit's racialized beliefs should come out well-founded. He concludes that my analysis applies only to a specific cultural configuration of racism, one found in a stretch of American history in which anti-black racist assertions are rare, but anti-black racist practices are rife.

In reply, even if Whit's beliefs were maintained in part by explicit assertions of the sort Begby envisions, they would still be maintained in larger part by social practices and other aspects of the cultural milieu that make those assertions seem so natural to Whit. We thus have a case of one factor getting washed out by a large set of other factors. In comparison, if a creature had 19 eyes and one of them didn't see well, its vision could still be overall okay, as the distortions brought on by the bad eye would wash out among the information taken in by the other 18 good eyes. Similarly, even if one factor, assertion, is a good-making feature of the route by which Whit's belief is formed and maintained, there could be a host of other factors distinct from assertion that contribute to the maintaining of his belief, such as the naturalness of his friends' racialized beliefs. And in my example, as I initially described it, there would still be such a host of factors, even if Whit sometimes heard people assert the content of his racialized belief.

It is not especially realistic to suppose that no one ever asserts anti-black racist sentiments, in any stretch of American history. Since assertions about social groups only seem reasonable against a background social setting that provides trappings of reassurance, it's plausible that those trappings carry the most weight in maintaining beliefs like Whit's. If Whit had different friends, different books or a different pattern of curiosity, his racialized beliefs would sit less easily with him. But part of the scenario is that Whit lacks precisely things (see Ch. 10, p. 82).²

2. *Reply to Ghijzen*

2.1 *Grounds for epistemic charge*

What features of perceptual experiences might explain what makes them rationally assessable? I consider two options. The Inferential Ground hypothesis says that perceptual experiences are rationally assessable when and because they are conclusions of inferences. The Phenomenal Ground hypothesis says that they are rationally assessable when and because they belong to a subject's overall outlook on the way the world is.

Ghijzen is suspicious of both answers. In response to the Phenomenal Ground hypothesis, he asks 'why would phenomenal character have any role to play in grounding the rational status of perceptual experience?'. He considers my answer that such experiences belong to the subject's overall outlook, because they have a presentational phenomenal character and that since one's overall outlook redounds well or badly on a subject, so do

2 Thanks to Endre Begby for discussion.

the parts of that outlook. Ghijsen then objects that unconscious reactions could also form part of one's outlook and concludes that presentational phenomenal character is irrelevant to grounding rational status.

I agree with Ghijsen that unconscious beliefs contribute to one's outlook. Why would the fact that both unconscious beliefs and perceptual experiences contribute to one's outlook show that the presentational phenomenal character is irrelevant?

Ghijsen suggests that on my view, both unconscious beliefs and perceptual experiences contribute to one's overall outlook in virtue of their representational content. If having this kind of content is sufficient for grounding epistemic charge, he reasons, then there is no obvious role left for phenomenal character to play. Ghijsen construes phenomenal character as something 'over and above' representational content, and, therefore, a potential ground for epistemic charge distinct from the feature that plays this role in unconscious belief. But presentational phenomenal character is too closely tied to representational content to provide a distinct feature in the way Ghijsen envisions. When your perceptual experience presents you with the property roundness, for example, it attributes the property to something, and your experience is correct only if something is round.³

Ghijsen's complaint about the Inferential Ground hypothesis is that it won't classify demon-induced perceptual experiences as rationally assessable, so long as those experiences are caused by an external manipulator rather than by any inferential process internal to the subject. Ghijsen thinks (and thinks it is plain) that if any perceptual experiences are irrational, demon-induced experiences are. He writes:

[O]nce one allows that perceptual experiences themselves already rebound on our rationality, then these types of demon-induced experiences should look exactly like the kind of experiences that would rebound badly on our rationality.

Why would demon-induced experiences rebound badly on a subject, when the subject's own mental capacities played no role in bringing about those experiences?

The conceit of *The Rationality of Perception* is that processes that occur within a subject's own mind and are of her own mental doing are paradigms of rationally appraisable processes. If such processes culminate in perceptual experiences, then those perceptual experiences are appraisable as well. In contrast, demon-induced experiences bypass the reasoning capacities of a subject entirely.

Ghijsen's criticisms of the two hypotheses are related. Both criticisms draw on and extend his earlier point (Ghijsen 2016) that the phenomenal character

3 For further discussion of the relationship between presentational phenomenal character and representational content, see Siegel 2010, Ch. 3.

of experiences does not bestow them with the epistemic power to be unjustified justifiers. If phenomenal character can't bestow any power to justify without needing justification, he reasons, it can't ground rational standing either. And since proponents of the Rationality of Perception thesis agree that presentational phenomenal character isn't enough to make perceptual experiences into unjustified justifiers, he thinks they should go farther and hold that such phenomenal character can't provide justification in the new evil demon scenario, where the experiences are induced by a demon rather than resulting from the subject's own inferences or other cognitive machinations.

We thus have two points of major disagreement, both central to the disputes between internalist and externalist approaches to justification. First, Ghijsen thinks presentational phenomenal character plays no epistemic role in perceptual justification, whereas I think it does. Second, Ghijsen thinks that degrees of reliability correlate with degrees of justification and ground rational standing, whereas I think they don't.

On the second debate, it is hard to see how mere degrees of reliability could ground the rational standing of anything. The height of a seedling reliably indicates how long it has been in the earth, but the plant does not thereby have a rational standing, and neither does its height. When reliabilists claim that degrees of justification are measured by degrees of reliability of a belief, they are often taking as understood that beliefs redound well or badly on a subject and aiming to explain how justified or unjustified those beliefs are.

On the first debate, Ghijsen is right to ask what gives presentational phenomenal character such extraordinary power to bestow justification. It might seem that if anything has the power to bestow justification, it would be the fact that normally, perceptual experiences facilitate knowledge of the environment (Peacocke 1999), rather than phenomenal character itself, when construed as a property of experiences that supervenes on the internal state of the subject's mind, rather than on that plus the mind's relation to things in the environment.

My reason for favouring the Phenomenal Ground hypothesis is that it accommodates the seemingly simple fact that seeing the mustard in your fridge can give you reason to think that it is there, while also accommodating the potential adverse effects of wishful or fearful seeing. My reason for rejecting phenomenal conservatism is that it gives the wrong results in key cases of cognitive penetration. In any of the three versions recapped by Ghijsen, the Phenomenal Ground hypothesis preserves the idea that perceptual experience easily carries the weight we naively think it does in providing justification, while respecting the ways that our own mental processes can detract from its role.

2.2 Modulating epistemic power

A second strand in Ghijsen's critique focuses on whether inference can modulate the epistemic powers of experience. He agrees that the aetiology of

experiences can modulate epistemic power, but disagrees that the notion of inference helps explain how this modulation occurs.

His initial criticism of my account of inference is that it fails to predict that experiences he considers ‘intuitively [epistemically] good’ end up with that status.

In his example, drawn from Lyons’s (2011) discussion, an ill-founded but true belief about the prevalence of snakes puts a subject on the lookout for them, leading her to notice more snakes than she would without any such belief. Ghijsen observes that since the prior belief that there are snakes around helps cause an experience, if one then strengthens one’s original belief on the basis of snake experiences, the transition will have the same components as a circular inference. My account of inference, Ghijsen says, lacks the resources to explain why these snake experiences are not inferential responses to the prior belief that there are snakes around, and, therefore, why the beliefs formed on the basis of those experiences are not as poorly justified as circular inferences.

When I discuss this case in Chapter 6, I take it as a datum that directing spatial attention is not part of the role of premise states in inference. Directing spatial attention is a way of starting off a perceptual process, and one’s pattern of attention need not affect how that process unfolds.⁴ In Lyons’s example, the fact that your attention has been directed downward does not affect what you see when you look at the ground. If you had looked at the same place on the ground because you were shy or because the light from the sun was too bright, you would see the same snakes as the ones seen by Lyons’s character who is on the lookout for snakes. The fact that this vigilant person’s experiences are shaped by what he is actually seeing leads Ghijsen to say that those experiences are ‘intuitively good’.

Ghijsen worries that I haven’t said enough to justify my assumption that the influence of prior belief on spatial attention falls outside the subject matter of theories of inference. A reductive theory of inference would purport to identify features essential to inference that are missing here. But this method of delineating the subject matter is not open to me, Ghijsen observes, since I claim to illuminate the nature of inference without offering any reductive analysis of it.

My reply is that we infer from information we have already, whereas Lyons’s case involves taking in new information about the location of particular snakes and drawing conclusions exclusively from that new information. Suppose I want to know whether you’ve reached the top of the stairs, so I look up the staircase, see that you are at the top and form the belief that you are now one floor above me. Do we need a theory of inference to tell us whether or not I have inferred that you reached the top of the stairs from my

4 In contrast, feature-based attention can affect how it unfolds, as Carrasco (2011) and others have argued.

desire to know where you are? This case does not seem to be a borderline case between inference and non-inference, and it is exactly analogous to Lyons' snake case. I am drawing the conclusion that you are at the top of the stairs from my visual experience. My desire preceded and helped cause my experience, but it does not enter into my reasoning at all.

Because cases like Lyons's snake example or my staircase example fall clearly outside the subject matter of theories of inference, we don't have to rely on the resources of those theories to determine whether or not those cases belong to the phenomena that these theories aim to illuminate. A sceptic who holds there is no such thing as epistemically appraisable inference would take issue with any purported division of mental processes into inferential ones and non-inferential ones. But their criticisms would apply to any philosophical theory that takes an inference as an epistemically interesting phenomenon, not just to the inferentialist version of the Rationality of Perception thesis.

In a different version of the snake case, the prior belief would both direct attention and combine with perceptual information one gets when one attends, to yield the experience with the content: there's a snake nearby. In contrast to the role of prior beliefs in Lyons's snake case, this kind of role is a paradigm of inference to perception.

A further criticism from Ghijsen concerns this kind of inference. According to Ghijsen, 'sensory stimuli are always going to underdetermine the conclusion' that shapes a percept, and so every case of inference to percepts will count as a case of poor inference, unless it is supplemented with additional assumptions. I describe how Jill's fear could make her conclude that Jack is angry from her experience of Jack's blank stare, where the conclusion takes the form of an anger experience. If the transition is an inference, I explain, then it manifests Jill's sensitivity to the rational relationships as Jill sees them between Jack's blank stare and her background assumptions that link his blank stare to anger.

Ghijsen charges that my analysis of this case 'simply stipulates how to deal with [it] ... rather than arguing that the inferential account best explains what is going on in the case'. Since Jill is a fictional character, there is no independent fact of the matter about what psychological mechanisms produce her experience. Jill's scenario is an occasion for describing what various routes to experience could be like, rather than an explanandum for theories of how experiences of anger can arise.

In describing the various different routes to Jill's experience that could fill out the fiction, Ghijsen contrasts merely causal routes from Jill's fear to her experience with inferential ones and asks how my theory can distinguish between a merely causal, non-inferential route by which she moves from a blank-stare experience of Jack's face to an experience of his face as angry and an inferential route in which the conclusion is not warranted by the inputs. My answer is that distinctively inferential responses are found in both

epistemically faulty inferences and epistemically good ones but not in merely causal transitions. The difference comes from the way in which the subject responds to the blank-stare experience. Because my theory is non-reductive, there is no further analysis of what kind of response this is, though we can characterize such responses indirectly in several ways, as I do in Chapter 5. If Ghijsen thinks any non-reductive theory of inference would be inadequate, then that is where the issue needs to be joined, and his objections to an inferentialist version of the Rationality of Perception thesis will be derivative of that more general concern.

3. *Reply to Samoilova*

Samoilova argues that the framework of epistemic charge can help analyse cases where high stakes, moral considerations, or other contextual factors raise the amount of justification needed for knowledge. Samoilova calls the level of justification needed in these cases the *justificatory highline*. Consider Dretske's (1981) classic pair of cases involving the perception of birds. In Case 1, you see a duck in a pond, with markings distinctive of Gadwall ducks, such as a white patch on wing. As it happens, it really is a Gadwall duck. In some versions of this situation, you could know that it's a Gadwall duck on the basis of your visual experience together with your ability to recognize this kind of duck.

Case 2 is just like Case 1, except you learn that an ornithologist has a hypothesis that Siberian Grebes have started migrating through the area, and those birds are hard to distinguish from Gadwall ducks. The crucial difference between them is a mark on the belly, but from your perspective you can't see the belly of the duck because it is swimming. The ornithologist's hypothesis gives you some reason to think that the bird you are seeing is a Siberian Grebe rather than a Gadwall duck.

Dretske argues that in Case 2, your knowledge of the ornithologist's hypothesis has raised the bar for knowing that it's a Gadwall duck. If you checked its belly, you could know that it's a Gadwall duck, but as things stand, you don't know, even though if you hadn't learned about the ornithologist's hypothesis, you would know.

If Dretske is right, one explanation for why you don't know in Case 2 is that a single set of resources gives you less justification in Case 2 than it gives you in Case 1. These resources are your visual experience and your beliefs about what Gadwalls look like. If in Case 2 you added to those resources a further visual experience in which you saw the duck's belly (and saw that it looks as a Gadwall's belly would look, rather than looking the way a Grebe's belly would look), you could meet the higher bar for knowledge.

But when we focus on the visual experiences you actually have in Case 1 and Case 2, for all we've said, the epistemic potency of those experiences is the same before you look at the duck's belly and afterward. What would help

you reach the highline is a further experience that tells you what's on its belly – not any augmented epistemic power of experience or any additional epistemic charge that would go with it.

Dretske's example does not illustrate Samoilova's main claim, but in other types of cases, the notion of epistemic charge could help analyse what enables someone to reach a justificatory highline. We need a pair of cases that vary in how much justification is needed for knowledge, and then we need to hold constant a phenomenal type of experience across a pair of cases, as well as any relevant background beliefs. We can then consider whether the extra justification called for in the high-stakes case could be provided by the experience, due to the route by which the experience was formed in the subject's mind. Samoilova claims that there are such cases, and her central observation is that if experience could provide the extra justification that's needed, then the change in justification provided by the same experience across cases could be expressed as differences in epistemic charge.

Here are two examples.

First, imagine a judge whose job is to watch a series of divers dive into a swimming pool and assess how well-executed their dives are. The judge watches the divers carefully from the time they position themselves at the edge of the diving board until they hit the water, noting their posture at each stage and their overall trajectory as they flip and turn and twist. For each dive, there's a paradigmatic way to execute it, and the judge assesses how closely each diver's trajectory follows the contours it is meant to follow.

We can imagine that part of the judge's expertise consists in knowing how the diver's body should be positioned at each moment and how various positions and micro-movements interact. For instance, a twist at time t can be well-executed only if the diver was aligned properly at the earlier time t^* , so registering the proper alignment at t^* helps the judge assess how well executed the twist is at time t . Like many tasks, a complex dive will involve a host of such interactions. And they provide the judge with ample opportunity for inferences as she watches the dive unfold – both inferences to beliefs and inferences to experiences. For instance, once she sees that the diver has left the board with enough power to reach an optimal height, but not so much as to exceed it, she can infer that the twist will finish at the optimal time so that she can hit the water with an untwisted body. And this basis for inference is in principle available to her experience of the dive at the next moment.⁵

- 5 I am told by a gymnastics expert that in gymnastics competitions, judges can sometimes be heard to gasp a few moments before other spectators see that the gymnast has made a bad error. This phenomenon is evidence that the judges are perceiving cues to the mishap that non-experts cannot see. In this video, a gymnastic coach sees the start of a mishap early on and saves a gymnast from error (and quite possibly injury). To a non-expert the beginning of the flip looks perfectly fine. <https://www.youtube.com/watch?v=mjb6cTD59Ic>. Thanks to gymnastics judge and philosopher Lauren Davidson for discussion and for the example.

We can compare the judge's experience of a diver who is about to hit the water with a novice's experience of the same diver. The judge and the novice might both experience the diver's body as configured in exactly the same way, and their perspective might be such that their experiences are the same. But the judge's experience could have more epistemic power than the novice's, if the way in which it is caused by the diver's actual configuration is informed by the judge's knowledge of how various micro-features of the diver's movement interact. In a high-stakes situation, the epistemic power of the judge's upgraded experience could meet the Samoilovan highline, while the novice's experience would not. As Samoilova points out, this difference could be reflected in the amount of epistemic charge belonging to each experience. It's easy to picture a high-stakes situation here. Whether a diver wins a competition could depend on the judge's verdict, and the verdict could in turn depend on how the dive looks to the judge when she sees it.

In suggesting that epistemic charge can help analyse different epistemic contributions of experience in cases where the standard for knowing is higher than usual, Samoilova emphasizes that inference can modulate epistemic charge of experience. But to illustrate the phenomenon she brings into focus, in which experiences can be potent enough to reach a justificatory highline, any modulator of epistemic charge would work. This brings me to a second example of two experiences that are phenomenally the same but differ in epistemic charge, with the result that one of them could meet a justificatory highline that the other would not.⁶

Consider a chess match, in which a chess expert is playing one game and a novice is playing a separate game. By chance, at one point in their respective games, expert and novice face the same configuration on the board, and in response, they both make the same move. The novice just happened to notice the move, whereas the chess expert, like chess experts generally, perceives the board as chunked into patterns poised for attack or defence (Chase and Simon, 1973). The move made by both expert and novice was evident from focusing on a 'chunk' of the board that highlighted the pieces poised for attack. It's this spatial chunk that the novice happened to notice.

Here, it's the same visual experience that gives expert and novice grounds for thinking that the move that they both end up making is a good move. What is the content of their shared experience of the board? One option is that their experiences present the affordance of moving a piece. A different option is that the experiences present only 'thin' content, such as the spatial configuration of pieces.⁷ Either way, the expert's experience has that content thanks to a long-term process of perceptual learning in which she learns to chunk configurations of pieces into a single unit (this process is called 'unitization' in

6 Thanks to Zoe Jenkin for discussion of both this case and the results in Chase and Simon 1973.

7 On thin content, see the debate between Siegel and Byrne (2017).

the theory of perceptual learning), whereas the novice simply hits on the unit by chance.⁸

If perceptual learning modulates the amount of epistemic charge the expert's experience has, compared with the novice's experience of the same structure, then it's a further question whether it is modulated by inference. It could instead be modulated by recognitional skill. If so, that modulation would illustrate a non-inferentialist version of the Rationality of Perception.

*Harvard University
Cambridge MA 02138, USA*

References

- Carrasco, M. 2011. Visual attention: the past 25 years. *Vision Research* 51: 1484–525.
- Chase, W. and H. Simon. 1973. Perception in chess. *Cognitive Psychology* 4: 55–81.
- Connolly, K. 2017. Perceptual learning. In *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta, Summer 2017 edn. <https://plato.stanford.edu/archives/sum2017/entries/perceptual-learning/>.
- Dretske, F. 1981. The pragmatic dimension of knowledge. *Philosophical Studies* 40: 363–78.
- Ghijsen, H. 2016. The real epistemic problem of cognitive penetration. *Philosophical Studies* 173: 1457–75.
- Haslanger, S. 2018. Cognition as a Social Skill. *Australasian Philosophical Review* 2. Forthcoming.
- Lyons, J. 2011. Circularity, reliability, and the cognitive penetrability of perception. *Philosophical Issues* 21: 289–311.
- Peacocke, C. 1999. *Being Known*. Oxford: Oxford University Press.
- Siegel, S. 2010. *The Contents of Visual Experience*. New York: Oxford University Press.
- Siegel, S. and A. Byrne. 2017. Rich or thin? In *Current Controversies in Philosophy of Perception*, ed. B. Nanay, 57–80. New York: Routledge.

8 On unitization and other mechanisms of perceptual learning, see Connolly 2017.