

Handout #10: Dretske

1. A Pragmatic/Reductionistic Conception of Understanding the Understanding

“If you want to know what intelligence is, you need a recipe for creating it out of parts you already understand” (491).

Example: (a) To understand what a cake is you need to know how to make a cake. (b) To know how to bake a cake, you need to be able to make a cake from ingredients that do not include cakes among them.

Against Dennett: The same is true of minds: “Recipes for thought can’t have interpretive attitudes or explanatory stances among the ingredients—not even the attitudes and stances of others. That is like making candy out of candy—in this case, one person’s fudge out of another person’s caramels. You can do it, but you still won’t know what candy is” (492).

Question: What can we assume or adopt as “parts” if we’re to construct a mind in such a way as to demonstrate our knowledge of how a mind works?

Dennett’s Homuncular Functionalism: Mental functions as complexes of simpler functions. If you want to know how a human understands language or sails a boat or proves a theorem, you need to decompose that task—accomplished, as it is, by the human using her mind as a whole or significant parts of it— into simpler tasks—accomplished by parts of her mind or lesser parts of it. By in turn decomposing those tasks into simpler ones we may arrive at a level at which the component tasks do not require any intelligence because they are so simple. Perhaps we will need to explain the mechanical details of how these simple tasks are executed. But either way we will have explained how people execute the complex tasks with which we began.

Dretske: We don’t need to be able to explain thought in terms of things that have no intentional or representational properties of any kind. “What we are trying to understand...is not intentionality, per se, but the mind. Thought may be intentional. But that isn’t the property we are seeking a recipe to understand. As long as the intentionality we use is not itself mental, then we are free to use intentionality in our recipe for making a mind” (492).

2. Natural Representations and their Relational Natures

“Some objects are constituted by their relationships to other objects.”

Analytic examples: Rembrandts, \$100 bills, cousins

“Some recipes require a special cook,” If a painting wasn’t created by Rembrandt it isn’t a Rembrandt even if it is qualitatively identical to one. If a bill wasn’t printed by the US mint, it isn’t a \$100 bill, even if it is qualitatively identical to one. “This is why you can’t build my cousin in your basement while my aunt and uncle can.”

An example: The compass. It tracks the north pole not the polar bears even though the polar bears only live at the north pole.

“Talk about what instruments and gauges indicate or measure creates the same kind of intensional (with an “s”) context as does talk about what a person knows or believes. Knowing or believing that that is the north pole is not the same as believing that that is the habitat of the polar bears even though the north pole is the habitat of the polar bears. . . . What one is describing with these intentional terms [e.g. indicating the north pole], is, therefore, in this sense, an intentional state of the instrument” (493).

Dretske’s Relational Realism About Mental Representation: “The mind, I think, is like that. . . . To say that the compass indicates the direction of the arctic pole is to say that the position of the pointer depends on the whereabouts of the pole. This dependency exists whether or not we know it exists, whether or not anyone ever exploits this fact to build and use compasses. The intentionality of the device is not like the intentionality of words and maps, borrowed or derived from the intentionality (purposes, attitudes, knowledge) of its users. The power of this instrument to indicate north to or for us may depend on our taking it to be a reliable indicator (and, this, on what we believe or know about it), but its *being* a reliable indicator does not depend on us” (493).

Dretske Naturalistic Attitude toward Aboutness: There is no need to naturalize the notion of aboutness or intentionality. It is already a naturalistic notion we can use to establish a physicalist understanding of the mind. “It exists wherever you find dark clouds, smoke, tree rings, shadows, tracks, lightning, flowing water, and countless other natural conditions that *indicate* something about how the rest of the world is constituted” (emphasis added, 493).

Questions: Do you share Dretske’s realist intuitions? Does x only supply the information that p if we use x to indicate that p or is there information in a world without minds? Can we use an unreduced or unanalyzed notion of information (so understood) to supply a truly reductive theory of mind or must we reduce the notion of information to something (e.g.) fully mechanical or biochemical?

Further questions: If naturalist theorists of mind can just help themselves to an unanalyzed or undefined notion of indication or representation or aboutness, how are they to distinguish thinking about something or having a belief about someone from tree rings, smoke and the like?

Let us say that a tree’s rings contain information about the age of that tree in a thin sense that does not entail any awareness or knowledge.

Question: Does mental representation distinguish itself from other forms of representation by the kinds of use to which information is put by the organism whose mind “contains” that information in the thin sense at issue?

3. Misrepresentation

According to Dretske, **what we need to build** (or know how to build) **if we are to understand the mind is something that** cannot only represent states of itself (as a tree’s rings represent its age) and represent things outside itself (as a smoke represents the fire or combustion at its source) but something which **can misrepresent things or incorrectly represent things** by indicating that things are such and such or so and so when they are not.

“For meaning or content, the what-it-is one thinks, is. . . independent of the truth of what one thinks” (494).

Fodor's Disjunction Problem: Lots of things cause an animal's nervous system to assume a structure R. "How can something's being R mean that something is F, when something's being F is only one of the things that cause R?"

Which of the many events involved in the generation of a neural token which we describe as an experience of a red surface is the object represented by that experience? Which of the many events involved in the generation of a thought we describe as a thought about Paris is the object represented by that thought?

Dretske: "Anything that can misrepresent something as being F is, of necessity, something whose meaning is independent of its causes, something that can mean cow even when it is caused by a horse on a dark night. It is therefore something whose meaning is less than the disjunction of conditions capable of causing it" (494).

The ability of compasses, thermometers and the like to misrepresent is explained by the functions for which we use them. We use the compass to indicate the North Pole so we can orient around that direction in order to navigate. That's why the compass is misrepresenting if it indicates N when pointed south. Dretske says we have given the device the job of indicating north, just as we give the mercury in a thermometer "the job" of indicating the temperature. "This is why thermometers can, while paper clips cannot, "say" something false about temperature" (494).

4. Natural Functions

Question: Can something (e.g. a state of an organism's mind) have a representational job in the sense in which the mercury in a thermometer has the job of indicating temperature) even if no one has given it that job?

Dretske: "If an information-carrying element in a system could somehow acquire the function of carrying information, and acquire this function in a way that that did not depend on our intentions, purposes and attitudes, then it would thereby acquire...the power to misrepresent the conditions it had the function of informing about" (495)

Phylogenetic or Evolved Function: Might the fitness enhancing function of a trait establish its natural representational function?

(a) Suppose X carries the information that p in the thin sense in which tree rings carry information about the age of trees.

(b) Suppose X is an adaptation insofar as X persists in a population because organisms with X had a fitness advantage over organisms without X and X therefore came to predominate in that population over time.

(c) Suppose, moreover, that X wouldn't have bestowed a fitness advantage on organisms who had it unless p were true, when X indicated as much.

If (a)-(c) obtain, is that enough to conclude that indicating that p is the function of X?

Dretske: "If the heart has the function of pumping blood, if that is why it is there, then by parity of reasoning, the senses...might have an information-providing function, the job of "telling" the animal in whom they occur what it needs to know in order to find food and mates and avoid

danger. If this were so, then, the natural function of sensory systems would be to provide information about an organism's optical, acoustic, and chemical surroundings" (495).

Ontogenetic or Learned Function: Suppose X tokens pursuit/avoidance of Y because this is necessary for the survival of organisms in which X occurs.

Question: Is this sufficient to conclude that X represents Y as good/bad for the organisms in question?

Dretske's Organic Relationalism: "If the only natural functions are those provided by evolutionary history and learning, then, no one is going to be able to build a thinker of thoughts, much less a mind, in a laboratory. This would be like building a heart, a real one, in your basement" (496).

Questions: Why shouldn't we instead say that things can get their functions in a number of ways including the intentions of a creator or the purposes of a possessor? On this understanding an artificial heart is a real heart because it is created to pump blood and/or it in fact pumps blood for the patient who has had that heart implanted?

5. Action Guiding Representations v. Reaction Guiding Representations

"I expect Darwin to help us understand why people blink, reflexively, when someone pokes a finger at their eye, but not why they (deliberately) wink at their friend. There are probably internal representations (of objects approaching the eye) involved in the blink reflex, representations that have an evolutionary origin, but these are not the sort of representations (beliefs, purposes, and intentions) at work in explaining why we wink at a friend or pack for a trip... Darwin is concerned with precisely those behaviors the explanatory mechanisms for which are genetically determined—precisely those behaviors that are not voluntary" (Dretske, 496).

So Dretske turns to ontogenetic functions and solves the disjunction and misrepresentation problem by invoking the idea of a "selectional process" by means of which a representation is recruited as a "determinant of system output."

Question: Does this idea allow us to distinguish the thought that, e.g., X is a cow, from the thought that X is a jersey cow, from the thought that X is an animal, when the thought in question is reliably triggered by and trained on Xs which are all three? Don't we still need a different response to the three classes? And what about classes with extensions that exactly co-vary? How can we use this apparatus to distinguish the thought that x is a creature with a heart from the thought that x is a creature with a kidney?

Dretske's answer: "If R, a COW indicator, gets its function of indicating cows by "exposure" to only Jersey cows, this does not mean that R means (has as its representational content) JERSEY COW. Whether it means COW, JERSEY COW, or perhaps simply ANIMAL, will depend..on the counterfactuals."

The counterfactual in question is whether R would still be triggered by a non-jersey cow. If so, Dretske might conclude that it is a cow indicator rather than a jersey cow indicator. And if R would still be triggered by something which is not a cow at all (e.g. a bull), then perhaps Dretske would conclude that R doesn't mean COW but instead means ANIMAL.

Questions: Why say that instead of saying that R means JERSEY COW and in the cases we've imagined the animal in whom R is triggered mistakenly represents a non-Jersey cow as a Jersey cow or mistakenly represents a sheep as a Jersey cow? Has Dretske really answered the disjunction problem?

Dretske's response: In cases of misperception there must be some impairment or risk of impairment in functioning. So if R means COW and is still regularly triggered by bulls, and the animal in whom R occurs tries to milk the bull, that animal won't succeed. So R will represent what triggers it as a means to successfully executing the animal's ends.

Question: Must fitness or success play this essential role in our analysis of representational content?