

Symposium on *The Rationality of Perception**

Introduction

Susanna Siegel

The Rationality of Perception concerns two epistemic questions. The first question concerns the interface between an individual's mind and their cultural milieu. Can the epistemic status of a culturally entrenched presumption be transmitted to an individual who unreflectively absorbs it?

The second question concerns the relationship within an individual's mind between their perceptual experiences, and their beliefs, desires, suspicions, fears, and other outlooks that influence them. Can the epistemic status of an individual's prior outlook be transmitted to a perceptual experience, when the outlook shapes the experience?

The Rationality of Perception makes the case that both kinds of epistemic transmission are possible. I start with normative concepts from analytic epistemology: well-foundedness and ill-foundedness. Traditionally, these concepts applied to beliefs held by an individual, but not to perceptual experiences of an individual, nor to culturally entrenched presumptions. I argue that the concept can be extended in both of these directions. These normative concepts can apply more widely than has traditionally supposed.

The bulk of *The Rationality of Perception* is devoted to the question about the intra-personal interface between perceptual experience of an individual and her prior outlook. The main thesis is:

Rationality of Perception (RP): Both perceptual experiences and the processes that give rise to them can be rational or irrational.

For example, suppose Jill fears that Jack is angry and suspects he is. When she sees him, he looks angry to her, due to her fearful suspicion. There are versions of this case, I argue, in which strengthening her fear in response to experience would be unreasonable.

Another example: Vivek is an excessively vain performer. To him, the faces in the audience range in their expression from neutral to pleased. Remarkably, no one ever looks disapproving. Vivek's vanity influences the character of his perception. Depending on his self-conception, how a scene looks to him will differ, even when all other conditions stayed the same.

* Many thanks to symposiasts Endre Begby, Harmen Ghijsen, and Katia Samoiloa for engaging so robustly with *The Rationality of Perception* and for writing such interesting and thoughtful responses to it.

What should Jill and Vivek believe, in response to their visual experiences? Could you blame Jill for strengthening her belief that Jack is angry, once she sees him? Not really. Yet since the problem originates in her own mind, what else besides her mind's own working is there to blame? In this context, the concept of blame has no clear application. To address what Jill or Vivek should believe, we need a different normative concept. And since some effects on perceptual experience by prior outlooks are epistemically good or harmless, whereas others are epistemically bad, what's needed is a framework for analyzing the epistemically relevant features of the routes to perceptual experience that shows how to draw this distinction.

A framework that can do this is inference. When we draw inferences from our beliefs to other beliefs, we reason from information we have already. When our inferences are epistemically good they redound well on us as rational subjects, and when they are epistemically bad, they redound poorly on us on rational subjects. There are two standard ways for an inference to be bad: the inputs to the inference can be unjustified, or the transition to the conclusion can be improperly constrained.

The inferentialist version of the Rationality of Perception thesis is that perceptual experiences can be the conclusions of epistemically evaluable inferences. Just as the conclusions of inferences from beliefs to other beliefs redound well or poorly on us as rational subjects, according to Inferentialist RP, the same is true of perceptual experiences that are conclusions of inference. Inferentialist RP distinguishes epistemically better and worse routes to perceptual experience by better and worse inferences. In this way, it finds a more extensive scope for epistemically evaluable inference than is traditionally assumed.

Epistemic charge

If the main theses of *The Rationality of Perception* are correct, then perceptual experiences are among the mental phenomena that redound well or badly on the subject's rational standing, and that their epistemic status can be modulated by inferences. My label for the property that measures the epistemic power and rational standing of an experience is *epistemic charge*. Just as electrical charge can be conducted from one material to another, epistemic charge can be transmitted between experiences and other mental states. An example of an epistemically charged mental state is a belief that is either justified or unjustified. As justification is usually construed, beliefs themselves are the bearers of this epistemic normative status.

The Rationality of Perception thesis says that perceptual experiences can, in themselves, manifest an epistemic status. They are not merely an enabling condition for other mental states to manifest such a status. Nor are they merely contributors to determining the epistemic status of beliefs as either well or poorly justified, though they play this role as well. In addition to this role, they manifest an epistemic standing in themselves. Epistemic charge is a way to measure the epistemic standing of an experience. And just as beliefs that are very well justified can transmit their justification to subsequent beliefs formed on their basis, similarly

experiences that are highly epistemically charged can transmit their charge to subsequent beliefs formed on their basis.

According to the Inferentialist version of RP, epistemic charge can be modulated by inference. Poor inferences can reduce the epistemic charge an experience has below a baseline, relative to which beliefs based on the experiences give you good reason to believe that things are the way the experience presents them as being. For instance, if you want to know whether there is any mustard in the fridge, you can open the door to look. You can discover that a fear of running out of mustard is unfounded, by opening the fridge and seeing a jar of mustard there. Here, your experience has a baseline amount of epistemic charge. By contrast, in cases of perceptual hijacking, perceptual experiences may end up with less than a baseline amount of epistemic charge, and therefore with less epistemic power to provide justification for subsequent beliefs formed on their basis.

Replies to Begby, Ghijsen, and Samoilo Susanna Siegel

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 having a 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:

“the ill-foundedness doesn't simply detract from the justification, it renders the resultant belief irrationally and **epistemically blameworthy**.” [page number to be inserted]

“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.” [page number to be inserted]

“to 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.” [page number to be inserted]

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 analyze 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 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.

By 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 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 "his main 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 suggests 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 working 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 comes to operate. In these ways, individual minds are not epistemically separable from the bad-making features of his 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 analyzed independently of the operation of

their minds. This consequence erases robs us of the resources we need to analyze 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 employed in analytic epistemology only within the minds of individuals. My analysis of Whit's case takes an epistemic notion designed for analyzing 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 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 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 favor 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 favor 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 by weight). In that scenario, the culturally entrenched presumption could start out ill-founded, and end up well-founded.

By locating the epistemic badness 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 that makes the notion of blame and culpability

ill-suited as tools of normative analysis. By contrast, the notion of ill-foundedness does better at connecting and 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 with the epistemology of testimony between an individual and their social milieu. 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).

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 interpersonal testimony. If your mother tells you the water is unsafe to drink and you believe her, but her belief in that the water is dangerous is ill-founded, then according to the view I grant for the sake of argument, you could end up a well-founded belief. My point in granting this position is to show that the considerations that favor 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. By contrast, neither of these factors apply 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 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 could still be maintained in part by the social practices and other aspects of the cultural milieu that make those assertions seem so natural

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

to Whit. We thus have a case of one factor getting washed out by a large set of other factors. By comparison, if a creature had nineteen 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 eighteen 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 these things.

2. Reply to Ghijsen

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.

Ghijsen 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 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, unconscious beliefs and perceptual experiences both 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. Ghijzen 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 Ghijzen 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.²

Ghijzen’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. Ghijzen thinks (and thinks it is plain) that if any perceptual experiences are irrational, demon-induced experiences are.

“[O]nce one allows that perceptual experiences themselves already redound on our rationality, then these types of demon-induced experiences should look exactly like the kind of experiences that would redound badly on our rationality.” [page numbers to be filled in later]

Why would demon-induced experiences redound 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. By contrast, demon-induced experiences bypass the reasoning capacities of a subject entirely.

Ghijzen’s criticisms of the two hypotheses are related. Both draw on and extend his earlier point (Ghijzen 2016) that the phenomenal character 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, Ghijzen thinks

² For further discussion of the relationship between presentational phenomenal character and representational content, see Siegel (2010) chapter 3.

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 favoring 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 the 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 etiology 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.³ 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 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 involving 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 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 skeptic 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.

³ Feature-based attention can affect how it unfolds, as Carrasco (2011) and others have argued.

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 the places here likely to be to yield the experience that there's a snake nearby. By 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 stimulation 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 makes 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 goes 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 more general concerns.

3. Reply to Samoilova

Samoilova argues that the framework of epistemic charge can help analyze cases where high stakes, moral considerations, or other contextual factors raise the amount of justification needed for knowledge. Samoilova calls this level of justification needed in these cases the *justificatory highline*. Consider Dretske's classic pair of cases involving the perception of birds (1981). 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 starting 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 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 psychological 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 analyze 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 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.

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 the exact 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 judge's upgraded experience could meet the Samoilova highline, while 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 analyze 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 meets a justificatory highline whereas the other does not.⁴

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 defense (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 spatial configuration of pieces.⁵ Either way, the expert comes to have that content through a process of perceptual learning in which she learns to chunk configurations of pieces into a single unit (this process is called "unitization" in the theory of perceptual learning), whereas the novice simply hits on the unit by chance.⁶

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's 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.

Bibliography

M. Carrasco (2011) Visual attention: The past 25 years. *Vision Research* 51:1484–525.

W. Chase and H. Simon (1973) "Perception in Chess" *Cognitive Psychology* 4: 55-81.

K. Connolly (2017) "Perceptual Learning" *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.). URL = <https://plato.stanford.edu/archives/sum2017/entries/perceptual-learning/>.

⁴ Thanks to Zoe Jenkin for discussion of both this case and the results in Chase and Simon (1973).

⁵ On thin content see the debate between Byrne and Siegel (2017).

⁶ On unitization and other mechanisms of perceptual learning, see Connolly (2017).

F. Dretske (1981) The pragmatic dimension of knowledge. *Philosophical Studies* 40 (3): 363--378.

H. Ghijzen (2016). The real epistemic problem of cognitive penetration. *Philosophical Studies*, 173(6):1457-1475.

J. Lyons (2011) "Circularity, reliability, and the cognitive penetrability of perception". *Philosophical Issues* 21(1):289-311.

C. Peacocke (1999) *Being Known*. Oxford: Oxford University Press.

S. Siegel (2010) *The Contents of Visual Experience*. New York: Oxford University Press.

S. Siegel and A. Byrne (2017) "Rich or Thin?" In B. Nanay, ed. *Current Controversies in Philosophy of Perception*. New York: Routledge. 57-80.