

A guide to Shoemaker's 'Self-knowledge and inner sense'

Jeff Speaks
Philosophy 519

March 5, 2004

1	Self-knowledge as 'object perception'	2
1.1	Perception according to the object perception model	2
1.2	Problems with the neutral conditions	3
1.2.1	The lack of an organ of introspection	3
1.2.2	The lack of sensory 'appearances'	3
1.3	Problems with the object-model conditions	4
1.3.1	Knowledge of the self	4
1.3.2	Knowledge of beliefs, desires, and intentions	5
1.3.3	Knowledge of sensations	6
1.3.4	The role of attention in knowledge of sensations and intentional states	6
1.4	Result of rejecting the object-perception model	7
2	Self-knowledge as 'broad perception'	8
2.1	Perception according to the broad perceptual model	8
2.2	The concept of self-blindness	9
2.3	Self-blindness and bodily sensations	13
2.4	Self-blindness and perceptual experiences	14
2.5	Self-blindness and the will	16
2.6	Self-blindness, beliefs, and desires	17
2.6.1	The argument from self-ascription	17
2.6.2	The argument from acting 'on the basis of' belief	18
2.6.3	The argument from the possibility of deliberation	19
3	Shoemaker's positive picture	19
3.1	The problem	19
3.2	The functionalist view of mental states	19
3.3	Shoemaker's functionalist explanation of self-knowledge	20
4	The problem posed by phenomenal character	21
4.1	Perception and its objects	22
4.1.1	The 'simple view'	22
4.1.2	Literal projectivism	23
4.1.3	Figural projectivism	23
4.1.4	Shoemaker's positive view of phenomena properties and qualia	23

4.1.5	An objection to Shoemaker’s view of phenomenal properties	25
4.1.6	A Shoemaker-ian view of color	26
4.1.7	Shoemaker’s view on the relationship between phenomenal feel and representational content	26
4.2	Phenomenal properties and self-knowledge	27
4.2.1	The case against functionalism about perceptual representation	27
4.2.2	Shoemaker’s defense of functionalism	27

The guiding question of Shoemaker’s three lectures is: can our self-knowledge be understood as the result of a kind of sense perception — a kind of inner sense? This question is important for a number of reasons.

First, the idea that self-knowledge is akin in some ways to perception is, in a way, the most intuitively appealing explanation of our ability to know our own mental states. Doesn’t your awareness that you are seeing, say, a red after-image *seem* similar to perceptual awareness? Also, it’s important to note that our knowledge of our own mental states is knowledge of a class of contingent facts; and the standard explanation of knowledge of contingencies is via appeal to perceptual experience.

Second, if perceptual models of self-knowledge fail — as Shoemaker thinks they do — we’ll need to seek some other explanation of our self-knowledge. And, if Shoemaker is right, the reasons behind the failure of perceptual models of self-knowledge point toward an explanation of self-knowledge which reveals a lot, not only about the nature of self-knowledge, but also about the nature of mental states more generally.

It’s important to note at the beginning an ambiguity in the use of the word ‘introspection.’ According to one meaning, ‘introspection’ just stands for whatever process it is by which we gain knowledge of our own mental states. On this interpretation, the claim that self-knowledge is possible because of our introspective access to our own mental states is a trivial one. This is not, of course, the sense in which Shoemaker is asking whether our self-knowledge can be understood as the result of introspection. Rather, in such questions, ‘introspection’ is given a substantial interpretation, according to which to say that we have self-knowledge on the basis of introspection is to say that we have self-knowledge on the basis of a kind of perception of our own mental states.

As Shoemaker recognizes, this immediately raises the questions: how does the introspection theorist claim that self-knowledge is like ordinary perceptual knowledge? And what does the introspection theorist think that ordinary perceptual knowledge is like, anyway? The fact that these questions are pressing ones is the reason why Shoemaker spends so much time explaining the ‘object perception’ and ‘broad perceptual’ views of perception and introspection. The idea is that we need to make the claim that self-knowledge is akin to introspection precise by spelling out two different views of, or species, of perception; we’ll then be in a position to see whether self-knowledge can be understood according to either.

1 Self-knowledge as ‘object perception’

The first version of the perceptual, or ‘inner sense’ explanation of self-knowledge that Shoemaker discusses is what he calls the ‘object perception’ model. Roughly, the idea is this. Everyone agrees that perception gives us access to/represents/makes us aware of both objects and states of affairs. E.g., my visual experience might now presents to me both an object — a glass — and a state of affairs involving that object — the state of affairs of the glass being on the table to my left. One question in the philosophy of perception is: which of these is more fundamental? Are we presented in the first instance with objects, or with states of affairs? Shoemaker does not address this question directly. Rather, he lays out both views of perception, and asks whether introspection can be understood in terms of either.

First, he considers the view that our inner sense functions according to the ‘object perception’ model.

1.1 Perception according to the object perception model

According to the object perception model, perceptual experiences satisfy eight criteria, which I group below into three categories:

Neutral conditions

- [1] Sense perceptions involve an organ of perception under one’s voluntary control.
- [2] Sense perceptions involve the occurrence of sense experiences or sense impressions distinct from the object perceived.

Object-model conditions

- [3] Sense perception provides awareness of facts only via awareness of objects.
- [4] Sense perceptions provides information about objects which enables the perceiver to re-identify them.
- [5] Sense perceptions provide awareness of extrinsic properties only via awareness of their intrinsic properties.
- [6] One can shift one’s attention between objects perceived without changing what one perceives.

Independence conditions

- [7] Sense perceptions are causally produced by the objects of perception.

[8] Sense perceptions typically provide awareness of objects and states of affairs which exist independently of their being perceived.

These eight conditions are not, as the three categories suggest, on equal footing. The first two ‘neutral conditions’ give features of perceptions without taking much of a stand on what perceptual experience or its objects are like; the ‘object-model conditions’ state features of perceptions specific to a view of perception which gives awareness of objects primacy over awareness of states of affairs; the ‘independence conditions’ again abstract from the object perceptual model to focus on features of the relationship between perceptual experiences and their objects (whatever they happen to be).

1.2 *Problems with the neutral conditions*

Despite the fact that Lecture I is primarily devoted to the object perceptual model — and hence to the question of whether self-knowledge is explained by our possession of an inner sense satisfying the object-model conditions [3]-[6] — he begins by raising a few questions about the neutral conditions, [1]-[2].

1.2.1 *The lack of an organ of introspection*

Obvious problem for the claim that self-knowledge is ‘just’ another kind of perception. The case of proprioception.

Difficult to see how important this objection is, because it is not obvious why someone who holds a basically perceptual model of introspective knowledge should think that feature [1] carries over from the straightforwardly perceptual to the introspective case.

1.2.2 *The lack of sensory ‘appearances’*

Claim [2] about ordinary sense perceptions has it that in cases of sense perception, there is, in addition to the object perceived, an appearance of that object. Why this claim seems plausible; the example of a coin which ‘looks elliptical.’ Why you might want to resist it: questions about what appearances might be.

Shoemaker claims that there is a problem with taking introspection to have feature [2]; as he says, “No one thinks that in being aware of a sensation or sensory experience one has yet another sensation or experience that is “of” the first one, and constitutes its appearing to one in a particular way.” (254-5)

I agree that it seems odd to think of introspection as exhibiting this feature; but what’s the argument against this view? There’s no obvious way to generate a regress here.

An argument suggested by footnote 3: appearances are perspectival; but there can be no such thing as different introspective perspectives on a single mental entity. (This is not obvious . . . can you have different perspectives on a pain?)

1.3 Problems with the object-model conditions

The meat of Lecture I, though, is the criticism of the view that introspection can be something characterized by the object-model conditions, [3]-[6]. Rather than discussing each condition in turn, Shoemaker organizes the discussion by considering in turn different sorts of objects of self-knowledge: (i) the self, (ii) ‘intentional’ mental states like beliefs, desires, and so on, and (iii) sensory experiences.

1.3.1 Knowledge of the self

Hume’s denial that we have any awareness of a self, apart from awareness of specific mental acts and properties. Why, even if you reject Hume’s subsequent denial of the existence of a self apart from such acts and properties, you might take his line of argument to show that our awareness of a self does not fit the object perception model very well.

Shoemaker’s main argument here is that awareness of the self fails to satisfy feature [4], which requires that perception of objects provide information which will enable a perceiver to re-identify them. The line of this argument is not easy to follow, but goes something like this: in the case of perception, perceptual experiences of objects *explain* our ability to re-identify and track those objects over time. So, if self-knowledge works like perception, we should expect our knowledge of ourselves to explain our ability to re-identify, and track, our selves over time. But in fact self-knowledge does not work this way; it presupposes rather than explains our ability to re-identify ourselves.

One problem in understanding this argument is in understanding what it is to *identify* an object. For purposes of this argument, I think that we can take it that to identify an object is more or less to *refer* to it, either in language or in thought. (Referring to an object might be a matter of thinking or expressing an object-dependent thought about it.)

Why does our access to the self not work this way? Shoemaker’s idea is that we do not need to know properties of the self that we can use to re-identify it; there’s no need for such criteria of identification, since the self is ‘immune to error through misidentification.’

Objection: sometimes we do need to identify ourselves by using properties of ourselves, as when we identify ourselves in a mirror. Shoemaker’s reply: even in such cases these identifications rest on a more fundamental ability to identify ourselves, without knowledge or perceptual awareness of distinguishing properties. Marx brothers example, p. 258.

1.3.2 Knowledge of beliefs, desires, and intentions

The main problem with a perceptual model of our awareness of beliefs, desires, and intentions comes not with feature [4] of the object-perceptual model, but with feature [5], which holds that in perception we become aware of the extrinsic properties of objects via awareness of their intrinsic properties.

Given the truth of the kind of externalism which Shoemaker assumes, the content of a belief, intention, or desire is an extrinsic property of it. So how could we have a quasi-perceptual awareness of the contents of our beliefs? Presumably via awareness of some of their intrinsic properties. But, as Shoemaker says, this is not very plausible, for at least two main (related) reasons: (i) there is no obvious candidate for an intrinsic property of a belief or desire to which we have introspective access; and (ii) we typically identify beliefs by their contents; if the object perception model and [5] were correct, we should expect it to be possible to identify a belief and misidentify it as a belief with a content other than its actual content. But this does not seem possible.

Some notes on this argument:

1. It is important to be clear about the kind of externalism Shoemaker is assuming. The important point in the argument is not that, e.g., beliefs are extrinsic properties of agents, though that is true as well. We are discussing the object perception model; if this model is correct, then we should expect our awareness of our beliefs, intentions, etc. to be awareness directly of these mental objects, rather than awareness of beliefs as properties of agents. (The latter would be more in keeping with the broad perceptual model, which takes the primary objects of perception to be states of affairs — like *me believing that so-and-so* — rather than objects.) So the important point is that, according to this version of externalism, the contents of beliefs are extrinsic properties *of those beliefs*. This comes from Shoemaker's own version of externalism + functionalism, which holds that beliefs are internal states whose contents are fixed (at least partly) by relations between those states and facts external to the believer.
2. This is grist for the mill of the narrow content theorist, who will hold that the narrow contents of beliefs are both intrinsic properties of beliefs and what we have introspective access to.
3. Is it really impossible to misidentify the content of a belief or an intention? One can certainly mistakenly think that one believes *p*, and then later realize that one really believed all long a closely related proposition *q*. There is some temptation to say that this is not a case in which one has correctly identified a belief and mischaracterized its content, but rather one in which one mistakenly thought that one had a given belief that one does not, and then later correctly identified the real belief. (Compare: a case in which one has correctly identified a person but falsely attributed a property to them, and a case in which one misidentifies a person: one thinks that they are someone that they are not.) But it is not obvious that we *must* describe things this way.

1.3.3 *Knowledge of sensations*

Shoemaker rightly points out that the object-perception model is strongest here; things like after-images, patches of color, and pains do seem more ‘object-like’ than other mental entities. It is not implausible to think at first that we become aware of facts about our mental likes through a kind of quasi-perceptual awareness of these ‘mental objects.’

Shoemaker’s main argument against the object-perception model here is that if we construe our awareness of, say, after-images on the object-perception model, we are thereby construing illusions and hallucinations on an act-object model. More simply: we are taking hallucinations to be cases in which we are aware of a mental object. But, says Shoemaker, the Argument from Illusion takes us from this claim to the sense datum theory of perception: the view that the immediate objects of our awareness are always mental items, (loosely speaking) our own ideas.

The question of our awareness of sensations raises questions about qualia, and whether what we are aware of in being aware of a sensation is exhausted by the representational properties of the relevant sensory experience. Shoemaker will return to this topic later.

1.3.4 *The role of attention in knowledge of sensations and intentional states*

According to the object-perception model of perception, we are able to shift our attention from one to another object of our perception with a view toward getting more information about one of those objects. Does this part of the object-perception model, which is [6] of the object-model conditions, carry over to the case of self-knowledge?

A first thought is that it applies to our knowledge of sensations — since we can attend to one sensation or bodily experience rather than another, but not to beliefs, desires, and other intentional states. One might think this because, as Shoemaker says, it seems that attending to a belief is simply a matter of thinking the proposition which is the content of the belief. (You might also do things like think about the source of the belief; but this is not supposed to be a kind of attending to the belief.)

But, Shoemaker says, there is reason to doubt the application of the object-model’s view of the relationship between attention and perception to even our knowledge of our own sensations. The reason: when we shift our attention from one sensation to another (in the case of veridical experience) we are really just shifting our attention from one object of perception to another. The experiment to back this up: (1) Hold both hands before your face, and shift your attention from one hand to the other. (2) Hold both hands before your face, but this time shift your attention from your visual experience of one to your visual experience of the other. (3) Now compare (1) and (2): did you do anything different? Shoemaker’s thought is that we should answer (3) in the negative. But, if so, there is really no such thing as shifting your attention from one sensation to another; or, if there is, it is just a matter of shifting

one's attention from one object of perception to another.

Objection: what about cases of non-veridical perceptions (such as after-images) and bodily sensations (such as pains)? In these cases one can't 'explain away' attention shifts in the same way.

Shoemaker's reply: strictly speaking, there are no such things as after-images, and so no mental objects like that to be the objects of attention (this is forced on us by the argument from illusion, and the desire to avoid a sense-datum theory of perception). After-images are not mental objects, but the "non-existent intentional objects of outer sense." (266) Then we can give the same reply as above: the shifts in attention in question are just shifts in attention to objects of ordinary perception, not any special shift in attention among mental objects.

Objection 1: talk of attending to non-existent objects, even if we call them 'intentional', is of dubious coherence.

Objection 2: how does this handle bodily sensations, like pains? These are not cases of perception, so we cannot explain away shifts in attention between pains, itches, etc. as shifts in attention between different objects of perception.

Reply to objection 2: bodily conditions are in fact a kind of perception of bodily conditions (pain might be, e.g., a kind of perception of bodily damage).

What is an itch a perception of? (Maybe whatever bodily condition it covaries with under normal conditions.)

1.4 Result of rejecting the object-perception model

Shoemaker thinks that the object-perception model of self-knowledge brings with it a number of objectionable features; one of the main payoffs of rejecting that model is that we then need not accept its consequences. Among these are:

1. The object-perception model incorporates an act-object model of self-knowledge; which means that for every instance of self-knowledge (of the most fundamental kind) there will be a corresponding mental object. This sticks us with mental objects like after-images and sense-data.
2. This model also encourages us to think of the objects of self-knowledge as simple: they must be able to be, so to speak, taken in at a glance. But this seems to count against functionalism: the view that mental states are defined by their causal connections to other mental states. How could we find that a given mental object was a sensation of a dog or a belief about grass if the identities of those mental states depended on a complex web of connections to other mental states? (The problem is made worse, as Shoemaker notes, by the fact that many of these connections will be merely counterfactual.)

Most people take functionalism to be the only plausible going model of the mind. Quick discussion of identity theories and behaviorist theories and their failings.

As Shoemaker says at the end of Lecture I, he wants a non-observational model of self-knowledge. But before the ground is cleared for this, we need to consider a model of self-knowledge based on a different stereotype of perception than the object-perception model.

2 Self-knowledge as ‘broad perception’

2.1 Perception according to the broad perceptual model

Shoemaker begins Lecture II by stating two opposed theses of the relationship between mental states and our knowledge of those states.

The first *transparency* thesis, which Shoemaker calls Cartesianism, holds that mental states are strongly ‘self-intimating’: it is the nature of mental states to be known, and so we cannot be in a mental state without knowing that we are.

The second incorporates the final two conditions which were a part of the object-perceptual mode, which I called there the *independence conditions*. Condition [7] stated that perceptions are causally produced by the objects of perception, while condition [8] stated that the objects of perception can exist independently of their being perceived; the analogous theses about self-knowledge hold that beliefs about one’s mental states are causally produced by those mental states, and that those mental states can exist independently of their being known by the agent in the mental state.

These claims about perception are very plausible; the analogous points are more contentious about self-knowledge. We should be clear about two points at the start: (i) these claims conflict with Cartesianism; and (ii) the proponent of the view that self-knowledge is made possible by a kind of inner perception should endorse [7] and [8] as applied to self-knowledge.

Shoemaker calls the view of self-knowledge which endorses these independence conditions and opposes Cartesianism the *broad perceptual model*.

Doesn’t the object-perception model of self-knowledge also endorse [7] and [8]? Yes. Shoemaker is taking the object-perception model to have been discredited by the arguments of Lecture I, discussed above. The broad perceptual model consists only of claims that are essential to the view that introspection is a kind of inner sense: the neutral conditions [1] and [2], and the independence conditions [7] and [8], without the object-model conditions:

Neutral conditions

- [1] Sense perceptions involve an organ of perception under one’s voluntary control.
- [2] Sense perceptions involve the occurrence of sense experiences or sense impressions distinct from the object perceived.

Independence conditions

- [7] Sense perceptions are causally produced by the objects of perception.
- [8] Sense perceptions typically provide awareness of objects and states of affairs which exist independently of their being perceived.

The objections to the perceptual view of self-knowledge in Lecture I all turned on object-model conditions (the quick objections based on [1] and [2] at the outset aside). We have now to ask whether there are problems with the very idea that self-knowledge is a kind of perception, or whether there are merely problems with cashing that idea out in terms of one paradigm of how perception works.

2.2 The concept of self-blindness

Shoemaker thinks that the problems with the view that self-knowledge is made possible by a kind of inner sense run very deep. His plausible idea is that anyone who holds that self-knowledge is akin to perception must hold that the objects of self-knowledge — bodily sensations, perceptual experiences, intentional mental states, the self — must be at least conceptually independent of our knowing about them, just as the objects of perception are typically at least conceptually independent of our perceiving them. This seems right; if a kind of knowledge does not satisfy this feature, then it is difficult to see what the grounds might be for taking it to be based on a kind of sense (whether inner or outer).

His aim, then, in arguing against perceptual views of self-knowledge, is to argue that mental states are *not* conceptually independent of our knowing about them.

This raises an important question: what is required for one thing to be conceptually independent of another? Shoemaker, after all, grants the falsity of the Cartesian transparency thesis: he does think that mental states can exist without being known. So in what sense is there a conceptual connection between mental states and their being known?

Shoemaker is not completely explicit about this, but I think that it is fairly clear what he had in mind. Given the shape of his argument he would, I think, accept the following:

If *A*s are conceptually independent of *B*'s, then there can be no necessary truths connecting *A*'s and *B*'s.

This criterion does mark out an important distinction. Consider, for example, the connection between the properties of being human and of having a life span of less than 120 years. There is a fairly strong connection between these properties: if one has the former, then one's chance of having the latter is pretty high. But, it is natural to think, these properties are in a good sense independent: after all, it *could have*

been the case that the average human life span was 200 years. This is the kind of independence that the objects of perception have of their being perceived. There may be a good correlation between the property of being a purebred poodle and the property of being perceived by some human being; but the former property is still independent of the latter, because there is no incoherence or impossibility in the assumption that such a correlation should fail to obtain. In short: there are, in general, no necessary connections (even if there are reliable contingent connections) between the properties of objects perceived by human beings and the property of their being perceived.

If you grant this, as it seems one should, then Shoemaker's argumentative task is pretty clear: argue that there are some necessary connections between being in certain mental states and having knowledge that one is in those mental states. To argue successfully that this is the case will be enough to show that the relationship between our beliefs about our mental states and those mental states is quite different from the relationship between our perceptions and their objects.

Shoemaker's way of putting the thesis he is arguing for is that it is impossible that creatures should be *self-blind* with respect to one or another class of their mental states.

He defines self-blindness as follows:

- A* is self-blind with respect to a mental state type *M* if and only if
- (i) *A* has the conceptual capacities to conceive of *M*'s, &
 - (ii) *A* lacks introspective access to *M*'s.

A number of points about this definition.

1. As Shoemaker notes, this definition leaves it open that *A* might be able to find out about whether he is in an *M* by means other than introspection by, for example, observing their own behavior.
2. To understand the notion of self-blindness, we need already to have a grip on the idea of introspective access. This needn't be problematic; we do have the intuitive idea that we can find out about our own mental states via some method not suitable for finding out about the mental states of others. But it is important to note that Shoemaker's discussion here presupposes some demarcation of introspective access which is not explained.
3. A more critical point: why is Shoemaker entitled to restrict attention to agents with the conceptual capacities to think about the mental states in question? What is wrong with the proponent of the broad perceptual model just showing that sensations are conceptually independent of our knowing about them by simply citing an example of a lower animal which has sensations but lacks any beliefs about them?

Reply on Shoemaker's behalf: the definition of 'conceptually independent' given above says that there must be *no* necessary truths connecting the two classes of facts in question. Showing that there is some necessary connection between some

occurrences of mental states and some beliefs about those mental states is enough to show that the two are not conceptually independent!

Objection: this is simply not good enough. By restricting one of two domains of facts artificially, it then becomes too easy to show that nothing is conceptually independent of anything else. Take our previous example: plausibly, the property of having a life span less than 120 years and the property of being human are conceptually independent. But it is a necessary truth that for all x , if x is human and is short-lived by our standards, then x 's life span was less than 120 years. So it must be that, after all, these properties are not conceptually independent — there is a necessary connection between being human and having a life span of less than 120 years! But surely this is not plausible; the connection between the two properties is purely contingent, and we have just manufactured a necessary truth by artificially restricting the facts that can be considered. But this is just what Shoemaker is doing.

Better but less precise reply on Shoemaker's behalf: necessary truths require some explanation. If there is a necessarily true conditional running from a class A of facts to the B facts, then there must be some explanation of why this is so. One standard kind of explanation will be in terms of the nature of the A facts. E.g., it is a necessary truth that if something is a canine, then it is an animal. Plausibly, this is because being an animal is part of what it is to be a dog. Now suppose that there is a necessary truth running from the property of being in a certain kind of mental state + having certain conceptual capacities to having knowledge of that mental state (in at least some cases). If we take the form of explanation of this necessary truth which seems appropriate to the case of canines and animals, then we will end up saying that it is part of the nature of the given class of mental states – be they pains, or beliefs, or perceptual experiences, or whatever – to be known by certain kinds of subjects. *And this is just the result that Shoemaker wants.* If this is true of mental states, then it seems that the relationship between mental states and their being known is quite different than the relationship between the objects of perception and their being perceived (because of the plausibility of the independence conditions [7] and [8]). And this is enough to show that self-knowledge cannot be understood just as a kind of inner perception.

This does not mean that there cannot sometimes be other explanations for the necessary truth of a conditional of this form. Two, in particular, come to mind.

The first of these has not been discussed so far. Consider the following plausibly necessary truth: if x is to the west of y , then y is to the east of x . Plausibly, if this is to be explained in terms of the nature of either the relational property of being to the west of something or of being to the east of something, it is to be explained in terms of both: the two properties are interdependent. This is a disanalogy with the canine/animal case. But this paradigm of necessary truth would also be enough to suit Shoemaker's purposes; if mental states and their being known about are interdependent, then this too would be enough to show that the relationship between those mental states and their being known about is quite different than that between the objects of perception and their being perceived.

The second alternative explanation of a necessary conditional seems a bit more wor-

rying for Shoemaker; this explanation is exemplified by the conditional considered above:

Necessarily, for all x , if x is human and is short-lived by our standards, then x 's life span is less than 120 years.

Here we are imagining that someone is trying to use the truth of this claim to show that there is a conceptual connection between being human and having a life span of less than 120 years; to that extent it is on par with Shoemaker's use of the conditional

Necessarily, for all x , if x has mental states of type M and has the conceptual capacities to think about M 's, then x 's has introspective access to M 's.

to show that there is a conceptual connection between being in a mental state and having a certain sort of access to knowledge about it. The problem here is that we do not think that the first conditional is to be explained in terms of the nature of being human; rather, it is to be explained in terms of the restriction put on instances of that property to human beings short-lived by our standards. Can the proponent of perceptual models of self-knowledge not appeal to the same sort of explanation of Shoemaker's conditional? If so, then it is in the nature of having certain kinds of conceptual capacities that one have introspective access to certain sorts of mental states, and *not* in the nature of certain mental states to be introspectively accessible.

Let's recap so far. If Shoemaker's conditional – which is equivalent to the claim that self-blindness is impossible – is true, and if we accept a plausible claim about the explanation of necessary truths, then we are left with two possible views:

- (i) It is in the nature of certain mental states to be introspectively accessible (in certain cases).
 - (ii) It is in the nature of having the conceptual capacities to conceive of (certain kinds of) mental states that one have introspective access to those mental states.
- (i) will certainly suit Shoemaker's purposes in arguing against the perceptual model of self-knowledge – would (ii) suit his purposes as well?

This is not obvious to me.

The structure of the argument of the rest of Lecture II is to consider, for a host of different mental states, whether self-blindness with respect to that mental state is possible. Shoemaker's answer is in each case the same: no.

2.3 Self-blindness and bodily sensations

Shoemaker acknowledges that we can be unaware of pains: examples include a wounded soldier who does not notice that he is in pain until a battle is over, or

an athlete who does not notice the pain of an injury until after the game. But this does not show the possibility of self-blindness: self-blindness requires not just the possibility of occasional failures of introspection, but rather the complete absence of introspective capacities.

The structure of the argument against self-blindness in the case of pain is the same as in later sections: try to imagine someone who is self-blind in the required way, and then show some absurdity in the supposition.

Does someone who is self-blind with respect to pain find pains unpleasant?

Possibility 1: the creature does find the pains unpleasant, but is just unaware that they do so.

Obj. 1: Finding the pains unpleasant could not have any of its usual behavioral consequences in this case. If the person, e.g., went to the closet to get aspirin, the creature would not be able to give any account of why it was doing these things. It would therefore cease to be an agent in a certain sense (with respect to its pain behavior). We are supposed to doubt then whether we should really say that the agent in question finds the pains unpleasant.

Obj. 2: if the creature does find the pains unpleasant, then the creature is not self-blind to pain at all; rather, such a creature is self-blind with respect to certain beliefs about pain.

Possibility 2: the creature does not find the pains unpleasant; but the pains are unpleasant nonetheless.

Obj.: Shoemaker objects that this view “makes the connection between a state’s being unpleasant and its subject’s wanting to be rid of it a purely contingent connection.” (275) But if this were so, then it should be possible that, e.g., most of “what we take to be innocent pastimes” could be extremely unpleasant, even though we aren’t able to notice this aspect of them. Shoemaker summarizes this possibility as: “His pains hurt, but they don’t hurt *him*.” Shoemaker thinks that this is absurd, and so he does not think that this kind of self-blindness is a genuine possibility. It follows that the connection between a state’s being unpleasant and its subject’s wanting to be rid of it is more than contingent.

Possibility 3: the pains in question are not unpleasant (and the creature does not find them to be unpleasant).

Obj.: Shoemaker never considers this possibility, but the idea is supposed to be that mental events which were neither unpleasant nor found to be unpleasant would not be pains at all.

Shoemaker thinks that these points about pains carry over to other bodily sensations — itches, tingles, etc.

2.4 *Self-blindness and perceptual experiences*

Shoemaker's question in §III is: Could someone have perceptual experiences just like ours, and yet not have these experiences available for introspection? (Note that the question is not whether someone could have states which registered information from the environment reliably without having access to the information registered; this is not only possible, but presumably widespread.) As Shoemaker uses the term, a perceptual experience will be a matter of things appearing/looking a certain way; we are asking whether we can have perceptual experiences in this sense without having introspective access to how things appear to us.

We can suppose that the agent in question gathers information as usual from objects in her environment; she can also make judgements on the basis of her interaction with the environment. We can suppose that she is able to make judgements not only like 'There's a tree over there' but also 'I see a tree over there.' (She can make the latter judgement because she can make judgements about the environment on the basis of information gathered, and can know which organ the information was gathered through.) It's just that she has no introspective access to how the tree appears to her.

We must suppose that the agent is sensitive to information from the environment in the same ways that we are. The first question is then: what judgements will such an agent make in cases of illusion? E.g., if she is standing in front of a white wall illuminated with red light, what judgement will she make about the color of the wall?

Possibility 1: she will use background beliefs about the perceptual situation to arrive at beliefs in the normal way; in this situation, she will combine her knowledge of the fact that the wall is bathed in red light with her knowledge of the way that the wall is presented to her in perception, to arrive at the perceptual judgement that the wall is white.

Obj.: this supposes that she has knowledge of the way the wall looks to her, which is just to suppose that she is not self-blind with respect to her perceptual experiences.

Possibility 2: she will use background beliefs (about, e.g., the color of the light projected on the wall) to arrive at the right judgement about the color of the wall; but she will not do this via an inference from this belief along with the way the wall appears, but will rather be 'hard-wired' to immediately arrive at the normal perceptual judgement on the basis of input through her eyes and the relevant belief. (Such an agent will, as Shoemaker notes, often be able to infer claims about perceptual experience from the perceptual judgement and the background beliefs, but this knowledge will be inferential and so not an instance of introspective access to how things appear to one.)

Obj. 1: This sort of agent will not always make the same judgements about her perceptions as we will. Suppose, e.g., that we put her in a position where she does not know anything about the perceptual conditions under which she is viewing the wall: whether, e.g., she is viewing it under red light, or normal light, or She will, like a normal person, be unsure what color the wall is. But she will also be unable

to say how the wall looks to her.

Reply to obj. 1: How does this show that self-blindness is impossible? It shows us that self-blindness would be detectable (at least in subjects with normal verbal abilities); but I don't see a straightforward argument from this kind of case to Shoemaker's conclusion that this individual would fail to have perceptual experiences in the same sense that we do. At least, I don't see any such argument which would not just assume the conclusion.

Obj. 2: Suppose that we took someone who is blind (and so presumably has no visual experiences). If we linked up their brain and optical system in such a way that their behavior would be just like a self-blind person of type 2, would this be enough to make it the case that they had perceptual experiences just like ours, but just was not aware of this?

Reply to obj. 2: The opponent of Shoemaker should not say that this person is self-blind — but she doesn't have to. She can claim that it is possible for someone to have perceptual experiences without introspective access to them without claiming that this sort of surgery on a blind person would turn them into such a self-blind agent. (Shoemaker would respond, perhaps, that this is implausible, since the two might be 'functionally identical.' But the proponent of the possibility of perceptual self-blindness is unlikely to accept the premise that having perceptual experiences is a functionally definable property.)

Possibility 3: she will not use background beliefs at all, but will always make perceptual judgements on the basis of the information taken in by the relevant sense organ. So in the case she described above, she would judge that the wall was red.

Shoemaker does not discuss this possibility; it is not obvious to me what is impossible or absurd about it. To be sure, this sort of person would form perceptual judgements in a very different way than we do, and perception would play a very different role in her life than in ours; but we need some argument from these claims to the conclusion that she would not have perceptual experiences in the same sense that we do.

One is also tempted to make some sort of evolutionary argument: why would someone have perceptual experiences that were not introspectively accessible? What role could such experiences play in helping the person get around in the world which 'appearanceless' information gathering could not? But, pretty clearly, this sort of 'why would' question does not count for very much when what is in question is the conceptual possibility of a self-blind agent.

Conclusion: I do not see that Shoemaker has given us any reason (in this section) to discount the possibility of perceptual self-blindness.

2.5 Self-blindness and the will

Shoemaker's first argument against the possibility of self-blindness with respect to the will — i.e., against the possibility of someone undertaking intentional actions

without being able to know which actions she is trying to perform — involves verbal behavior.

It runs as follows: (i) If someone is self-blind with respect to their own will, then they should be able to tell us that they were unaware of particular exercises of their will. (ii) But either the act of telling us this will be an intentional action, or it won't be. (iii) If it is not, then it is no evidence that the agent is self-blind. (iv) If it is, then she is telling us that she is self-blind “with respect to the very activity she is engaged in.” (279) But this is incoherent; someone would then be making an assertion which ‘says of itself’ that it is not an intentional assertion.

I'm not sure that this argument is so convincing. Briefly: (i) is not obvious; there are kinds of evidence other than self-ascription, as the case of the agent self-blind with respect to her own perceptual experiences shows. Further, it may be possible to imagine an agent who consistently reports self-blindness with respect to past actions, and is unable to consider present ones.

Shoemaker further argues as follows (280): many plans for action incorporate both verbal and non-verbal behaviors coordinated in pursuit of one goal. So an agent who was self-blind with respect to her own will, but otherwise just like us with respect to the will, would exhibit verbal and non-verbal behaviors combined in this way, but with verbal claims professing introspective ignorance of what is going on interspersed. But in such a case we would not say that the author of the professions of ignorance was the agent of the coordinated plan; we would either say that what appeared to be a plan really failed to be intentional action, or that the plan was the intentional action of another ‘person’ in the same body with the author of the professions of ignorance.

Either way, Shoemaker thinks, we do not have a genuine case of self-blindness. In the first case, we do not have an agent who is performing intentional actions, and hence do not have an agent self-blind with respect to those actions. In the second case, the agent who is blind with respect to intentional actions is not the agent of those actions — and hence is not *self*-blind.

Reply: Again, here Shoemaker seems to be supposing that an agent self-blind with respect to her own actions must always be telling us this while the actions are going on. Is the case obviously so absurd if we imagine the agent only professing ignorance of past actions (in the sense of saying that he did not know while performing such-and-such that he was doing so)?

In the end, Shoemaker's conclusion seems very plausible — we would not count an agent as the author of intentional actions unless she knew, at least in standard cases, which intentional actions she was pursuing. But it is not obvious to me that his example provides further support for this intuition.

2.6 Self-blindness, beliefs, and desires

Shoemaker gives three different arguments against the possibility of self-blindness with respect to beliefs and desires (here I'll focus on the case of belief).

2.6.1 The argument from self-ascription

Argument. Suppose that a rational agent with the same conceptual capacities as a normal human being is self-blind with respect to his beliefs. If he has the concept of belief, he will answer the same to 'Is it true that p ?' and 'Do you believe that p ?' (Always?) But he will also know that an affirmative answer to the latter will have the same truth value as 'I believe that p ', which will cause him sometimes to volunteer the latter, and so ascribe a belief to himself.

Reply 1. What is this supposed to show? Shoemaker says (283) that it shows a connection between the ability to give verbal expression to beliefs and the ability to self-ascribe them. But what does this have to do with self-blindness?

If this argument were on the pattern of Shoemaker's earlier arguments, the conclusion would be that the agent turns out not to be self-blind after all. But this does not seem to be obvious, because it is not obvious that the agent imagined does really lack introspective access to the contents of his beliefs.

Reply 2. And of course this argument only goes through if we assume that the self-blind person has the concept of belief. – Shoemaker will reply by insisting that self-blindness is not to entail any cognitive deficiency, and so that this assumption is legitimate. — But this seems not to be enough. What Shoemaker ultimately wants is an argument that it is in the nature of beliefs to be known. What we are given is always an argument that some conditional of the form

For any x , if x has beliefs and meets condition C , x has introspective access to his beliefs.

But this will entail the intended conclusion only if the necessity of the conditional cannot be explained in terms of the nature of condition C as well as the nature of the property of having beliefs. But here the condition C is possession of the concept of belief; and it is not completely implausible that it is part of the conditions necessary for possession of this concept that one have introspective access to one's beliefs. If this is right, then Shoemaker's argument here could go through, and self-blindness in the imagined case *could* be impossible, without this showing anything about the nature of belief.

2.6.2 The argument from acting 'on the basis of' belief

Argument. Take a self-blind agent who believes p . Since we are supposing that she is rational, we must suppose that she will take p as a premise in her reasoning and will

act on the assumption that it is true. She will further know that when one acts on the assumption that a proposition is true, one acts just as one acts when one believes that proposition. So she will know that she is planning to act just as one who believes p would act. Given that she also knows that it can be useful to inform other people (presuming them cooperative) of one's courses of action, she will sometimes (when in the presence of such people) find that it advances her plans to say 'I believe p .'

Reply. Again, what does this show? It is not obvious that the agent in question really fails to be self-blind, and so it is not obvious that we have exposed some kind of absurdity in the supposition that an agent be self-blind with respect to her own beliefs.

Shoemaker says that this reply misconstrues the point of the argument:

"The reasoning for pointing out that such reasoning is available is not to suggest that it regularly goes on in us – obviously it doesn't – but rather to point out that in order to explain the behavior we take as showing that people have certain higher-order beliefs, beliefs about their first-order beliefs, we do not need to attribute to them anything beyond what is needed to give them first-order beliefs along plus normal intelligence, rationality, and conceptual capacity." (284)

It is not obvious what Shoemaker is saying here. Following are two interpretations:

Interpretation 1. We say that our self-blind agent is self-blind because he must go through this process of reasoning to arrive at beliefs about his own beliefs. But such an agent could simply learn to pursue the same kinds of courses of action without self-consciously reasoning through the steps. But his pursuing these actions would be enough for us to attribute to him the belief that he believes p , and so would be enough for us to arrive at the thesis that he is not self-blind, after all. If introspection were only possible through a kind of inner sense, it would be impossible for him to acquire this introspective ability in this way.

Interpretation 2. As in the above interpretation, we note that the agent in question could learn to behave just as we do simply by not going through these kinds of arguments to arrive at beliefs about his own beliefs. But if this is all that is required for this sort of action, isn't it bizarre to suppose that evolution endowed us with an extra inner sense, which is after all redundant in that we could act just as we do without it? (see the discussion on the top of p. 285)

2.6.3 The argument from the possibility of deliberation

Argument. Rational creatures can modify their beliefs in response to evidence; so we must imagine that self-blind creatures are also capable of this. But sometimes it is not obvious how evidence should require us to change our beliefs. In such cases, we must evaluate arguments involving the evidence. But which arguments are relevant to modification of one's beliefs depends on what one believes. (In one obvious kind

of case the arguments may, e.g., include arguments whose conclusion is the negation of something one believes, and whose premises include the evidence along with other propositions that one believes.) So in order to be fully rational one must have the introspective access to one's beliefs required to choose which arguments are relevant to modification of one's beliefs in the light of new evidence.

See reply 2 to the 'Argument from self-ascription.'

3 Shoemaker's positive picture

3.1 *The problem*

Shoemaker's deep reason for rejecting perceptual accounts of self-knowledge is that such accounts treat mental phenomena as independent of our knowing about them in a way that they are not. Because self-blindness with respect to mental states is impossible, Shoemaker thinks, we need some conception of mental states and our knowledge of them which explains this impossibility rather than making it difficult to understand.

Not enough to regard mental states as 'self-intimating' in some mysterious way.

The basic idea: if it is in the nature of mental states to be known, then an epistemology of self-knowledge should be based in the philosophy of mind. We need a view of what it is to have a pain, or a belief, or to try to perform an action, which can explain the fact that it is impossible to be in such states without knowing that one is in them.

3.2 *The functionalist view of mental states*

Behaviorism and the intuition that beliefs are not patterns of actions, but internal states which cause patterns of actions.. Identity theories and the problem of multiple realizability. The idea behind functionalism. Why it is plausible to think that functionalism is the only plausible materialist view of the mind.

A sample functionalist theory for perceptual beliefs.

3.3 *Shoemaker's functionalist explanation of self-knowledge*

The investigation of self-knowledge places two constraints on our account of the mental states to which we have privileged access: (i) our account must explain why complete ignorance of such states is not possible, and (ii) it must explain how ignorance of individual pains, beliefs, acts of the will, etc. is possible.

Shoemaker gives us the following two-stage (functionalist) picture of mental states:

According to the functionalist, any mental state M will be defined by a causal role C . (Think of causal roles as complex properties.) The idea is then that a physical state x of an agent is a token of M iff x has C . We will then, in Shoemaker's terminology, call x the *core realization* of M (in that agent). As he puts it, the mental state will come and go as the core realization comes and goes.

For x to have C — for a physical state of the brain, e.g., to instantiate a certain causal role — is for that state to cause and be caused by certain things, and presumably also to have dispositions to cause and be caused by others. But there will also be facts about the agent in question which allow x to have that causal role.

Suppose, e.g., that x is standardly caused by visual contact with red things; suppose further that it represents redness in virtue of having this causal role. There will presumably be a class of facts about the way that, e.g., state x is related to the eyes and the rest of the visual system which allow x to play that causal role. In this case, we can imagine that if those facts failed to obtain, then the agent would no longer be disposed to be in state x when in visual contact with red things. All facts related to the core realization in this way, along with the core realization itself, will be a part of what Shoemaker calls the *total realization* of a mental state (in an agent).

So: the core realization of a mental state in an agent is a proper part of the total realization of that mental state in that agent.

Further, the notions of core and total realizations are not only relative to a mental state type, but also to an agent. Different agents can have different core realizations and total realizations of the belief that, e.g., sled riding is fun.

Given this framework for thinking about the natures of mental states and their physical realizations, how can we explain the impossibility of self-blindness?

Shoemaker gives us two possibilities:

1. The total realization of a mental state m and the belief that one is in m might overlap (without the two having the same core realization). Then it would be the case that being in m tended to cause one to have the belief that one is in m . If this were a necessary truth — if it were true that for any possible agent, the total realizations of m and the belief that one is in m overlapped — then, Shoemaker says, “it would be of the essence of that state that under certain circumstances it tends to give rise to that belief.” (288)

How, you might ask, could this be a necessary truth? It's not easy to see how this could happen without an overlap in functional roles, and an overlap in core realization. (Given the example he uses, though, Shoemaker seems not to think this.)

2. A mental state m and the belief that one is in m could have exactly the same core realization for an agent, while the total realization of the former could be a proper part of the total realization of the latter. Shoemaker thinks that this is a plausible story about first and second order beliefs.

A confusing part of this claim: if core realization is defined by functional role, then

does the claim that m and the belief that one is in m have the same core realization entail — implausibly — that they have the same associated functional roles? No. They could have different functional roles implemented by the same physical state; but for the state to play one of the causal roles it may require, e.g., a connection to some other state which the other causal role does not require.

But could it be a necessary truth that two states have the same core realization without their having the same associated functional roles (and so, according to the functionalist, being the same mental state)?

Should we be worried about a priori arguments which yield results about how the brain must work?

4 The problem posed by phenomenal character

The topic of Shoemaker's third lecture is the question of whether the functionalist account of self-knowledge he expounds at the end of Lecture II can account for our knowledge of aspects of our own perceptual states. Shoemaker has already argued that self-blindness with respect to such states is impossible; but it may seem that certain aspects of perceptual experience make his functionalist explanation of this fact implausible here.

Shoemaker is not very upfront about what these aspects are. But the basic idea is that the fact that perceptual experiences have a certain 'phenomenal feel' may make it unlikely that they can be given a functionalist definition. And if they cannot be functionally defined, then it does not look as though Shoemaker's positive picture will be applicable to them.

Discussion of spectrum inversion as a case which might make one doubt functionalism.

Since the question here is about introspective access to perceptual experiences, Shoemaker spends most of the Lecture discussing not the nature of self-knowledge, but the nature of perception: before discussing the nature of self-knowledge of perceptual experiences, he thinks that we'll have to get clear about what aspects of perceptual experience we have introspective access to. And that in turn will require getting clear on the nature of perception. That's where we'll begin as well.

4.1 Perception and its objects

Shoemaker says at one point,

"I have deliberately used the vague phrase 'the phenomenal character we are confronted with,' which leaves it unspecified what is supposed to *have* this phenomenal character — the external objects perceived, or our subjective experiences of them. And the problem is, in part, about how to eliminate this vagueness." (294)

Shoemaker's idea is that our self-knowledge of our own perceptual states amounts to knowledge of, to use another vague phrase, the way things *seem* or *feel* to me when I am having a perceptual experience. It is natural to think of these seemings or feelings as acquaintances with certain properties. But then we must ask: are these properties properties of external objects, or merely properties of our own subjective experiences? Shoemaker thinks that our answer to this question will affect our account of self-knowledge of perceptual experiences.

He canvasses several possibilities:

4.1.1 *The 'simple view'*

Perhaps the simplest view is just that in cases of veridical perception, the qualities or properties with which we are confronted in perception are, so to speak, revelations of properties of external objects: redness is a property of external things which is just as we perceive it to be.

Shoemaker does not say much about this, but he does say the following:

“But reflection on the disparity between the scientific and the manifest image makes inescapable the conclusion that . . . the phenomenal character we are confronted with in color experience is due not simply to what there is in our environment but also, in part, to *our* nature, namely the nature of our sensory apparatus and constitution. . . . How *could* the phenomenal character we are confronted with be solely determined by what is in the environment, if what is there in the environment is anything like what science tells us is there?” (294)

There are perhaps two different arguments contained in this passage. (i) It is obvious that colors cannot be properties of the surfaces of things, because our perceiving them depends on facts about our sensory systems as well as facts about the surfaces of things. (ii) Science tells us that the surfaces of things do not have the properties that our experience attributes to them.

(i) is manifestly a bad argument, bad enough that it is difficult to interpret Shoemaker as making it. Compare: the shape of a key, and the shape of the keyhole into which it fits.

(ii) may also not seem particularly convincing; it depends on one's views about the relationship between the properties that there are and the properties revealed by science, and there is more than one way to think about this. One can think that all the properties of objects are determined by the properties revealed by science while thinking that colors are properties of the surfaces of objects and are not among the properties studied by science.

4.1.2 *Literal projectivism*

The subjectivist view that color properties are properties not of external objects, but only of our experiences of those objects. It's just that in experience we make a characteristic mistake: we project these properties of our experiences onto the objects we are perceiving.

Shoemaker objects that this view faces a dilemma. On the one hand, colors may be properties of parts of the visual field, which are then projected onto spatial regions of the world; but in this case it lapses into sense datum theory. On the other hand, colors may be properties of events of experiencing the world; but in this case it is incoherent, since we are at a loss to understand how colors, which are only conceivable as properties of extended things, could be a property of such events.

4.1.3 *Figural projectivism*

The subjectivist view that says that there are two kinds of properties at work here: (i) properties of experiences, in virtue of which they 'seem' the way they do; and (ii) the properties experiences represent objects in the world as having.

Advantages of the view: it avoids the supposed implausibility of objectivist views, and also avoids the incoherence associated with literal projectivism, since there are no properties of experiences which are 'projected' onto the spatial regions of the world.

Disadvantages: it secures this second advantage only at the cost of making it the case that properties of type (ii), which are the properties that our experiences represent objects as having, are *not instantiated anywhere*. They are not properties of objects in the world (on pain of lapsing into objectivism about color) and they are not properties of experiences (on pain of lapsing into literal projectivism).

4.1.4 *Shoemaker's positive view of phenomena properties and qualia*

Recall again that 'phenomenal qualities' are the properties with which we are confronted in experience. We have now briefly considered a few views about what these properties could be: simple properties of the surfaces of objects (the simple view), properties of our experiences (literal projectivism), or properties of nothing at all (figurative projectivism). Shoemaker has found all of these alternatives wanting.

Shoemaker thinks that the properties our experiences represent objects as having are not properties of our experiences, but of objects. But they are a special kind of property of objects: the property of causing certain qualia under normal conditions.

A natural question is why he holds a view like this. To answer this question, we'll need to be clearer about what *qualia* are supposed to be.

It is not easy to figure out exactly what he's talking about here. To a first approximation, qualia are properties of experiences in virtue of which those experiences are

of a certain phenomenal type, or ‘seem’ or ‘feel’ a certain way. So, assuming that we have some fix on the ‘phenomenal feel’ of an experience, let’s employ the following claim about qualia to fix ideas:

Definition of qualia

Two experiences have the same phenomenal feel iff they instantiate the same quale.

These are not the properties which experiences represent objects as having, or which objects seem to have when experiencing them in a certain way. Qualia are properties of experiences, not properties which either are of or seem to be of the objects of experience.

Let’s see how we get from here to Shoemaker’s view. First, because he takes spectrum inversion to be possible, he thinks that

Possibility of spectrum inversion

Two objects can have the same color and yet cause different perceivers to have phenomenally different experiences of them (under identical perceptual conditions) because one perceiver is spectrum-inverted relative to the other.

From the Definition of Qualia along with the Possibility of Spectrum Inversion we get

Two objects can have the same color and yet cause different perceivers to have experiences with different qualia (under identical perceptual conditions) because one perceiver is spectrum-inverted relative to the other.

Further, he thinks that phenomenal properties, or the properties experiences represent objects as having, should satisfy the following requirement:

The phenomenal property/phenomenal likeness link

Two perceivers are having perceptions of the same phenomenal property iff the two perceivers are having experiences which are phenomenally alike (or have the same ‘phenomenal feel.’)

From this link plus the Definition of Qualia we get

Two perceivers are having perceptions of the same phenomenal property iff the two perceivers are having experiences with the same qualia.

Using these three principles and two corollaries we can get a better idea of the kinds of properties - the phenomenal properties - Shoemaker thinks we are confronted with in experience. Let Y be the quale that my experiences as of banana yellow have. Then the properties objects are represented as having in experience are properties like *the property of currently producing a Y -experience in someone viewing the object under normal lighting conditions* (298). More precisely, phenomenal properties may be thought of as obeying the following conditional:

For every phenomenal property P there is an associated quale Q^P such that, necessarily, an object has P (at a time) if and only if it is (at that time) causing someone perceiving that object under normal conditions (e.g., normal lighting conditions) to have an experience with Q^P .

Phenomenal properties are the properties with which experience ‘confronts us’. These are properties which objects seem to have. Qualia are properties of experiences; objects do not seem to have them. But the two are related: the quale of an experience determines the phenomenal properties it represents an object as having.

4.1.5 *An objection to Shoemaker’s view of phenomenal properties*

An immediate objection to Shoemaker’s view of the kinds of properties that experiences represent objects as having is that those properties are not monadic properties of objects. Surely when I have an experience as of a scarlet object I am representing that object as having a certain monadic property; I am not representing it as being related to me in a certain way, and certainly not as causing in me an experience with a certain quale.

Shoemaker replies that the way properties are represented in experience is not an infallible guide to their nature. The property of *being to the left of* is represented as dyadic, but is really triadic. The property of *feeling heaving* is represented as monadic, but is really dyadic (or triadic).

4.1.6 *A Shoemaker-ian view of color*

With this on the table, we can also conclude that Shoemaker does not think that colors are phenomenal properties — he does not think that colors are properties we are confronted with in experience. We know this because of the way he states the possibility of spectrum inversion: two agents could perceive objects *with the same color* and yet be having experiences with different qualia. Add to this the premise — which Shoemaker accepts — that neither of the two agents spectrum-inverted relative to each other is having a perception more or less veridical than the other, and we get the result that phenomenal properties are not colors.

The identification of colors with phenomenal properties defined as above would also lead (via the possibility of spectrum inversion) to the denial of ‘color incompatibilities.’

What, then, are colors? A plausible thought at this point for one who finds Shoemaker's line convincing would be that colors are not phenomenal properties, but dispositions to have certain phenomenal properties. Perhaps, for example, redness is the disposition to cause experiences with quale R in certain conditions to individuals whose visual system is set up in such and such a way... Alternatively, redness could be some physical property underlying such dispositions.

4.1.7 Shoemaker's view on the relationship between phenomenal feel and representational content

Shoemaker's view of perception, then, has us distinguish the following aspects of perceptual experience:

1. The event of perceiving.
2. Qualia, which are properties of such events which determine the phenomenal properties the event represents objects as having.
3. Phenomenal properties, which are properties of objects that our experiences represent them as having.
4. Colors, which are properties of objects constructed out of phenomenal properties.

What, we should ask, do we have access to in introspection? Shoemaker's plausible answer is: we have access to the phenomenal properties our experience represents objects as having.

An intuitive argument for this claim: when we try to focus on the phenomenal feel of an experience, 'there seems nothing for one's attention to focus on except the content of the experience.' (301)

An additional evolutionary argument for this claim.

This puts us in a position to raise the question of self-knowledge with respect to perceptual experience in a more precise way. Recall Shoemaker's functionalist explanation of self-knowledge at the end of Lecture II. If we are to apply that model to the case of self-knowledge of perceptual experiences, it seems that the relevant question is: can we give a functionalist analysis of the phenomenal properties represented in experience?

(We will be skipping §III of Lecture III; the gist of it is that Shoemaker thinks that the model he applies to perceptual experience is also applicable to the case of bodily sensations. So, if this is right, the explanation of self-knowledge should be the same in the two cases.)

4.2 *Phenomenal properties and self-knowledge*

4.2.1 *The case against functionalism about perceptual representation*

Shoemaker's discussion of perception started out with the consideration of a puzzle arising from a kind of dissonance between the scientific picture of the world and the properties our experience represents the world as having. Shoemaker claims to have gone some distance toward resolving this dissonance:

“Placing the phenomenal character in the representational content, and pointing out that the representational content includes the ascription of properties things have in virtue of their relations to experiences . . . removes the appearance of a disparity between the phenomenal character we are presented with and our beliefs about the objective nature of our external environment.” (304)

But this does not, as he notes, resolve the problem of the relationship between the phenomenal character we are presented with and our beliefs about the objective nature of our own minds. In particular, we may feel at a loss to explain how physical facts about our brains could give rise to experiences which present us with the phenomenal properties with which we are familiar.

But we can also present arguments which seem to show that a materialist functionalist view of the mind is false: (i) the knowledge argument; (ii) the explanatory gap; (iii) the inverted spectrum. Of these, (iii) seems the most convincing as an argument against functionalism.

The inverted spectrum argument. Why it seems to count against functionalism. Doubts about the argument.

Zombies and 'absent qualia' as a further argument against functionalism.

4.2.2 *Shoemaker's defense of functionalism*

Shoemaker puts his reply as follows:

“I think that the best response to these worries is to show that the existence of these apparent disparities between manifest and scientific image is just what an acceptable physicalist theory should lead us to expect. What I shall argue is that a broadly functionalist view, combined with physicalism, predicts that we will be presented in experiences with a phenomena character which is in a certain sense irreducible to its functional and physical underpinnings.” (304)

Shoemaker's defense of this claim may be broken down into the following steps (pp. 305-7):

- i. Any creature capable of perception will distinguish among certain stimuli and not among others.
- ii. These distinctions will typically not correspond to fundamental physical distinctions, but rather will be at least partially determined by the biological needs of the creature in question. (Recall S's view of phenomenal properties: they are properties of the external world, but are also 'anthropocentric.')
- iii. The states by which a creature represents these stimuli will also be more or less phenomenally similar to each other.
- iv. These similarity relationships will be functionally definable.
- v. Qualia are properties of these experiential states which determine these similarity relationships.
- vi. So differences and similarity amongst qualia will also be functionally definable.
- vii. But even if this were so, individual qualia would not be functionally definable were the quality space in question symmetrical.

Shoemaker glosses the conclusion of this line of thought as follows:

“If this last is right, we have an account, compatible with a materialist view of mind, which predicts that the phenomenal character of experience will be ineffable, and so irreducible, in a very strong sense. But that would be frosting on the cake ...” (307-8)

But it is difficult to see how this is supposed to help Shoemaker's position. He admits that we have introspective access to the phenomenal properties represented in experience. He further admits that individual functional properties are not functionally definable. Doesn't it just follow from this that his functionalist explanation of self-knowledge fails for perceptual states and bodily sensations? Why should it matter whether or not this failure is derivable from the supposition that functionalism + materialism is true?