THE DUALIST

Spring 1998 • Volume V • Number 1

Department of Philosophy
Stanford University

1500 - 1000 B.C.E. They are generally thought of as divine in nature and as unauthored.

2. The notion of sudden enlightenment or instant realization of the ultimate is not found directly in Sankara's ideology. It is, however, raised as an objection, expressed primarily from the Buddhist perspective.
3. I do not want to suggest that univocity will "take over" for agapeic astonishment and erotic perplexity. The movement between the excess of *Agape* and the lack of *eros* continues as the force motivating this process through the stages, not just that which begins the process.
4. *Lila* or the playfulness of the Lord is a complicated subject; volumes have been written about it. For our purposes, this playfulness, as well as the superimpositions it creates, is not ultimately real and in no way can be designated as characteristic of *nirguna*-Brahman.

REFERENCES

Desmond, William. *Being and the Between*. Albany: State University of New York Press, 1995.

Plato. *Symposium*. Trans. Alexander Nehams and Paul Woodruff. Indianapolis: Hackett Publishing Company, 1989.

Ramachan, Anantanand. *The Limits of Scripture: Vivekananda's Reinterpretation of the Vedas*. Honolulu: University of Hawaii Press, 1994.

Sri Sankaracarya. *Bhagavadgita with the Commentary of Sankaracarya*. Trans. Swami Gambhirananda. Calcutta: Advaita Ashrama, 1995.

——— (*BSB*). *Brahma-Sutra-Bhasya*. Trans. Swami Gambhirananda. Calcutta: Advaita Ashrama, 1996.

Upanisads. Trans. Patrick Olivelle. New York: Oxford University Press, 1996.

AN INTERVIEW WITH FRED DRETSKE

MAY 1998
Stanford University

THANKS TO Ned Block, Tyler Burge, Dan Dennett, Jerry Fodor, Keith Lehrer, and Ernie Sosa for their questions. I had fun thinking about them. Since the questions were mainly about what I now think, I had the experience—rare in philosophy—of being reasonably certain my answers were correct.

Ernest Sosa, Brown University
Fred, you are considered an early advocate of an approach to knowledge and skepticism that features the following elements:

> *(a) Appeal to subjunctive conditionals linking* P *and belief of* P *in understanding knowledge of* P
>
> *(b) (Consequent upon that first element* (a)*) denial of the closure of knowledge under flawless deduction*
>
> *(c) Use of a distinction between alternatives that are "relevant" and those that are not, in responding to skepticism.*

These elements have proved attractive and are increasingly prominent (though not always in combination). How much of the approach would you still advocate, and can you highlight what is distinctive of your own version?

I still accept all three elements. The use of a subjunctive conditional linking *P* and the belief that *P* (or the experience on which that belief is based) began in *Seeing and Knowing* (1969). In that book I took a thoroughly externalist (as it is now called) view of perceptual knowledge (denying, for instance, the necessity of justification). I also explored ideas that led to my later rejection in *Epistemic Operators* (1970) of closure for knowledge. For example, I denied that to see that *x* is *F*, it is necessary to see (or even be able to see) that *x* is *G*, where *x*'s being *G* is a known logical consequence of *x*'s being *F*.

The third element that Ernie mentions—that some alternatives (to knowing that *P*) are not the sort of thing one need be able to rule out on

evidential grounds in order to know that *P*—is pretty much a consequence of the first two doctrines. If knowledge really is a matter of getting oneself "connected" to the facts in the right way—by means of a relationship that can be expressed by a subjunctive conditional—then *being connected* to *Q* doesn't follow from *being connected* to *P* even when *P* logically requires *Q*. Irrelevant (to knowing *P*) alternatives are just those to which one isn't connected in being connected to *P*. So the pieces were all present in the 1969 book. What came later was a better understanding of how these pieces fit together.

What is distinctive about my own approach? Perhaps only the priority I give to perception. If we know anything about the world, then, surely, perception (seeing, hearing, smelling that *x* is *F*) is a way of knowing. If this is so, then, since (I assume) perception is also an important source of knowledge for children and animals (if Rover can't *see that* the stick landed here, not there, why does he run here, not there, to fetch it?) one is driven (I believe) to externalism about knowledge. I think it is hard to formulate plausible internalist conditions on knowledge when they have to fit Rover knowing where the stick is.

Ned Block, New York University

It appears that on your view, zombiehood is hereditary. Have I misinterpreted you? Suppose that George Steinbrenner's grandparents were all swamp-people, people formed by chance from particles from a swamp. Unlikely, yes, but possible, also yes. And suppose that no significant evolution was involved in the process by which Steinbrenner's grandparents and parents met each other and produced their offspring. On your theory, it would seem that Steinbrenner is therefore a zombie. Growing up in the Bronx playing stickball did not give his experiences any representational content. But when an outraged fan kicked him in the shins, did he not feel pain?

I don't know whether many readers will understand Ned's motivations for asking about the hereditariness of zombiehood. So let me supply a bit of background. I have a theory about the character of mental states (the qualities that, in Tom Nagel's phrase, define what it is like to be *S*) that makes these qualities depend on the history of *S*. Identical creatures with different histories can have different thoughts and experiences and, if the histories are different *enough*, one can have experiences while the other (a physically identical creature) lacks experiences entirely (i.e., is a zombie). All this sounds crazy to many (most?) people. Ned's question is his way of dramatizing how counterintuitive this doctrine is. Surely George feels pain when kicked in the shins.

Let me begin by saying that I agree with Ned about one thing. This

is all terribly counterintuitive. But then most people—on first hearing it—also find the idea that physically identical creatures could have different *beliefs* counterintuitive. Nonetheless, Putnam's and Burge's examples (involving Twin Earth) have convinced many of these people that thoughts, purposes, intentions, desires, and so on (propositional attitudes) do not supervene on the physical constitution of the agent. Despite the same present physical constitution, creatures with different causal histories could have different beliefs and desires. This being so, complaints about the counterintuitive character of a theory about phenomenal experience that makes *it* depend on history sound to me a little hollow coming from theorists (I assume Ned is one of them) who adopt (externalist) theories about judgment. After all, some theorists (Dennett, Armstrong, et al.) conceive of experience as just a special form of (implicit, covert) judgment. Hence, if judgments do not supervene on the current physical constitution of the judger, neither do experiences.

Ned has himself given examples that demonstrate to many people's (certainly to my) satisfaction that functionalism doesn't capture the qualitative aspects of experience. Functionally identical beings need not have the same kind of experience—indeed, one could be a zombie (having no experience at all) while the other did. So Ned is prepared to say that George Steinbrenner could have a functional duplicate who was not in pain—someone (looking just like Steinbrenner) hopping around crying in what seems like pain after being hit with a stick who was not really in pain. Frankly, I don't see a lot of difference between this result and the consequences of my own theory that Ned finds objectionable.

In answer to Ned's questions about the heritability of zombiehood, I make two points. No, zombiehood is not necessarily heritable since (on this view of experience) the descendants of zombies might not be zombies. Their internal states could acquire indicator functions that make them representations. This, I assume, is what happened in the evolution of consciousness: we (conscious beings) had (remote) ancestors who were not conscious. In the example Ned gives, however, we are asked to assume no significant selection has occurred, and, thus, that nothing in George has a biological function that the corresponding organ or state did not have in his parents. In this case, of course, George is a zombie if his ancestors were.

Jerry Fodor, Rutgers University

You have fixed more of my philosophical fractures than any of my other gurus, for which I am grateful forever. But there's one thing I still don't get, even after all these years. (I think it's pretty much what Chomsky asked Skinner.) If the content of a thought is determined by its history,

how can there ever be novel thoughts?

Jerry asks about the source of novelty. If history determines what we think (the content of thought), how can we have thoughts we've never had before. I think I'm going to give an answer to this question that is pretty dull. It is exactly what both Jerry and I (and maybe everyone else) already believes. I'll give it anyway.

What history determines are the concepts out of which beliefs are formed, not the beliefs themselves. What makes brain state b_1 mean *cat*, brain state b_2 mean *milk*, and brain state b_3 signify the relation *x likes y* is history. What makes b_1, b_2, and b_3 combine *today* in *S*'s head in a way that expresses the proposition that cats like milk (a thought *S* has never had before) isn't determined by history (at least not the same history that gives the individual concepts their content), but factors about current psychology and/or stimulus conditions. If everything is determined, then *S* is determined to think that cats like milk and so, in this sense, nothing is really novel. It is all predictable. But the novelty of thought—even in a deterministic world—arises from the fact that the (historical) conditions that give the elements of thought their meaning do not themselves determine which arrangements of those elements (the thoughts) will subsequently occur.

Keith Lehrer, University of Arizona
The account that you offer of perceptual knowledge and the content of belief is an exceptional achievement, but it leaves one wondering how to extend the account to theoretical knowledge and the content of beliefs about theoretical entities—that is, ones that one cannot perceive.

There are really two questions that Keith is asking—one about knowledge:

> (1) How do I extend my account of perceptual knowledge to theoretical knowledge?

And the other is about the content of belief:

> (2) Given my account of how we acquire the concept *cow* and *tree* (things we can perceive), how do I extend this to acquiring concepts for things we cannot perceive—e.g., *bacterium* or *neutrino*?

I don't think I need *extend* the account of perceptual knowledge to the domain of theoretical knowledge since the account of knowledge is the same in both cases: we know that *P* when we are caused to believe that *P* by the information that *P*. Sometimes (as with cows) the objects we come to know about are visible (audible, etc.); in other cases (bacteria) they aren't. In the latter case, instruments (microscopes, etc.) are needed. The seeing, hearing, etc. is indirect. But whether the information is delivered directly or indirectly, the knowledge is the same: information caused belief.

The answer to the second question is harder. I don't really know the answer. When asked this question I typically give the usual empiricist reply accompanied by a lot of hand-waving. Concepts for non-perceptual objects are (somehow!) manufactured out of the observational concepts we already have plus whatever syntactical machinery is available (don't ask me where this comes from) to combine concepts in various ways. I develop the concept of *object too small to be seen* (which applies to objects I cannot see) out of concepts like *object, small, seen,* etc.—concepts that apply to objects I can see.

Dan Dennett, Tufts University
In "Differences that Make No Difference" (Philosophical Topics, *1994, the special issue on my work), you defend a version of non-epistemic seeing which, you noted, challenged my account in* Consciousness Explained. *I responded with a challenge of my own (same issue, "Get Real" in the section entitled "Dretske's Blind Spot"), based on some of the recent research on change blindness (which I had predicted in* Consciousness Explained*): your non-epistemic seeing is not interestingly different from what happens on the inert wall of an unoccupied camera obscura. You need to add* some *sort of uptake, I claimed, to justify saying that details are "in consciousness" in the sense that they are not also on the wall in the camera obscura. Do you still think that there is a theoretically interesting sense of non-epistemic seeing that does not just collapse into something like "ephemeral cortical activity was the result of retinas being irradiated by light from" (or alternatively into something that is a non-ordinary variety of epistemic seeing after all)?*

Dan Dennett's and my disagreements (we also have many areas of agreement) go back nearly thirty years, so we aren't likely to settle them now. Each time we have a go at each other, I come away thinking I've finally put my point in such a convincing way that Dan will *have to* agree. Apparently he feels the same way about his ripostes. We keep disappointing each other.

He asks whether I still think there is a theoretically interesting sense of non-epistemic seeing that doesn't just collapse into ephemeral cortical activity (or some version of epistemic seeing). Well, yes I do,

but a lot depends on what one finds "theoretically interesting."

Dan thinks the richness of the outside world, in all its ravishing detail, does not enter our experience of that world. I think a lot of this richness, this detail, does enter our experience. Not all of it, of course, but more of it than we notice or respond to. Maybe conditions can be arranged in which we *would* respond to (non-epistemically) perceived differences and qualities, but in most—perhaps all—casual encounters (think of glancing at a crowded room of people or a shelf full of books) we see more than we can know, more than it is possible to *cognitively* process. This is why I think a small child sees (exactly) five fingers before she knows how to count, before she knows the difference between five and four (or six). This is why I think seeing a Ø is so much different (subjectively speaking) from believing (judging, knowing) that there is a Ø.

Dan is worried that my view makes the claim that *S* (non-epistemically) sees a Ø collapse into (something like) the claim that the Ø is visible. Not quite, but close. Given my bad eyesight, there are some visible things that I cannot see—even when I'm looking right at them. They are too small for *me* to see, but that doesn't mean they are invisible. Others can see them.

The results of the experiments Dan describes are absolutely neutral between his view and mine. The only way they can be interpreted as supporting Dan's claim that *S* cannot see *X* unless there is some cognitive uptake (some identification or recognition of *X*) is if cognitive uptake is taken as criterial for seeing *X*. But this, of course, is exactly what—given our differences—he cannot do. I am, in fact, quite happy with the results of these experiments. They show what I would expect to be shown, what I take to be obvious—viz., that a lot of what we see we don't notice. They do not show that we do not see things unless we notice them. They only show that *if* we see these things in a non-epistemic way (as I claim we probably do), epistemic seeing does not always accompany non-epistemic seeing. That, though, was my claim from the very beginning.

Tyler Burge, UCLA

Earlier in your career, you wrote a fine book entitled Seeing and Knowing. *That book was primarily concerned with the epistemology of perception. Do you see yourself as having shifted more to issues in metaphysics and philosophy of mind since writing that book? Or do you see your work on intentional content as continuous with the concerns of that book? What changes in doctrine have you made since writing* Seeing and Knowing*?*

You have said several times that you agree with others in thinking that your representationalist account of consciousness has certain implausible features, but that you persist with the account because it is the only way you see of remaining a materialist—or naturalist—about the mind. Could you explain what you take materialism or naturalism to be, and why you regard it as a doctrine to die for—or at least to go to great lengths to defend?

Tyler Burge asks whether my concerns have shifted since I wrote *Seeing and Knowing* (1969). Yes, they have shifted—and in just the way Tyler suggests. My research took a turn away from epistemology and toward the philosophy of mind and metaphysics when I started worrying more about what made something a belief than what made a belief knowledge, more about what made something an experience than how experiences figures in perception. Aside from these shifts of interest, though, there hasn't been much change in my epistemological views (see my answer to Ernie Sosa). I still maintain the strong externalism I defended in *Seeing and Knowing* and for (roughly) the reasons I gave then.

I do think my account of conscious experience (in *Naturalizing the Mind*, 1995) has strongly counterintuitive consequences. I persist with it anyway because, as Tyler notes, it seems to me to be the only way of closing (what has come to be called) the explanatory gap in a naturalistically acceptable way. How else could one explain how various states of gray matter can make the person in whom that gray matter resides experience (be aware of) the color orange, movement, or middle C when there need be no sounds, colors, or movements inside (or even—during hallucination—outside) the head of the person undergoing these experiences? These properties are not properties of the experience (since experiences are in the head and nothing in the head has—or needs to have—these properties). Neither need they be properties of anything outside the head (e.g., hallucination). They must, it seems, be properties the experience in some way represents things as having, properties the experience stands for or is about. They must, in other words, be intentional properties of the experience. When I ask myself what, in the natural (non-mental) world I am familiar with that can be about properties even when nothing has those properties I find only representational artifacts like instruments, language, and signs—things that come to be about what they are about in virtue of having a certain kind of history. This history (usually bound up with the intentions and purposes of designers, makers, and users) gives a device (a pointer on a gauge, a symbol, a sign) a meaning that is independent of what properties it has and (during misrepresentation) independent of what properties sur-

rounding objects (those it is about) have. This meaning is tied up with what the device has the function of indicating. That is why speedometers can be about (i.e., represent the vehicle in which they are placed as going) 60 mph when nothing (including the speedometer) has this property.

So I am driven to thinking of experiences as being something like the states of a measuring device. The difference is that the functions—unlike those of a measuring instrument—do not come from us. They come from natural selection and learning. The fact that this is the only way I can imagine some state of a physical device (I take the brain to be a physical device) being about (representing) properties that nothing has may only show that my powers of imagination are weak. Maybe, in a hundred years, we will know more and understand better how brains can perform this marvelous trick. Maybe. But I have to philosophize with what I'm given, and today I don't see any alternatives to the one I've given. And if I'm going to do philosophy at all, if I'm going to try to figure out how the mind could work in terms I already understand (and this is all I mean by *naturalism*), I have to work with what I'm given.

THE DUALIST would like to thank Fred Dretske for taking the time to participate in this interview. Professor Dretske is the Bella and Eloise Mabury Knapp Professor of Philosophy and Department Chair at Stanford University. He is the author of the books Seeing and Knowing *(1969),* Knowledge and the Flow of Information *(1981),* Explaining Behavior *(1988), and* Naturalizing The Mind *(1995). His research has been in epistemology—in particular, the area of visual perception—and the philosophy of mind. His recent work centers on the nature of intentional action and perceptual experience.*