

# THE EPISTEMOLOGY OF PAIN\*

Fred Dretske

Many people think of pain and other bodily sensations (tickles, itches, nausea) as feelings one is necessarily conscious of. Some think there can be pains one doesn't feel, pains one is (for a certain interval) not conscious of ("I was so distracted I forgot about my headache."), but others agree with Thomas Reid (1785: 1, 1, 12) and Saul Kripke (1980: 151) that unfelt pains are like invisible rainbows: they can't exist. If you are so distracted you aren't aware of your headache, then, for that period of time, you are not in pain. Your head doesn't hurt. It doesn't ache. For these people (I'm one of them) you can't be in pain without feeling it, and feeling it requires awareness of it. The pain (we have learned) may not bother you, it may lack the normal affective (as opposed to sensory) dimension of pain, but it can't be completely insensible. That wouldn't be pain.

Whether or not we are necessarily conscious of our own pains is clearly relevant to the epistemology of pain. If there are--or could be--pains of which we are not aware, then there obviously is an epistemological problem about pain. How do we know we are not always in pain without being aware of it? Pain becomes something like cancer. You can have it without being aware of it. Not only that. You can not be in pain and mistakenly think you are. As a result of misleading external signs, you mistakenly think you are in pain that you haven't yet felt. The epistemology of pain begins, surprisingly, to look like the epistemology of ordinary physical affairs. So much for Descartes and the alleged first-person authority about our own mental life.

One could go in this direction, but I prefer not to. I prefer to simplify things and minimize problems. I therefore focus on pains (if there are any other kind) of which we are aware, pains, as I will say, that really hurt. For those who conceive of pain as something of which one is necessarily aware, then, I am concerned, simply, with pain itself: when you are in pain, how do you know you are? For those who think there are, or could be, pains one does not feel, I am concerned with a subset of pains--those one is aware of, the ones that really hurt. When it hurts, how do you know it does?

If this is the topic, what is the problem? If we are talking about things of which we are necessarily aware, the topic is things that, when they exist, we know they exist. There may be a question about how we know we are in pain, but that we know it is assumed at the outset. Isn't it?

No it isn't. To suppose it is is to confuse awareness of things<sup>1</sup> with awareness of facts. One can be aware of an armadillo, a thing, without being aware of the fact that it is an armadillo, without knowing or believing it is an armadillo--without, in fact, knowing what an armadillo is. There is a sensory, a phenomenal, form of awareness--seeing or in some way perceptually experiencing an armadillo--and a conceptual form of awareness--knowing or believing that it (what you are experiencing) is an armadillo. We use the word "awareness" (or "consciousness") for both. We are aware of objects (and their properties) on the one hand, and we are aware that certain things are so on the other. If one fails to distinguish these two forms of awareness, awareness of x with awareness that it is x, one will mistakenly infer that simply being in pain (requiring, as I am assuming, awareness of the pain) requires awareness of the fact that one is in pain and, therefore, knowledge. Not so. I am assuming that if it really hurts, you must feel the pain, yes, and feeling the pain is

awareness of it, but this is the kind of awareness (thing-awareness) one can have without fact-awareness of what one is aware of--that it is pain--or that one is aware of it. Chickens and maybe even fish, we may suppose, feel pain, but in supposing this we needn't suppose that these animals have the concept PAIN. We needn't suppose they understand what pain is sufficiently well to believe (hence know, hence be aware) that they are in pain. They are aware of their pain, yes. It really hurts. That is why they squawk, squirm and wiggle. That is why they exhibit behavior symptomatic of pain. But this does not mean they believe or know that they are in pain. It is the chicken's feeling pain, its awareness of its pain, not its belief or awareness that it is in pain that explains its behavior. The same is true of human infants.<sup>2</sup> They cry when they are hungry not because they think (much less know) they are hungry. They cry because they are hungry or, if you prefer, because they feel hungry, but they can be and feel hungry without knowing it is hunger they feel or that they are feeling it.

So I do not begin by assuming that if it hurts, you know it does. In fact, as a general rule, this is false. Maybe we, adult human beings, always know when it hurts, but chickens and fish (probably) don't. Human infants probably don't either. I'm asking, instead, why we, adult human beings, always seem to know it and if we really do know it, how we know it. What is it about pain—unlike, say, cancer—that confers epistemological authority on its possessors?

When I see a pencil, the pencil doesn't depend on my awareness of it. Its existence, and its existence as a pencil, doesn't depend on my seeing it. When I stop looking at it, the pencil continues to exist in much the way it did when I saw it. Pencils are indifferent to my attentions. There is the pencil, the physical object of awareness, on the one hand, and there

is my awareness of it, a mental act of awareness, on the other. Remove awareness of the pencil, this relational or extrinsic property, from the pencil, and one is left with the pencil, an unchanged object of awareness.

This can't be the way it is with pain since pain, at least the sort I am concentrating on here, unlike pencils, is something one is necessarily aware of. Remove the act of awareness from this object of awareness and this object ceases to be pain. It stops hurting.

So in this respect pains are unlike pencils. Unlike a pencil, the stabbing sensation in your lower back cannot continue to exist, at least not as pain, when you cease to be aware of it. If it continues to exist at all, it continues to exist as something else. When the marital relation is removed (by divorce, say) from a husband, the relation that makes him a husband, you are left with a man, a man who is no longer a husband. If you remove the awareness relation from a stabbing sensation in your lower back, the relation without which it doesn't hurt, what are you left with? A stabbing event (?) in your lower back that doesn't hurt? That isn't painful? What sort of thing could this be?<sup>3</sup>

To understand the sort of thing it might be think about something I will call a crock.<sup>4</sup> A crock (I stipulate) is a rock you are visually aware of. It is a rock you see. A crock is a rock that stands in this perceptual relation to you. Remove that relation, as you do when you close your eyes, and the crock ceases to be a crock. It remains a rock, but not a crock. A crock is like a husband, a person whose existence as a husband, but not as a man, depends on the existence of a certain extrinsic relationship. Cocks are like that. They look just like rocks. They have all the same intrinsic (non-relational) properties of rocks. They look like rocks because they are rocks, a special kind of rock, to be sure (one you are aware of), but a rock nonetheless.

Should we think of pains like crocks? When you cease to be aware of it, does your pain cease to exist as pain, but continue to exist as something else, something that has all the same intrinsic (and other relational) properties of pain but which requires your awareness of it to recover its status as pain in the way a rock has all the same intrinsic (and most relational) properties of a crock and requires only your awareness of it to regain its status as a crock? Is there something--let us call it protopain--that stands to pain the way rocks stand to crocks? Is a stabbing pain in your lower back merely a stabbing (?) protopain in your lower back that you happen to be aware of? Under anesthesia, is there still protopain in your lower back, something with all the same intrinsic properties you were aware of when feeling pain, but something that, thanks to the anesthetic, no longer hurts because you are no longer aware of it?<sup>5</sup>

If we understand pain on this model--the crock model--we have a problem. How serious the problem is depends on how much of a problem it is to know one is aware of something. To appreciate the problem, or at least threat of a problem, think about how you might go about identifying crocks. When you see a crock, it is visually indistinguishable from an ordinary rock. Rocks and crocks look alike. They have the same intrinsic, the same observable properties. They differ only in one of their relational properties. They are, in this respect, like identical twins. This means that when you see a crock, there is nothing in what you are aware of, nothing in what you see, that tells you it is a crock you see and not just a plain old rock. So how do you figure out whether the rock you see is a crock and not just a rock?

I expect you to say: I know the crocks I see are crocks because I know they are rocks (I can see this much) and I know I see them. So I know they are crocks. Assuming there is

no problem about recognizing rocks, this will work as long as there is no problem in knowing you see them. But how do you find out that you see, that you are visually aware of, the rock? There is, as already noted, nothing in what you see (the crock) that tells you that you see it. The crock would be exactly the same in all observable respects if you didn't see it--if it wasn't a crock. Just as husbands differ from men who are not husbands in having certain hidden (to direct observation) qualities, qualities one cannot observe by examining the husband, crocks differ from rocks in having a certain hidden quality, one that can't be observed by examining the crock. When you observe a crock, the relational property of being observed by you is not itself observed by you. Just as you must look elsewhere (marriage certificates, etc.) to find out whether the man you see is a husband, you must look elsewhere to find out whether the rock you see is a crock.

But where does one look? If one can't look at the crock to tell whether it is a crock, where does one look? Inward? Is introspection the answer? You look at the rock to see whether it is a rock, but you look inward, at yourself, so to speak, to find out whether it is that special kind of rock we are calling a crock. If we understand pain in your lower back on the crock model--as a condition (in the lower back?) you are aware of--we won't be able to say how you know you have lower back pain until we understand how you find out that this condition in your lower back is not just a condition in your lower back, but a condition in your back that really hurts, a back condition that you are aware of.

This is beginning to sound awfully strange. The reason it sounds so strange is that although awareness (at least awareness of objects) is a genuine relation between a person and an object, it is, we keep being told by philosophers, an epistemologically transparent<sup>6</sup> or self-intimating relation. When S is aware of something, S knows automatically, without the need

for evidence, reasons, or justification, that he is aware of it. S can be married to someone and not realize he is (maybe he has amnesia or he was drunk when he got married), but he can't be aware of something and not know he is. If this were so, there would be no problem about knowing the rocks you see are crocks since you can easily (let us pretend) see that they are rocks and you would know immediately, without need for additional evidence, in virtue of the transparency of awareness, that you see (are visually aware) of them. So anything one sees to be a rock is known, without further ado, to be a crock. The fact that makes a rock a crock--the fact that one is aware of it--is a transparent, self-intimating fact for the person who is aware of it. That is why there is no epistemological problem about pain over and above the familiar problem of distinguishing it from nearby (but not quite painful) sensations--e.g. aggressive itches. That is why we don't have a problem distinguishing pain from protopain. Whatever it is, exactly, we are aware of when we are in pain, we always know, in virtue of the transparency of awareness, that we are aware of it. When we have a pain in our back, therefore, we always know it is real pain and not just protopain.

This nifty solution to our problem doesn't work, but it comes pretty close. It doesn't work because awareness is not transparent or self-intimating in this way. Animals and very young children are aware of things, but, lacking an understanding of what awareness is, they don't realize, they don't know, they are aware of things. A chicken is visually aware of rocks and other chickens without knowing it is. That is why animals and young children--even if they know what rocks are (they probably don't even know this much)--don't know the rocks they see are crocks. They don't know they see them. They don't know they are aware of them. If awareness is a transparent, self-intimating relation, it is so for only a select class of people. It certainly isn't so for children and animals.

So if we are going to appeal to the transparency of awareness in the epistemology of pain, we must be careful to restrict its transparency to those who understand what awareness is, to those capable of holding beliefs and making judgments about their own (and, of course, others) awareness of things.<sup>7</sup> We need, that is, a principle something like:

T: If S understands what awareness is (i.e., is capable of holding beliefs and making judgments to the effect that she is aware of things), and S is aware of x, S knows she is aware of x.

This sounds plausible enough, but we have to be careful here with the variable “x”. What, exactly, does it mean to say that S knows she is aware of x? If x is a rock, must S know she is aware of a rock? Clearly not. S can see a rock and not know it is a rock and, therefore, not know she is visually aware of a rock. She thinks, mistakenly, it is a piece of cardboard. She might not even know it is a physical object. She thinks she is hallucinating. So what, exactly, does principle T tell us S knows about the x she is aware of?

Nothing. Except that she is aware of it. Awareness of objects makes these objects available to the person who is aware of them as objects of de re belief, as things (a this or a that) he or she can have beliefs about. Since, however, none of these additional beliefs you have about x need be true for you to have them, you needn’t know anything about x other than that you are aware of it. To illustrate, consider the following example. S sees six rocks on a shelf. She sees them long enough and clearly enough to see all six. When S looks away for a moment, another rock is added. When S looks back, she, once again, observes the rocks long enough and clearly enough to see all seven. She doesn’t, however, notice the difference. She doesn’t realize there is an additional rock on the shelf. She sees--and is, therefore, aware of--an additional rock on the shelf, but she doesn’t know she is. S is aware of something (an additional rock) she doesn’t know she is aware of.



Does this possibility show, contrary to T, that one can be aware of an object and not know it? No. It only shows that one can be aware of something additional without knowing one is aware of it under the description "something additional." Maybe, though, one knows one is aware of the additional rock under the description "the leftmost rock" or, simply, as "one of the rocks I see" or, perhaps (if she doesn't know it is a rock), as "one of the things I see." If all seven rocks are really seen the second time, why not say the perceiver knows she is aware of each and every rock she sees. She just doesn't know they are rocks, how many there are, or that there are more of them this time than last time. But she does know, of each and every rock she sees, that she is aware of it.

If we accept this way of understanding "S knows she is aware of x," there may still be a problem about the intended reference of "x" in our formulation of transparency principle T. Suppose S hallucinates a talking rabbit with the conviction that she really sees and hears a talking rabbit. S mistakenly thinks she is aware of a talking white rabbit. She isn't. There are no white rabbits, let alone talking white rabbits, in S's vicinity. What, then, is S aware of? More puzzling still (if we assume she is aware of something), what is it that (according to T) S knows she is aware of? Is there something, something she can (perhaps mentally) pick out or refer to as that, that she knows she is aware of? If so, what is it? Is it something in her head? A mental image? If so, does this image talk? Or does it merely appear to be talking? Does it have long ears? Or only appear to have long ears? Is S, then, aware of something that has (or appears to have) long white ears and talks (or sounds as though it is talking) like Bugs Bunny?

Knowing what lies ahead on this road (viz., sense-data) many philosophers think the best way to understand hallucinations (dreams, etc.) is that in such experiences one is

not aware of an object at all--certainly nothing that is white, rabbit-shaped, and talks like Bugs Bunny. Nor is one aware of something that only appears to have these properties. It only seems as though one is. Although there appears to be an object having these qualities, there actually is no object, certainly nothing in one's head, that has or even appears to have<sup>8</sup> the qualities one experiences something as having.<sup>9</sup>

This way of analyzing hallucination, however, seems to threaten T. S thinks she is aware of something--a talking rabbit, in fact--but she isn't. She isn't aware of anything. So while hallucinating, S's belief that she is aware of something is false. This seems to show that, sometimes at least, S can't tell the difference between being aware of something and not being aware of something. Why, then, suppose, as T directs, that S always knows when she is aware of something?

What we need to understand in order to sidestep this kind of objection to T is that the "x" S knows she is aware of needs to be interpreted liberally. It needn't be a physical object. It needn't be a mental object (a sense-datum) either. It can be a property or a set of (appropriately "bound" together) properties the subject experiences something as having. In hallucinating a talking white rabbit S is conscious of various sensory qualities: colors, shapes, tones, movements, orientations, and textures. These are qualities S experiences (perceptually represents) something as having, qualities S is conscious of in having this hallucinatory experience. According to the intended interpretation of T, it is such qualities, not some putative object that has (or appears to have) these qualities, that S is aware of and (in accordance with T) knows she is aware of. The difference between an hallucination of a talking white rabbit and a veridical perception of one isn't--or needn't be--the phenomenal (sensory) qualities one is aware of. These can be exactly the same. The experiences can be

subjectively indistinguishable. In one case one is aware of something that has the qualities, in the other case not. But in both cases the subject is aware of, and in accordance with the intended interpretation of T, knows she is aware of, the qualities that make the experiences that kind of experience.

So much by way of propping up T. What we are left with may appear contrived and suspiciously *ad hoc*. Nonetheless, it or something close to it seems to do the job. It explains how one can, without additional epistemic effort (beyond what it takes to identify rocks), know that crocks are crocks. More significantly for present purposes, it also explains why someone who understands what awareness is, someone who is cognitively developed enough to think she is in pain, can't be aware of protopain (the cluster of qualities she is aware of when in pain) without knowing she is aware of them and, therefore, without knowing she is in pain.

Is T true? If it is, why is it true? What is it about awareness, or perhaps the concept of awareness, or perhaps the having of this concept, that yields these striking epistemological benefits?

The fact that, according to T, S must not only be aware of x, but also understand what awareness is (an understanding animals and infants lack) in order to know--*gratis*, as it were--she is aware of x tells us something important. It tells us the knowledge isn't constitutive of awareness. It tells us that awareness of x doesn't consist of knowing one is aware of it. The truth of T--if indeed it is true--isn't what Fricker (1998) calls an Artifact of Grammar. There are some mental relations we bear to objects in which it seems plausible to say knowledge is a component of the relationship. Memory of persons, places, and things is like that. For S to remember her cousin (an object), S needn't remember that he is her cousin (maybe she never knew this), but she must at least remember (hence know) some facts about her cousin—that,

he looked so-and-so, for instance, or that he wore a baseball cap.<sup>10</sup> Memory of persons and things, it seems reasonable to say, consists in the retention (and, therefore, possession) of such knowledge about them. Awareness of objects and persons, though, isn't like that. You--or, if not you, then chickens and children--can be aware objects without knowing they are. So the knowledge attributed in T is not the result of some trivial, semantic fact about what it means to be aware of something. It isn't like the necessity of knowing something about the people you remember. If T is true, it is true for some other, some deeper, reason.<sup>11</sup>

Perhaps, though, it goes the other way around. Although a (lower level) awareness of something (a rock) doesn't have a (higher level) belief that one is aware of it (the rock) as a constituent, maybe the higher level belief that you are aware of it (a belief animals and young children lack) has awareness as a constituent. Maybe, that is, awareness of x is a relation that holds between x and whoever thinks it holds.<sup>12</sup> If this were so, then a belief that you are aware of something would always be true. According to some theories of knowledge, then, such a belief would always count as knowledge. Whoever thinks they are aware of something knows they are because thinking it is so makes it so. So they can't be wrong. So they know.

This possibility would be worth exploring if it really explained what we are trying to explain--viz., why, when we are aware of something, we know we are. But it doesn't. The fact (if it were a fact) that awareness of something is, somehow, a constituent of the (higher order) belief that one is aware of something would explain why the higher order belief, if we have it, is always true--why, if we believe we are aware of something, we are. But it would not explain what we are trying to explain, the converse: why, if we are aware of something, we always believe (thus, know) we are. The proffered explanation leaves

open the possibility that, when we are aware of something, we seldom, if ever, believe we are and, therefore, the possibility that, when aware of something, we seldom, if ever, know we are.

So if T is, somehow, an artifact of grammar, a truth vouchsafed in virtue of the concept of awareness, it is not in virtue of the belief being a constituent of the awareness or *vice versa*. If you always know when you are in pain, you know it for reasons other than that the belief (that you are in pain) is a constituent of the pain or the pain is a part of the belief. The pain and the belief that you are in pain are distinct existences. The problem is to understand why then, despite their distinctness, they are, for those who understand what pain is, apparently inseparable.

Chris Peacocke (1992; and earlier, Gareth Evans, 1982: 206) provides a way of understanding our possession of concepts in which the truth of T can be understood as somehow (to use Fricker's language) an artifact of grammar without supposing that it is to be understood in terms of the knowledge being a constituent of the awareness or *vice versa*. Concepts not only have what Peacocke (1992: 29) calls attribution conditions-- conditions that must be satisfied for the concept to be correctly attributed to something. They also have possession conditions, conditions that must be satisfied for one to have the concept. To have a perceptual (what Peacocke, 1992: 7, calls a sensational) concept for the color red, for instance, he says that a person must, given normal circumstances, be able to tell, just by looking, that something is or isn't red. She must know the concept applies, or doesn't apply, to the things she sees. Possessing the concept RED requires this cognitive, this recognitional, ability. Those who lack this ability do not have the concept RED.<sup>13</sup>

Adapting this idea to the case of awareness, it might be supposed that a comparable cognitive ability is part of the possession conditions for AWARENESS. Although (as the case of animals and young children indicate) knowledge isn't part of the attribution (truth) conditions for awareness (S can be aware of something and not know she is), an ability to tell, in your own case, authoritatively, that you are aware of something may be a possession condition for this concept. You don't really have the concept, you don't really understand what it means to be aware of something, if you can't tell, when you are aware of something, that you are aware of it. This is why, in the antecedent of T, an understanding (of what awareness is) is required. Awareness is transparent for those who possess the concept of awareness because its transparency, the ability to tell, straight-off, that one is aware of something, is a requirement for thinking one is aware of something. If you can think you are aware of something, then, when you are aware of something, you know you are in that special authoritative way required for possession of the concept

This strikes me as a plausible--if not the only possible--explanation of why T is true. Regrettably, though, it doesn't take us very far. It is, in fact, simply a restatement of what we were hoping to explain--viz., T: that those who understand what awareness is know, in virtue of having this understanding, when they are aware of something. It does not tell us what we were hoping to find out—the *source* of the epistemological ability required for possession of this concept. If, to have the concept AWARENESS, I have to know I'm aware of everything I'm aware of, how do I acquire this concept? What is it that gives me the infallible (or, if not infallible, then near-infallible) powers needed to possess this concept and, thereby, a capacity to *think* I'm aware of something?<sup>14</sup>

One doesn't explain infallible--or, if not infallible, then authoritative--application of a concept by saying that infallibility (or authoritativeness) in its application is a condition for possessing the concept. That might be so, but that doesn't help us understand where the authority comes from. It merely transforms an epistemological question—how do we know something is X--into a developmental question--how do we manage to believe that something is X. If, in our skeptical moods, we are suspicious about infallibility or first-person authority, then requiring it as a necessary condition for possessing a concept does nothing to alleviate our skepticism. It merely displaces the skepticism to a question about whether we in fact have the concept—whether we ever, in fact, believe we are aware of something. It is like trying to solve an epistemological problem about knowing you are married by imposing infallibility in believing you are married as a requirement for having the concept MARRIED. You can do this, I suppose, but all you really manage to achieve by this maneuver is a kind of conditional infallibility: if you think you are married, you know you are. But the old question remains in a modified form: do you think you are? Given the beefed-up requirements on possessing the concept MARRIED, it now becomes very hard—for skeptics impossible--to think you are married.

We began by asking how one knows one is in pain. Since pain, at least the kind of pain we are here concerned with, is a feeling one is necessarily aware of (it doesn't hurt if you are not aware of it), this led us to ask how one knows one is aware of something. We concluded, tentatively, that the kind of reliability in telling you are aware of something required for knowledge (that you are aware of something) must be a precondition for possessing the concept AWARENESS, a precondition, therefore, for thinking you are aware of something. That explains why those who think they are aware of something know

they are. By making reliability of judgment a possession condition for the concept of awareness, we have transformed our epistemological problem into a developmental problem, a problem about how one comes to possess the concept of AWARENESS or, indeed, any concept (like pain) that requires awareness. How do we manage to *think* we are in pain?

This doesn't seem like much progress. If we don't understand how we can make reliable judgments on topic T, it doesn't help to be told that reliability is necessary for making judgments about topic T.

But though it isn't much progress, it is, I think, some progress. If nothing else it reminds us that the solution to some of our epistemological problems, problems about how we know that a so-and-so exists, await a better understanding of exactly what it is we think when we think a so-and-so exists, a better understanding of what our concept SO-AND-SO is. It reminds us that questions about how we know P may sometimes be best approached by asking how we manage to believe P.<sup>15</sup> This is especially so when the topic is consciousness and, in particular, pain. Understanding how we know it hurts may require a better understanding of what, exactly, it is we think (and how we manage to think it) when we think it hurts.



**REFERENCES**

- Armstrong, D. 1961. *Perception and the Physical World*. London; Routledge and Kegan Paul.
- Armstrong, D. 1962. *Bodily Sensations*. London; Routledge and Kegan Paul
- Bilgrami, A. 1998. Self-knowledge and resentment. In *Knowing Our Own Minds*, Crispin Wright, Barry Smith, and Cynthia Macdonald, eds. Oxford: The Clarendon Press, 207-242.
- Burge, T. 1985. Authoritative self-knowledge and perceptual individualism. In R. Grimm and D. Merrill, eds. *Contents of Thought* (Tucson, AZ; University of Arizona Press), 86-98.
- Burge, T. 1988. Individualism and self-knowledge. *Journal of Philosophy*, 85: 649-63.
- Chalmers, D. 1996. *The Conscious Mind*. New York: Oxford University Press
- Dretske, F. 1983. The epistemology of belief," *Synthese*, 55.1: 3-19.
- Dretske, F. 1995. *Naturalizing the Mind*. Cambridge, MA; MIT Press.
- Dretske, F. 1999. The mind's awareness of itself. *Philosophical Studies*, 1-22. Reprinted in Dretske, F. *Perception, Knowledge, and Belief* (2000), Cambridge University Press, 158-177.
- Dretske, F. 2003. How do you know you are not a zombie? in *Privileged Access and First-Person Authority*, edited by Brie Gertler, Ashgate Publishing Co. Also published in Portuguese, Conference on Mind and Action III, Lisbon, Portugal, 2001. Edited by João Sàáguand
- Evans, G. 1982. *The Varieties of Reference*. Oxford; Oxford University Press.
- Fricke, E. 1998. Self-knowledge: special access versus artifact of grammar--a dichotomy rejected. In *Knowing Our Own Minds*, Crispin Wright, Barry Smith, and Cynthia Macdonald, eds. Oxford: The Clarendon Press, 155-206.
- Gallois, A. 1996. *The World Without, The Mind Within*. Cambridge; Cambridge University Press.
- Kripke, S. 1980. *Naming and Necessity*. First published in 1972. Cambridge: Harvard University Press.
- Lycan, W. 1996. *Consciousness and Experience*. Cambridge, MA; MIT Press.

Peacocke, C. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press: A Bradford Book.

Pitcher, G. 1971. *A Theory of Perception*. Princeton; Princeton University Press.

Putnam, H. 1981. Brains in a vat. *Truth and History*. Cambridge, Cambridge University Press: 1-21.

Reid, T. 1785. *Essays On The Intellectual Powers*.

Shoemaker, S. 1996. Self-knowledge and "inner sense." In *The First Person Perspective and Other Essays*. Cambridge; Cambridge University Press: 224-245.

Siewert, C. 1998. *The Significance of Consciousness*. Princeton: Princeton University Press.

Siewert, C. 2003. Self-knowledge and rationality: Shoemaker on self-blindness. In *Privileged Access: Philosophical Accounts of Self-Knowledge*, Brie Gertler, ed. (Ashgate Press: Aldershot, England): 131-146.

Tye, M. 1995. *Ten Problems of Consciousness*. Cambridge, MA; MIT Press

Wright, C. 1998. Self-knowledge: the Wittgensteinian legacy. In *Knowing Our Own Minds*, Wright, C. Smith B. C., and Macdonald, C., eds. Oxford University Press, pp. 13-45.

## ENDNOTES

---

\*A version of this paper was first given as a keynote address at the 7<sup>th</sup> annual Inland Northwest Philosophy Conference (INPC) on Knowledge and Skepticism, Washington State University and the University of Idaho, April 30-May 2, 2004. I am grateful to the audience there for helpful and constructive discussion.

<sup>1</sup> By "things" I mean spatio-temporal particulars. This includes, besides ordinary objects (houses, trees, and armadillos), such things as events (births, deaths, sunsets), processes (digestion, growth), conditions (the mess in his room), and states (e.g., Tom's being married). Events occur at a time and in or at a place (the place is usually the place of the objects to which the event occurs). Likewise for states, conditions, processes, and activities although these are usually said to persist for a time, not to occur at a time. So if one doesn't like talking about pains as objects and prefers to think of them as events (conditions, activities, processes) in the nervous system, that is fine. They are still things in my sense of this word. For more on property-awareness and object-awareness as opposed to fact-awareness see Dretske 1999.

<sup>2</sup> It is for this reason that I cannot accept Shoemaker's (1996) arguments for the "transparency" of pain--the idea that pain is necessarily accompanied by knowledge that one is in pain. Even if it is true (as I'm willing to grant) that pains (at least the pains of which we are aware) necessarily motivate certain aversive behaviors, I think it is an over-intellectualization of this fact to always explain the pain-feeler's behavior in terms of a desire to be rid of her pain (a desire that, according to Shoemaker, implies a belief that one is in pain). I agree with Siewert (2003: 136-37) that the aversive or motivational aspect of pain needn't be described in terms of conceptually articulated beliefs (that you have it) and desires (to be rid of it). In the case of animals (and young children), it seems to me implausible to give it this gloss. Maybe you and I go to the medicine chest because of what we desire (to lessen the pain) and think, (that the pain pills are there), but I doubt whether this is the right way to explain why an animal licks its wound or an infant cries when poked with a pin.

<sup>3</sup> Daniel Stoljar and Manuel Garcia-Carpintero (on two separate occasions) have asked me why I think there is anything remaining when I subtract awareness from pain. Why isn't subtracting awareness from pain more like subtracting oddness from the number 3 rather than subtracting married from a husband? My reason for thinking so is that when we are in pain there is something we are aware of that is ontologically distinct from the pain itself--e.g., the location, duration, and intensity (i.e., the properties) of the pain. These are among the qualities that give pain its distinctive phenomenal character, the qualities that make one pain different from another. They are the qualities that make a splitting headache so different from a throbbing toothache. Take away awareness of these qualities and, unlike the number 3 without oddness, one is left with something--the qualities one was aware of.

<sup>4</sup> I introduced crocks as an expository device in Dretske 2003.

<sup>5</sup> This way of thinking about pain (and other bodily sensations) is one version of the perceptual model of pain (Armstrong 1961, 1962; Dretske 1995; Lycan 1996; Pitcher 1971; Tye 1995) according to which pain is to be identified with a perceived bodily condition (injury, stress, etc.). Under anesthesia the bodily injury, the object you are aware of when in pain, still exists, but since it is no longer being perceived it, it no longer hurts. It isn't pain. I say this is "one version" of a perceptual theory because a perceptual model of pain can identify pain not with the perceived object (bodily damage when it is being perceived), but the act of perceiving this object, not the bodily damage of which you are aware, but your awareness of this bodily damage. In the latter case, unlike the former, one does not perceive, one is not actually aware of, pain. When in pain, one is aware of the bodily injury, not the pain itself (which is one's awareness of the bodily injury). I do not here consider theories of this latter sort. As I said at the outset, I am concerned with the epistemology of pain (sensations in general) where these are understood to be things of which one is conscious. If you aren't (or needn't be) aware of pain, there are much greater epistemological problems about pain than the ones I am discussing here.

<sup>6</sup> Transparency as here understood should be carefully distinguished from another use of the term in which it refers to the alleged failure (or at least difficulty) in becoming introspectively aware of perceptual experience and its properties. In trying to become aware of the properties of one's perceptual experience, one only seems to be made aware of the properties of the objects that the experience is an experience of, the things one sees,

---

hears, smells and tastes. One, as it were, “sees through” the experience (hence, transparency) to what the experience is an experience of.

<sup>7</sup> For careful formulations along these lines see C. Wright (1998), Fricker (1998), and (for "self-intimating") Shoemaker (1996). Chalmers 1996, pp. 196-97 describes awareness as an epistemologically special relation in something like this sense, and Siewert 1998 (19-20, 39, 172) suggests that mere awareness of things (or failure to be aware of things) gives one first person warrant for believing one is (or is not) aware of them. I take it that even animals and children have the warrant. They just don't have the (warranted) belief.

<sup>8</sup> Nor appears to have these qualities because to suppose that S was aware of something that merely appeared to have these qualities would be to introduce an appearance-reality distinction for mental images. What is it (a part of the brain?) that appears to be a talking white rabbit? This seems like a philosophically disastrous road to follow.

<sup>9</sup> If this sounds paradoxical, compare: it can appear to S as though there is a fly in the ointment without there being a fly who appears to be in the ointment.

<sup>10</sup> I do not argue for this. I'm not even sure it is true. I use it simply as a more or less plausible example of a relation we bear to objects that has, as a constituent, factual knowledge of that object.

<sup>11</sup> This is why functionalism (about the mental) is of no help in explaining why T is true. Even if awareness (of an object) is a functional state, one defined by its causal role, its role cannot include the causing of belief that one is aware of something.

<sup>12</sup> This echoes a Burgian (Burge 1985, 1988) thesis about belief--that the higher order belief that we believe p embodies, as a constituent, the lower order belief (that p) that we believe we have. This echo is pretty faint though. The major difference is that awareness of an object is not (like a belief) an intentional state. It is a genuine relation between a conscious being, S, and whatever it is she is aware of. It may be that believing you believe p is, among other things, to believe p, but why should believing you are aware of something be, among other things, awareness of something? Can you make yourself stand in this relation to something merely by thinking you do? It is for this reason that Bilgrami (1998) thinks that the constituency thesis (as he calls it) is only plausible for intentional states like belief (desire, etc.) that have propositional "objects".

<sup>13</sup> It isn't clear to me what concept they have—or even whether they have a concept—if they do not have this ability at the requisite (presumably high) level of reliability but are, nonetheless, more often right than wrong in describing something as red. If they don't have the concept RED, what are they saying? What, if anything, are they thinking? Nothing? I take this to be the problem David Chalmers was raising in the discussion at the INPC conference. I ignore the problem here for the sake of seeing how far we can get in the epistemology of pain by requiring a level of reliability (of the sort needed to know) in the capacity to believe.

<sup>14</sup> As I understand him, this is basically the same point Gallois (1996) is making against Peacocke's account of why (or, perhaps, how) we (those of us who have the concept of belief) are justified in believing that we believe the things we do. See, in particular, Gallois 1996, p. 56-60.

<sup>15</sup> In Dretske (1983) I argued that the condition (relating to justification, evidence, or information) required to promote a belief that x is F into knowledge that x is F is also operative in our coming to believe that x is F (in acquiring the concept F). Roughly, if something's being F isn't the sort of thing you can know, it isn't the sort of thing you can believe either. This, of course, is the same conclusion Putnam (1981) reaches by considering brains in a vat.