

What do we learn from experience?

Wolfgang Schwarz <wo@umsu.de>

Draft, April 10, 2006

Abstract. I try to figure out what information we acquire when we observe our environment. First I argue that it is just phenomenal information. Then I notice that this is incompatible with physicalism. So I give an alternative, and even better answer that turns out to be incompatible with dualism.

1 Learning by looking

There is a bucket in the corner of the room. You wonder if it's empty, so you go and look inside: there's water in it.

Our visual experiences provide us with information about our environment. They put us in a position to rule out possibilities we couldn't rule out before, like the possibility that the bucket is empty. Let's figure out exactly which possibilities your experience rules out.

Suppose you are undecided whether the bucket contains water (H_2O) or some other substance XYZ which looks exactly like water, but has a different chemical composition. Will you be able to rule out one of these two possibilities by looking inside the bucket?

No. By assumption, H_2O and XYZ look exactly the same, so you cannot decide whether something is H_2O or XYZ just by looking. Hence the information your perception provides to you is not that you are confronting a bucket full of water, even if you *are* confronting a bucket full of water. Perhaps it is that you are confronting a bucket full of watery stuff?

Well, suppose you are in a lab where very sophisticated holograms of liquids are produced. You are unsure whether the bucket in the corner contains some real liquid or some machinery that makes it look exactly as if it contained the liquid. Can you rule out one of these two possibilities by looking inside the bucket?

Again, no. In ordinary cases, we can tell whether something is a hologram or the real thing because holograms differ visually from real things in many ways. But if the possibility under consideration is that we confront a hologram that looks *exactly* like the real thing, there is no hope of finding out which it is just by looking. Hence the information you acquire by looking inside the bucket is not that you are confronting a

bucket full of liquid, or full of watery stuff. Not even that you confront any bucket at all, for you can't rule out that the bucket itself is a sophisticated hologram.

Of course in ordinary situations we mostly ignore the possibility that we might be surrounded by sophisticated holograms. We assume (defeasibly) that if something looks exactly like a bucket, then it really is a bucket. All I claim is that it is not our perception itself that rules out the skeptical alternative. To show this, we have to consider situations where the alternative is not already ruled out even before the perception arrives. We have to consider a situation where in the outset, the hologram hypothesis is a genuine open possibility. Only so can we find out if the perception is apt to rule out that possibility. And clearly it is not. In general, our perceptual experiences do not rule out – indeed, they increase the probability of – any possibility wherein we have just this kind of experience.

Suppose you are convinced that your glasses are set up so that whenever you look at a zebra, they stimulate your retina in the way looking at a bucket full of water normally would, and vice versa (i.e., whenever you look at a bucket full of water, they cause the stimulation pattern of looking at a zebra). Then you will take your visual experience of a bucket full of water as evidence that you are confronting a zebra. And if you are unsure whether you're wearing the deceptive glasses while confronting a zebra or whether you're wearing normal glasses while confronting a bucket full of water, your visual experience of the bucket will not help you to decide the matter.

2 Pure phenomenal content

So the information you acquire from seeing a bucket full of water is not that you're confronting a bucket full of water, or watery stuff, or that you are not confronting a zebra. If relevant alternative possibilities are brought up in which you have the same visual experience even though there is no bucket full of water in front of you, but rather a bucket full of XYZ, or a hologram, or a zebra, your visual experience cannot help you to rule them out.

Our perceptual experiences do not allow us to exclude any possibility wherein we have just those experiences. On the other hand, they do allow us to exclude any possibility wherein we do *not* have these experiences. That is why, when the only relevant possibility of having a visual bucket experience is that we are confronting a bucket, our experience allows us to conclude that we actually are confronting a bucket.

So we've found an answer to the question what information our perceptual experiences convey to us: they allow us to exclude *all and only* those possibilities wherein we do not have this kind of experience. The information an experience conveys to us is just that we are having this kind of experience. It locates us in the set of possible individuals who make this kind of experience. Let's call this the *pure phenomenal content* of our perceptions.

Pure phenomenal content by itself is pretty useless if we want to gain information about our environment. Our sense organs would be like a signal that lights up every now and then merely to indicate that it currently lights up. How do we get the information

we care about from pure phenomenal information?

The obvious answer is: by inferences. We assume that in ordinary cases, we have *this* kind of perception only when our environment is such-and-such. With these assumptions, we have already ruled out all possibilities wherein we have the bucket perception even though there is no bucket around. Given those prior assumptions then, the fact that the bucket perception occurs allows us to conclude that there really is a bucket before us.

I'm not, of course, talking about conscious inferences here. I'm not postulating psychological mechanisms at all, not even unconscious ones. When I say that we 'infer' that there is a bucket from our experience together with the assumption that in the present situation, bucket experiences go with bucket encounters, all I mean is that without that assumption, we would no longer take the experience as evidence that there is a bucket. (Presumably most of these assumptions are innate, for it is hard to see how we could have learned correlations between experiences and external states of affairs without already taking some experiences as evidence for external states of affairs.)

On this view, there is nothing in the experience itself that relates it to buckets. This aspect of its content is rather projected onto it by assumptions we make about when experiences of its kind normally occur. Hence what we assume is not that when we experience things to be such-and-such, they usually are such-and-such. This would presuppose that the experience already bears information about external affairs independently of those assumptions, so that we can refer to it as an experience of 'things being such-and-such'. In our assumptions, experiences are not individuated by their external content, but, presumably, by their phenomenal character: if we have an experience with phenomenal character *X*, then usually things are such-and-such.

So here's the idea: The information we acquire when we make a perception with a certain phenomenal character is just that we make a perception with this phenomenal character. From this we infer all kinds of facts about our surroundings by means of background assumptions that link perceptual experiences – individuated by their phenomenal character – to external states of affairs.

I like this view. Before I reject it, let me briefly defend it against three possible objections.

Objection: Surely it is preposterous to suggest that what we *see* is that we are having a certain visual experience, and not, say, that there is a bucket in the corner.

Reply: Granted. But the semantics of "seeing" is not my topic. I'm not trying to find some kind of proposition *P* such that "*X* sees (or perceives or whatever) that *S*" is true iff the embedded sentence *S* expresses the proposition *P*. For what it's worth, I believe that nothing could possibly satisfy that job description. Be that as it may be, this isn't what I'm up to. I'm trying to figure out what information we acquire when we make a perception.

Objection: This pure phenomenal content, despite its name, doesn't do justice to the phenomenology of perception: there is something about the raw feel of seeing a bucket that relates it to buckets; indeed, to get the phenomenology right, the bucket itself should be *part* of the experienced content.

Reply: Again, this is not my topic. I'm not trying to do justice to the phenomenology

of perception, if only because I don't understand the rules of this game. I'm just trying to figure out what information we acquire when we make a perception.

I don't believe that there are no other interesting conceptions of perceptual content, and maybe on one of these other conceptions, the bucket itself is part of the content. But whatever this other content is good for, it is not the information we conditionalize on when we make a perception.

Objection: The current proposal ignores the important distinction between veridical and non-veridical perceptions: it treats hallucinating a bucket on a par with wrongly inferring that there is one from the fact that somebody says so. But there is a difference: in case of illusion and delusion, the representational error is not merely in our inferences, but in the perception itself.

Reply: A quick (and by now probably unsurprising) reply would be that it isn't my aim to capture intuitions about veridical and non-veridical perceptions. But there's more to be said. First, I'm not sure the alleged distinction is really that clear: when I take a sophisticated bucket hologram to be a bucket, does the error lie in my perception or in my inferences? More importantly, whether we judge an experience to be veridical or deceptive depends, to some extent, on our knowledge and previous experience: after seeing many sticks in water, they don't really look bent any more, just as objects moving away from us don't really look as if they get smaller, even though perhaps they do so for very young children. The current proposal neatly explains this: an illusion is a case where our assumptions about which experiences go with which external states of affairs are false; these assumptions are affected by learning and previous experiences.

3 Physicalism

I'm inclined to believe that our perceptual experiences are brain states. Your experience of seeing the bucket in the corner is probably some pattern of activity X in your visual cortex. So if the information your perception conveys to you is that you have this kind of experience, then it conveys to you that you are in brain state X .

But this seems wrong. If you don't know that your experience is brain state X , or if you are undecided whether it's brain state X or brain state Y (or ectoplasma state Z), merely looking at the bucket will not help you to find out.

Consider the space of all possible individuals. Your perception allows you to locate yourself in a subset of this space. (For it allows you to exclude certain possibilities about yourself.) What distinguishes the individuals in this subset from the others? They all have an experience that is phenomenally indistinguishable from your current experience. But if there are no fundamental phenomenal properties, this cannot be the end of the story: whatever distinguishes the individuals in the subset from all the others must be something that can, in principle, be expressed in non-phenomenal terms. If, say, the phenomenal character of your experience is a certain physical property, then the relevant set of individuals is the set of individuals whose current state has that physical property. You could rule out that your state has some other physical property than that. But it seems that you can't, whatever precisely that property is.

So it appears that my proposal is incompatible with physicalism. Since I liked the proposal, I think this is bad news for physicalism.

The situation is reminiscent of the knowledge argument: the information our perceptions convey to us appears to be *phenomenal information* – the kind of information Mary lacks before her first encounter with colours. Interestingly, the ability hypothesis, that works reasonably well against the knowledge argument, doesn't work at all in the present case: we can hardly say that our perceptions do not convey to us any information at all, but merely teach us new abilities. Abilities are not the kind of thing you can draw inferences from or conditionalize on.

The other popular response to the knowledge argument – that Mary learns old facts in a new way – carries over better: the idea would be that the information our perceptions convey to us really is information about physical properties of our brain, but that this content is somehow presented to us in an intransparent way. By looking at the bucket, then, you can indeed rule out that you are in brain state Y , but you can rule out this possibility only under some phenomenal modes of presentation, not under the mode of presentation “I'm in brain state Y ”.

The idea that information can be available to us under different modes of presentation is an important insight. If you know the rules of chess, you have all the information you need to decide what is the best move in a given stage of the game. But it is a difficult task to bring that information into the right form. Strictly speaking, you don't acquire new information when you're told what is the best move: the possibilities you can thereby exclude are possibilities that were already excluded by things you knew beforehand. That is why you could, in principle, have figured it out all by yourself, just by thinking about what what you already knew. In general, if what you lack is merely a certain way of presenting (to yourself) information you already possess, then logical competence in principle suffices to get you there. If, by contrast, what you lack is genuine information, no amount of thinking about what you already know will help you.

Now consider the two sentences “there is water in the bucket” and “there is H_2O in the bucket”. Let's assume that they (or their present utterances) are true in exactly the same possible worlds. Thus in some sense they ‘have the same content’ or ‘express the same information’. But unless you already know that water is H_2O , their different impact on your beliefs is not merely a difference in modes of presentation. For suppose you are convinced that water is actually XYZ. Then by accepting “there is water in the bucket”, you will locate yourself in a region of logical space where there is XYZ in the bucket. This is clearly not the same region as the one where there is H_2O in the bucket. So in this case, we don't have the same region under different modes of presentation, but completely different regions to begin with. Another sign of this is that, unless you know that water is H_2O , no amount of logical competence will help you to find out that there is H_2O in the bucket from being told that there is water in the bucket. What you lack is real information, real exclusion of possibilities.

Let's return to phenomenal content. Can the difference between knowing that you have an experience with such-and-such phenomenal character and knowing that you are in brain state X be explained as a difference in mode of presentation? Unfortunately

not – at least not unless you already know that the relevant experience is brain state X : suppose you (wrongly) believe that it is brain state Y ; then having the experience will not cause you to locate yourself in the set of individuals with brain state X , no matter under what mode of presentation. On the contrary, it will cause you to locate yourself in the set of individuals with brain state Y . To find out that you're among the individuals with state X , what you need is not logical competence, but real information.

So the other response to knowledge argument also doesn't work in the present case (just as it doesn't work in the original case).

4 Pure non-phenomenal content

Let's start over again. Maybe I was too quick when I identified perceptual information with pure phenomenal information. Why doesn't your visual experience of the bucket allow you to exclude all possibilities where there is no bucket around? Because there are *skeptical scenarios* in which you have this very experience even though there is no bucket. So then what your experience should tell you is that *either* there is a bucket before you *or* you're in some kind of skeptical scenario. This is more than just phenomenal information.

Consider a possible situation where you have the bucket experience. It could be an ordinary situation where there is a bucket before you that causes your experience in the ordinary way. Or it could be a situation where there is a sophisticated hologram before you, or one in which you wear the deceptive glasses while confronting a zebra; or it could be a situation where somebody has rewired your brain to make you have bucket experiences when confronting a zebra. You know how to construct scenarios like these. Put them all aside. Are there still other cases, cases where you have the experience, there is no bucket before you, and yet you are not in some typical skeptical scenario? Could you have the bucket experience while looking at a zebra while there is simply no explanation of the error – no deceptive mirrors or glasses, no rewired brains, no drugs, nothing? Is it (epistemically) possible that zebras cause bucket experiences in just the ordinary way buckets actually cause bucket experiences?

This idea strikes me as odd. I've certainly never seen it brought up as a skeptical scenario. In skeptical scenarios there is always at least a hint at an explanation for the subjects's error, as when the perception is called a "halluzination", indicating that the causal link between environment and perception is somehow broken.

So perhaps your experience really tells you that either there is a bucket in front of you or you're in some ordinary skeptical scenario – not because it allows you to rule out other possibilities in which you have the experience while there is no bucket and yet you're not in an ordinary skeptical scenario, but simply because *there are no such possibilities* in the first place.

One can argue for a similar connection between bucket experiences and bucket beliefs: it is easy to imagine that someone regards her bucket experiences as evidence for the presence of zebras – she might think she wears deceptive glasses, or that her brain has been rewired. But can we coherently imagine that someone – or better, everyone – takes their bucket experiences to be completely natural signs for zebras in the way we

take them to be signs for buckets? Is it just a contingent fact about our world that we associate this kind of experience with buckets, and not zebras or houses or sandstorms?

If not, your experience tells you even more: it also tells you that people having this experience normally tend to believe that there is a bucket nearby.

This, then, is my second proposal: the information your bucket perception conveys to you is that you are in a state that is normally (barring skeptical possibilities) caused by bucket-like objects, and that normally causes people to believe that there is a bucket-like object nearby. I'll call this the *pure non-phenomenal content* of your experience.

I do not claim that it is intuitively obvious that this is the real content of our perceptions. First, I am still not interested in the semantics of "seeing" or our everyday notion of perceptual content, and I'm still not trying to get the phenomenology right. My proposal is only a proposal about what information we acquire by making a perception. Second, even as such, I do not think that the proposal is intuitively obvious or even plausible. At least to me, it sounds rather weird. But it *follows* from something I'm at least inclined to believe: that there is something incoherent about the idea that 1) zebras might cause bucket experiences in just the ordinary way we believe buckets cause bucket experiences (without any skeptical mechanisms in place), and also about the idea that 2) ordinary people might take their bucket experiences as natural signs for zebras just like we take them as natural signs for buckets.

It is tempting to suggest that pure non-phenomenal content is at least not *all* our perceptions convey: you *also* find out that you're having an experience with a certain phenomenal character. But this temptation must be resisted: when I say that it's incoherent to suppose that bucket experiences might be normally caused by zebras, what I mean is that it is incoherent that experiences with the phenomenal character of bucket experiences are normally caused by zebras. So the phenomenal character is not logically independent of the causal roles that make up non-phenomenal content, and so information about phenomenal character is not *additional* information.

Is it really impossible for experiences with the phenomenal character of bucket experiences to be normally caused by zebras? Couldn't anything cause anything? Consider a world like ours except that, by a fundamental law of nature, zebras cause just those retinal stimulations buckets cause in our world, say, by bending their surrounding spacetime to redirect the photons they reflect. Isn't that a world where zebras cause experiences with the phenomenal character bucket experiences have in our world? I'm not sure, but I don't really need to deny it. What I have to deny is only that this is an epistemic possibility for us: expressions like "the phenomenal character bucket experiences have in our world" might be rigid designators for some (physical) property of states that are normally caused by buckets in the actual world. If so, then we are mostly ignorant about the phenomenal character of our experiences: our experiences do not allow us to locate ourselves in the set of possible individuals whose present state has such-and-such phenomenal character. This also sounds odd, but maybe it is better than denying the possibility of the strange zebra world. As I said, I'm not sure.

Two final remarks. First, I've argued above that if pure phenomenal content is what our perceptions convey to us, then physicalism is false. Arguably, the converse is also true: if physicalism is false, then pure phenomenal content is all our perceptions convey

to us. For then phenomenal properties are logically independent of all other properties: *any* way the world might be is (epistemically) compossible with *any* phenomenal experiences. There are, for instance, (epistemically) possible worlds where everybody's phenomenal experiences are running two hours late: when people there eat breakfast and read the newspaper in the morning, they have the experience of still sleeping; later, when they have start working, they have the experience of eating breakfast and reading the newspaper, etc. If there are such (epistemic) possibilities, then clearly our perceptual experiences do not allow us to rule them out. Hence if physicalism is false, our experiences do not allow us to conclusively infer *anything at all* about the external world.

Second, both pure phenomenal content and pure non-phenomenal content are species of narrow content: content that never differs between intrinsic duplicates within the same world. If the question is what information we acquire when we perceive our environment, wide content is a *very* implausible answer. Wide content might fare better in other projects, in this 'getting the phenomenology' project perhaps, or in the semantics of "seeing": "Oscar sees that there is water in the bucket; Twosker sees that there is twater in the bucket". But the claim that we can find out whether some stuff is H₂O or XYZ just by looking, despite the fact that the two alternatives look exactly the same, sounds utterly incredible to me – though I realize that this contradicts many theories of mental content on which we do acquire the belief that there is H₂O (and hence not XYZ) around us just by looking.

Acknowledgements

I owe a lot here to discussions with David Chalmers and Robert G. Williams.