Nov 2nd Reading

For 'Speech Perception Contra Dretske'
(Justin Olaguer)

Dretske, F. 1997. Naturalizing the Mind. Cambridge: MIT Press.

Roughly, here is the central contention in a nutshell (124-126):

1 Seeming and Supervenience

The Representational Thesis is an externalist theory of the mind. It identifies mental facts with representational facts, and though representations are in the head, the facts that make them representations—and, therefore, the facts that make them mental—are outside the head. A state of the brain is an experience only if it represents the world in a certain way, and a state represents the world in this way, or so I have claimed, only if it has an appropriate information-carrying function. Since functions (whether systemic or acquired) have to do with the history of the states and systems having these functions, mental facts do not supervene on what is in the head. What is in heads A and B could be physically indistinguishable and yet, because these pieces of gray matter have had relevantly different histories, one is a representational system, the other is not; one is the seat of thought and experience, the other is not; one makes the person in whom it occurs aware of the world, the other does not.

For most philosophers these are unpalatable consequences. This is putting it charitably: for many they constitute a *reductio ad absurdum* of externalist theories of sensation. Even if one agrees with Putnam (1975) that "meanings ain't in the head" and finds Tyler Burge's (1979, 1982) examples in support of the social character of thought (a version of externalism) convincing—thus finds oneself accepting the idea that thought (or some thoughts [2]) have a content determined by external factors—few are willing to say the same about sensations. If two individuals are physically the same, then, surely, the fact that they live in different habitats, the fact that they have had different individual histories, or the fact that they evolved in different ways, is irrelevant to their current experiences. If one has a headache, so does the other. If one is having an auditory experience as of a piano being played (whether veridical or not), so is the other. Even if thought depends on the environment in which one exists (or existed), feelings do not. Theories that deny this are not *just* implausible, they are false. It is all very well for a theory to have materialistically attractive consequences. Other things

Not only does the Representational Thesis make *what* one thinks and feels externally determined, *that* one thinks and feels is likewise hostage to environmental and historical circumstances. Steve Stich's (1983) Replacement Argument dramatizes this fact. Imagine replacing a thinking-feeling being—you, say—with a duplicate, a ceptual experience. In §5 I describe the way it naturalizes intentionality and in §6 the way it provides a satisfying account of why, though experiences are in the head, one doesn't find them there. Other explanatory benefits are reserved for later chapters.

1 The Nature of Representation

The fundamental idea is that a system, S, represents a property, F, if and only if S has the function of indicating (providing information about) the F of a certain domain of objects.[1] The way S performs its function (when it performs it) is by occupying different states $s_1, s_2, \ldots s_n$ corresponding to the different determinate values $f_1, f_2, \ldots f_n$, of F. A speedometer (S) represents the speed (F) of a car. Its job, its function, is to indicate, provide information (to the driver) about, how fast the car is moving (F). When it is doing its job, its different states (pointer positions "24," "37," etc.) correspond to different car speeds (24 mph, 37 mph, etc.). Given the function of this instrument, each of its states is supposed to carry a different piece of information about the speed of the car: a registration of "37" is supposed to carry the information that the car is going 37 mph, "24" the information it is going 24 mph, and so on. The fact that the speedometer has a speed indicating function, and the fact that pointing at "37" means 37 mph are representational facts about the instrument and this state of the instrument. This is what the instrument was designed to do, what it is supposed to do, and, like any fallible system, it can fail to do what it is supposed to do. If a registration of "37" on a properly installed instrument fails to carry information about the speed of the car, or carries the same information that a registration of "24" carries, then it is not doing its job. The result, often enough, is misrepresentation.

The fact that the speedometer is connected to the axle by a cable that transmits information about speed, on the other hand, is not a representational fact about this instrument. It is a fact about a representational system, but it is not a representational fact. The instrument couldn't do its job without this cable, but the fact that there is such a cable does not imply that the device *has* a job—let alone has the job of providing information. For the same reason, the fact that a particular thermometer is filled with mercury, a metal whose volume serves to indicate temperature, is a fact about a representational device, but it is not a representational fact. A representational fact about S is a fact about what S is designed to do, a fact about what information it is supposed to carry. There are facts about representations—facts about their color, shape, material constitution, and mode of operation—that do not tell one anything about what information they are supposed to supply or, indeed, whether they are supposed to supply information at all.

On the representational account of mind I intend to give, the difference between representational facts and (mere[2]) facts about representations is the difference between the mind and the brain. Neuroscientists may know a great many facts about the brain. These facts may even turn out to be facts about mental representations—facts about experiences and thoughts. That does not make knowledge of such facts knowledge of the mind. Knowledge of the mind, of mental facts, is, according to the Representational Thesis, knowledge of representational facts, not (merely) facts about mental representations. One does not understand more about the representational life of a system by being told it contains mercury or has an orange pointer. Nor does one know what (or whether) a system represents by being told that it supplies information about speed. It is not whether it supplies information about speed that is important; it is whether it has the function of doing so.

Representation is here being understood to combine teleological with information-theoretic ideas. If the concept of representation is to do a useful job in cognitive science, if it is to be used, in particular, to illuminate the nature of thought and experience, it must be rich enough to allow for misrepresentation. It must include the power to get things wrong, the power to say that something is so when it is not so. This is what the teleology, the idea of something having an information-carrying function, is doing in the present theory. It captures the normative element inherent in the idea of representation. Since an object can retain a function even when it fails to perform it (think of the heart, the kidneys, and damaged instruments), a device can retain its indicator function—continue to represent (i.e., misrepresent) something as going 34 mph, for instance—even when things go wrong, even when it fails to provide the information it is its job to provide. There is information without functions, but there is no representation without functions.

Not all events that carry information have the function of carrying it. To use Matthen's (1988) example, the angle a column of smoke makes with the horizon carries information about wind speed, but that, surely, is not the function of the smoke. The smoke cannot misrepresent wind speed. We may be misled by the column of smoke. We may take it to indicate something it doesn't (thus ourselves misrepresenting wind speed), but the smoke does not misrepresent wind speed the way an anemometer can. This, incidentally, is why black-and-white television does not misrepresent the color of the sky while color television sometimes does. The power of the color television to misrepresent color lies in its color-depicting function. The images on the two screens can

be identical; yet, one misrepresents, the other does not. This is the difference between *not* representing color and *mis*representing color.

Something like this conception of representation is, I think, operative in many scientific approaches to mental representation. The terminology is not always the same, but the intuitions and theoretical purposes are similar. Anne Treisman (1992, p. 227), for example, speaks of representations as signals or events which have been assigned the "burden" (function?) of carrying meaning (information?). David Marr's (1982) well-known theory of vision assigns to early vision the function or task (Shapiro 1993) of depicting the spatial and chromatic properties of distal scenes. Gallistel (1990) describes the brain as representing aspects of the environment when there is a functioning isomorphism between the environment and the brain process that "adapts" the animal's behavior to it. Sensory organs and mechanisms are commonly described in terms of what they are "for." The semicircular canals of the middle ear are said to be for the detection of angular acceleration, the utricle and saccule for indication of linear acceleration, and the retina for encoding information about light for transmission to the brain. This way of describing perceptual mechanisms and processes is a representational way of thinking about them. The senses yield representations of the world, not just because they (when working right) deliver information about the world, but because that is their job. The senses, and the states they produce by way of performing their function, are thus evaluable in terms of how well they do their job.

Before applying these representational ideas to mental affairs, though, we need to know, in more precise terms, exactly what kind of representation a mental representation is supposed to be. If it is, indeed, the job of the various senses to provide information about the world—thus, on the present account, representing that world—what is the source of these functions, what information is it their job to provide, and how—in representational terms—do our experiences of objects differ from our thoughts about them? What, furthermore, is it about certain representations that makes the systems in which they occur *conscious* of what is being represented? Good questions. I begin my answers by explaining a number of pivotal distinctions: (1) natural vs. conventional representations; (2) representational states vs. representational systems; and (3) represented properties vs. represented objects. Armed with these distinctions we can give a preliminary taxonomy of mental representation and begin the argument that conscious experience is a species of natural representation.

2 Natural and Conventional Representations

The function of a system (or state) is what it is designed to do—what it is, by design, supposed to do. There are different sources of design. Each gives rise to a different kind of function and, thus, a different form of representation. One important difference (for our purposes) is the difference between naturally acquired and conventionally assigned functions—hence, the difference between natural and conventional representations.

The information-providing functions of measuring instruments, sensors, detectors and gauges are functions they get from us—their makers and users. We design them. We give them a job to do. We arrange things so that certain liquids, by their placement in transparent tubes adjacent to a calibrated scale, provide us with information about temperature. We call the resulting artifacts thermometers. Flag poles and (metal) paper clips, the volumes of which are proportional to temperature, also carry information about temperature. Their volume is also reliably correlated with temperature. That, though, is not their function. That is not what they are designed to do. As the names suggest, they have quite a different job. Though they carry the same information, flag poles and paper clips do not represent what thermometers represent. They do not represent anything. We haven't given them that kind of job.

When a thing's informational functions are derived from the intentions and purposes of its designers, builders, and users in this way, I call the resulting representations *conventional*. Representations that are not conventional are *natural*.

I assume that there are naturally acquired functions and, thus, natural representations.[3] I do not argue for this; I assume it. This view, I know, is not universally accepted (Dennett 1987, especially pp. 287–321 and Searle 1992, p. 52, for example, deny it). I return to this point in chapter 5, but for the present I follow the lead of Wright (1983, 1987), Kitcher (1993), Godfrey-Smith (1994), Millikan (1984), Neander (1991a, 1991b), Papineau (1993), Bennett (1976) and many others in supposing that bodily organs and mechanisms can, in the relevant sense, be designed to do a certain job—and, thus, have the function of doing it—without being designed by anyone to do it. Philip Kitcher (1993, p. 380) puts it thus: "one of Darwin's important discoveries is that we can think of design without a designer." The senses, I assume, have information-providing functions, biological functions, they derive from their evolutionary history.[4] As a result, perceptual (including proprioceptive) systems produce representations of those conditions (external or internal as the case may be) they have the function of informing about. The representations they produce by way of carrying out their informational functions have a content, something they say or mean, that does not depend on the existence of our purposes and intentions. This is why the senses—or, more precisely, the internal states (experiences, feelings) the senses produce by way of performing their function—have original intentionality, something they represent, say, or mean, that they do not get from us. That is why the perceptual representations in biological systems—unlike those in laptop computers, speedometers, and television sets—make the systems in which they occur *conscious of* the objects they represent.[5]

We have, then, the following preliminary classification.

Representational systems
Systems having indicator functions

Conventional
Indicator functions
are conventional
Instruments
Language

Natural
Indicator functions
are natural
Sensory systems
Concepts

Hereafter, when speaking of the Representational Thesis, I will mean the thesis that all mental states are *natural* representations. This makes the thesis a form of philosophical naturalism.

3 Representational Systems and Representational States

The above classification is crude. All mental representations (not to mention nonmental natural representations[6]) are

classified together. Thoughts and beliefs are classified with experiences. This is correct, as far as it goes. Both are forms of natural representation. For our purposes, though, it is important to distinguish seeing and hearing from knowing and believing. I can see Paul playing the piano and believe he is playing the piano, but the visual experience represents the piano playing in much different ways than does the belief. These are different kinds of mental representation. One can see or hear a piano being played without believing a piano is being played, and one can believe a piano is being played without seeing or hearing it being played. Seeing a piano being played is constituted, in part, by a visual experience, hearing by an auditory experience. Until these experiences occur one has not seen or heard the piano.[7] Experiences of piano playing do not require the concept of a piano (at least not in the same way as a belief or judgment requires it). They require no understanding of what a piano is or what it sounds like. Even mice can see and hear pianos being played.

Believing is something else. It requires the concept of a piano, some understanding of what a piano is. Mice who hear pianos being played do not believe pianos are being played. Their understanding is, I assume, too feeble to believe this even though their hearing is good enough to hear it.

All representations are representations of (purported) fact, but not all such representations are *conceptual* representations. A conceptual awareness of facts—a belief, judgment, or knowledge that the toast is burning—has a close tie with behavior. For those who have language, it normally brings with it an ability to *say* what one is aware of—that the toast is burning. This is not so with sensory awareness. One can see or smell (and, thus, be perceptually aware of) burning toast while having little or no understanding of what toast is or what it means to burn. "What is that strange smell?" might be the remark of someone who smells the toast burning but is ignorant of what toast is and what it means to burn. A mouse in the kitchen (the one who heard the piano being played) can smell, and thus have sensory awareness of burning toast but it will not (like the cook) be aware (i.e., believe) that toast is burning. The mouse will have little or no conceptual awareness of this event. It will smell the toast burning but not smell that it is, not smell it *as*, burning. The cook and the mouse differ in what they think about what they smell. They differ in how they regard the smell. They may also differ in how the burning toast smells to them. But they both smell it. They both experience it. They both have some kind of sensory representation of this event.[8]

The difference between experiences of *k* (as F) and thoughts about *k* (that it is F)—between sensory and conceptual representations of *k*—is generally clear enough when we describe ourselves as conscious of concrete objects (e.g., burning toast) and events (e.g., piano playing). When, however, we start describing ourselves as being aware of abstract objects—differences, numbers, answers, problems, sizes, colors—an ambiguity appears. I take a moment, therefore, to remark on this ambiguity since it would otherwise muddy the discussion to follow.

When we use an abstract noun or phrase to describe what we see, hear, or feel—what we are aware or conscious of—what is being described is normally a conceptual awareness of some (unspecified) fact. The abstract noun phrase stands in for some factive clause (Dretske 1993). Thus, to describe someone as seeing (being conscious of) the difference between A and B is to imply that the person sees (is conscious) *that* they differ. To describe someone as being

aware of the color of his (blue) shirt is to imply that he or she is aware *that* the shirt is blue—thus representing the color in some conceptual way (as blue, the color of the sky, etc.) It would be odd to describe a person as seeing (thus being aware) of the color of his shirt if the person did not, at some conceptual level, know what color the shirt was. Likewise, to be aware of the problem it isn't enough to see (experience) the thing that is the problem (e.g., a clogged drain). One has to see (the fact) that it is clogged. One has to conceptualize what one sees, the clogged drain, *as* a problem. Until one conceptualizes it this way, one is not (as we say) aware of the problem.

These differences are important in thinking about ways of representing properties and, thus, according to a representational theory, the properties we are conscious of. Since the topic is our experience of objects, not our beliefs about them, we must be careful not to describe an experience of a shirt's color as an awareness or consciousness of the shirt's color. For this form of words implies[9] a conceptual representation of the shirt's color, and this may not be present. A child or an animal might be visually aware of the shirt's color (their visual experience of the shirt being, as they say, suffused with blueness) without their knowing or thinking that the shirt is blue—without sorting (or having any disposition to sort) the shirt with other blue objects. If this is hard to imagine, think of the perceiver as an animal. Cats are not color blind just because they ignore differences in color.[10] One can experience blue (the shirt's color)—and, in this sense, be aware of blue—without being conceptually aware that anything is (or looks) blue.

I will try to keep these matters straight by distinguishing between sensory and conceptual representations of facts, between experiences of *k*'s blueness and beliefs or judgments that *k* (or something) is (or looks) blue. The word "phenomenal" is often used to describe this sensory mode of awareness, and I will sometimes use it. Phenomenal awareness is a mode of awareness that does not require—although it may, in fact, be accompanied by—conceptual awareness. One can be phenomenally conscious of a shirt's color, of a piano being played, and of burning toast without being conscious that anything is blue, that a piano is being played, or that something is burning.

What, then, in representational terms, is the basis of this ordinary—and, I hope, familiar—distinction between an experience of color, shape, and texture, and a belief or judgment about color, shape, and texture? In representational terms, what is the difference between an olfactory experience of toast burning and a belief or judgment that toast is burning?

Experiences (sensations) of burning toast and beliefs (thoughts) that toast is burning are representations, and all representations are particular (token) states or events. Nonetheless, token states have two different sources for their indicator functions. (1) A state may derive its indicator function—and, hence, its representational status—from the system of which it is a state. Call these *systemic* indicator functions (= functions_s) and the representations they give rise to systemic representations (= representations_s). If a system (e.g., a thermometer) is supposed to provide information about temperature, and ß is the state (e.g., mercury at such-and-such level) that is supposed to carry the information that the temperature is, say, 32°, then ß has the systemic function (function_s) of indicating a temperature of 32°. State ß therefore represents_s the temperature as being 32°. (2) A token state may, on the other hand, acquire its indicator function, not from the system of which it is a state, but from the *type* of state of which it is a token. No matter what ß systemically represents (what it is, by design of the system, supposed to indicate), it might acquire (or be given) a special, or a different, indicator function. If we print the number "38" at *this* point on the scale, then each and every time the mercury reaches *this* point, it *means* that the temperature is 38°. If we print the word "DANGER" at this (and higher) points, the mercury's rising to this point signifies danger. Call such functions (functions that state types are assigned or acquire independently of their—if any—systemic functions) *acquired* indicator functions (= functions_a). The indicator functions_a (hence, representations_a) of token states may be different from their systemic functions (what they represent_s). As we shall see, a token state might represent in both systemic and acquired ways, and what it represents_a need not be what it represents_s.

An example should make this difference clear. Suppose we have a simple speedometer mechanism that represents vehicle speed by registering the rotation of the axle. This mechanism was designed to be used in cars equipped with different sized tires. Since (at a given speed) the axle rotates more slowly with large tires than with small tires, the manufacturer left the job of calibrating the face of the instrument up to the user. If I use the instrument in a car with normal tires, I calibrate the dial in one way. If you use larger tires, you calibrate it differently. I put the number "50" at position ß of the pointer. You put the number "60" there. When our axles are turning at a rate corresponding to a pointer position of ß, my car is going 50 mph, your car is going 60 mph. So we calibrate differently by assigning different numbers to this pointer position. Pointer position ß, this state of the system, has the same systemic function in both cars—that of indicating an axle rotation rate of N rpm. It nonetheless, via different calibration, has a different acquired function in our two cars. In my car state ß represents$_a$ 50 mph, in your car it represents$_a$ 60 mph. State ß has the same indicator function$_s$ in all cars; what it represents$_a$ (about speed), however, varies from car to car.

To suggest (by way of tendentious description) the way this distinction is to be applied, we can describe individual systems in which this "perceptual" mechanism is installed as having the same "experiences" (viz., of an axle rotation of N) but as having different "belief's" (about speed). ß represents$_s$ (i.e., is an experience of) the same thing in all cars: an axle rotation of N. That is what this state means, what it represents$_s$, in both your car and my car. Given the information (about axle rotation) this system has the function of supplying, what it was designed (at the factory) to do, this particular state of the system has the *systemic* function of indicating an axle rotating at N rpm. Nonetheless, although this state represents$_s$ the same thing in both systems, it represents$_a$ something different in my car than it does in your car. It represents$_a$ 50 mph in my car, 60 mph in your car. Once calibrated, my speedometer, as it were, "sees" an axle rotation of N rpm as a speed of 50 mph. Your speedometer, differently calibrated, "sees" the same thing—an axle rotation of N—but sees it *as* 60 mph. Our speedometers have the same "experience" but the experience gives rise to different "beliefs."

Walker (1983, p. 246) describes two of Pavlov's dogs. One is conditioned to salivate when middle C is played on any instrument. The other is conditioned to salivate to a clarinet playing any note. Now imagine the two dogs hearing a clarinet playing middle C. They hear the same sound. Their experience of this sound may well be the same. Their response is also the same: both salivate. Yet, their responses

are mediated by different acquired representations. As a result of different learning, the dogs hear it differently—one (as we, not the dogs, might put it) hears it as middle C, the other as the sound of a clarinet. The way their experience represents$_s$ the sound may well be the same, but the way their experience represents$_a$ it is different.

As my description of these examples is intended to suggest, experiences are to be identified with states whose representational properties are *systemic*. Thought (conceptual states in general), on the other hand, are states whose representational properties are *acquired*. As a result, experiences have their representational content fixed by the biological functions of the sensory systems of which they are states.[11] How an experience represents$_s$ the world is fixed by the functions of the system of which it is a state. The quality of a sensory state—how things look, sound, and feel at the most basic (phenomenal) level—is thus determined phylogenetically. Since we inherit our sensory systems, since they are (at a fairly early age, anyway) hard-wired, we cannot (not easily anyway) change the representational$_s$ character of experience.[12] Through learning, I can change what I believe when I see *k*, but I can't much change the way *k* looks (phenomenally) to me, the kind of visual experience *k* produces in me. Experiences are, for this reason, modular in Fodor's (1983) sense. The way a belief represents the world, on the other hand, is ontogenetically determined. We can, through learning, change our calibration. We can change what we see something *as*—what we, upon seeing it, take it to be—even if we cannot, not in the same way, change what we see. This is why a representation$_s$ of *k* as red (a sensation of redness) is different from a representation$_a$ of *k* as red (a belief that *k* is red) even though both are representations of *k* as red.

The above example is a bit too simple-minded to reveal the full power of the represented$_s$-represented$_a$ distinction to illuminate the sensation-cognition (experience-belief) difference. Perhaps, therefore, a minor embellishment will better reveal its potential for capturing some of the structure of actual sensory systems (I am thinking mainly of constancy mechanisms).

Another manufacturer produces a more sophisticated speedometer. It is designed to be used on all cars, no matter what size tires they use. Pointer positions are determined by two sources of information: not just by the rate at which the axle is rotating, but by the height of the axle above the road surface. Since the height of the axle above the road surface provides a useful measure of tire size, these two sources of information are combined to yield a reliable measure of car speed no matter what size tires are used. The pointer positions of these instruments—driven by two sources of information—thus have a speed-indicating systemic function and are therefore calibrated at the factory. As their construction suggests, these more sophisticated instruments are designed to provide information, not about how fast the axle is rotating, but about how fast the car is going. Although the information-handling processes in these fancy instruments use information about axle rotation to generate a representation of speed, the final representation (pointer position) does not, in fact, indicate how fast the axle is rotating. The early phase in this informational process uses information about axle-rotation, but (by integrating it with information about elevation) it sacrifices this information in order to produce a final indication of speed. Information about axle rotations is thereby lost. This is an instrument

whose individual states (pointer positions) represent$_s$ speed, not axle rotation. This instrument, unlike the crude one described above, "experiences" speed (not axle rotation) even though the information-delivery process depends on information about axle rotation to generate this "experience."

But if speed is what this instrument "experiences" (= what its various states have the systemic function of indicating) what does it "believe"? This depends on what the states of this instrument represent$_a$, what they have acquired (or been assigned[13]) the job of indicating. Normally, of course, if the instrument was to be used as a speedometer, one would assign functions to its various states in the way that reflected their systemic functions. The state, ß, having the function$_s$ of indicating a speed of 50 mph, would be given this function$_a$ by printing the number "50" at this position of the pointer.[14] We are, however, free to assign any functions$_a$ we please. If we aren't particularly interested in exactly how fast the car is going, we could partition the dial into several sub-divisions and label them "slow," "medium," and "fast." Then ß (having the function$_s$ of indicating a speed of 50 mph) gets the same function$_a$ as the state _ (having the function$_s$ of indicating a speed of 48 mph)—the function$_a$, namely, of indicating a *medium* speed. This, in fact, might be the functions these states acquire if we used the device, not as a speedometer, but as a mechanism to control an automatic transmission. In this case, since we want the car to be in 3d gear at both 48 mph and 50 mph, we give these two states the same function$_a$. If the various states of the device acquire *these* functions$_a$, then it will end up "experiencing" 50 mph and 48 mph (these are the speeds these two states represent$_s$), but representing$_a$ them both in the same way (as a 3d gear speed). It experiences a difference between 50 mph and 48 mph, yes, but it, so to speak, "sees" both these speeds as medium (or 3d gear) speeds. Conceptually, this system "abstracts" from the sensory difference between 48 and 50 mph.

The last example is meant to be suggestive of the way innate sensory systems might have the function of delivering information about more or less continuous quantities (surface reflectances, speed, direction, frequency, etc.) while the cognitive system might "chunk" this information into parcels that are more useful to the individual system's needs and purposes. At the level of experience, I am sensitive to (i.e., can discriminate) all manner of differences in the light, sound, pressure, temperature, and chemistry of objects affecting my senses. I nonetheless have a limited conceptual repertoire for categorizing these sensory differences, for making judgments about (= representations$_a$ of) the differences I experience. I experience the differences (in the same way the speedometer described above registers a difference between 50 mph and 48 mph), but, at the conceptual level, I treat these differences as different instances of the same kind. At the sensory level I can discriminate hundreds of different colors. At the conceptual level I operate with, at best, a few dozen categories for the colors I experience. Learning calibrates me no more finely than I need to satisfy needs and carry out purposes. This is not to say, of course, that I could not "recalibrate" in a more discriminating way if the need arose. The information needed for such recalibration is already there in my experience.

The taxonomy, then, looks like this:

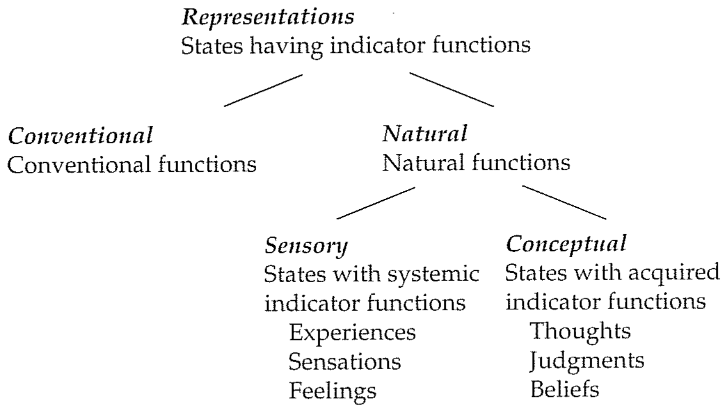

A few clarificatory remarks are in order:

(a) Experiences are representations$_s$, but, as noted earlier, not all representations$_s$ (not even all natural representations$_s$) are mental—let alone experiences. Experiences are those natural representations$_s$ that service the construction of representations$_a$, representations$_s$ that can be calibrated (by learning) to more effectively service an organism's needs and desires. They are the states whose functions it is to supply information to a cognitive system for calibration and use in the control and regulation of behavior.[15] Evans (1982, chapter 7, §4) expresses a similar idea when he describes internal states whose content depends on their phylogenetically ancient connections with the motor system. In order to qualify as conscious experiences, Evans requires that these content-bearing states serve as input to what he calls a "concept-exercising and reasoning system."[16] This is why, in our (second) speedometer example, the early stages in the information-delivery process (the ones carrying information about axle rotation, for example) are not appropriate analogues of experience. Though these early processes have a function$_s$—that of providing information about axle rotation (information the system needs to generate a representation of speed)—and therefore represent$_s$ axle rotation, this information is not available for calibration. The instrument *uses* this information, but doesn't, so to speak, make it available to calibrational (cognitive) processes. There is, then, no way this information can be given "conceptual" form. The information vanishes (absorbed into information about speed) before it reaches the relevant control structures that determine behavior (in this case, pointer positions). These are natural representations$_s$ that are not mental. Their function is to supply information to representations$_s$ that are.[17]

(b) Sensory systems have phylogenetic functions and are, therefore, comparatively modular in Fodor's (1983) sense: hard-wired, informationally incapsulated, and so on. This is not to say that the states (experiences) they produce are immune to changes in their representational status. We can use a pressure gauge (function$_s$ = delivering information about pressure) as an altimeter by recalibrating its face in "feet above sea level." Pointer positions representing$_s$ pressure now (also) represent$_a$ altitude. The gauge, as it were, experiences pressure *as* altitude. This same transformation occurs as we learn to identify and recognize objects and conditions we experience. We begin by hearing (experiencing) sounds and end by hearing (recognizing) words. We still hear sounds, of course, but, after learning, after the kind of calibration occurs that is involved in language learning, experiences acquire an added representational dimension.[18]

Sensory adaptation may be an instance of such recalibration. Just as we can recalibrate a speedometer when we

change the size of our tires (the state ß acquiring the function of indicating a speed of 60 mph, not, as it did with normal tires, a speed of 50 mph) so also can sensory systems recalibrate when there is a change in the information an experience delivers. After wearing spectacles that make everything "look" 30° to the left, touch "educates" vision and objects begin to "look" to be where they can be grasped. This, at least, is one way of describing the experiential change that accompanies adaptation. As we age, there is a gradual yellowing of the macula (spot on the retina) that changes the signals sent from the retina to the brain about the wavelength of light (example taken from Clark 1993, p. 170). Despite this constant change in the information representational$_s$ states carry, there is no corresponding representational$_a$ change: we still see blue and yellow, red and green, *as* blue and yellow, red and green.

(c) The states by means of which a representational system performs its informational functions have a structure that enables them to acquire functions without explicitly acquiring them. Once we assign "12" to a certain position of the clock hands, all the remaining indicator states (positions of the clock hands) receive an "implied" indicator function. When there is this kind of structure among a system's indicator states (as there is in our experience of color and sound, for instance), there is no reason to suppose that *each* representational$_s$ state has explicitly obtained its indicator function through some distinguishable evolutionary process. All the indicator states may receive an implied indicator function by one state explicitly acquiring it.

(d) Audition provides information, not only about pitch and intensity, but about timbre and (binaurally) the directional properties of sound. In the case of vision, there seems to be separate pathways, different systems, for information about color, location, form, and movement (Fischback 1992, p. 56; Zeki 1992, p. 71).

It would be more appropriate, then, to describe a given sense modality, not (as I have been doing) as having the function of indicating F, but as having multiple indicator functions. And there is no reason experiences "associated" with a given sense modality (smell or vision) might not exploit information from other modalities (see, e.g., Stein and Meredith 1993; Dretske 1981, pp. 145–147).

(e) Finally, the subjective quality of an experience, the phenomenal appearances, are the way experience represents$_s$ things to be. If I can use the speedometer analogy again, the first speedometer has axle-rotation qualia. The second has speed qualia. Once calibrated, though, both "describe" (i.e., represent$_a$) how things seem to them in the same way—as a vehicle speed of so-and-so many miles per hour. If we are talking about vehicle speed, things *seem* the same to them.

I am, with these brief remarks, trying to indicate how I intend to use the distinction between systemic and acquired representations. Justification for thinking of experience and thought in this way will come later. I am quite aware that questions about qualia, about what it is like to have an experience, about the way things seem, are complex questions that, often enough, straddle the sensation-cognition fence. We are sometimes talking about the properties objects are represented$_s$ as having (the phenomenal appearances), at other times about the ones they are represented$_a$ as having (what the phenomenal appearances prompt us to believe, or what we, on the basis of the experience, are inclined to take something as). I return to these tangled issues in chapters 3 and 5.[19]

4 Represented Properties and Represented Objects

Ordinary thermometers represent the temperature of whatever medium they are in. Put one in your coffee and it tells you how hot your coffee is. Hang the same thermometer on the wall and it tells you the temperature of the room. Put it in your mouth and it indicates something about you—whether you have a fever. It isn't the function of this instrument to say what it is—coffee, living room, or you—whose temperature it registers. The instrument says how hot *this* is, where *this* is whatever medium the thermometer is in. It does not represent the relationship (i.e., its being *in* this) that makes it *this* rather than *that* it represents the temperature of.

The same is true of all representational devices. Pressure gauges do not tell you which (if any) tire it is that they represent the pressure of. If you want to know which tire (if any) the gauge represents, which topic it is commenting on, you have to look, not at the gauge, but at the external connections between gauge and world that make it the right front tire, not the left rear tire, whose pressure it registers. Gauges do not supply—they do not have the function of supplying—information about these external connections. Representations have a sense (the properties they have the function of indicating) and, often enough, a reference (an object whose properties they represent), but the sense does not determine the reference. Two representations with the same sense can have different referents.

The difference between represented object and represented property, between the reference and the sense of a representation, is the same distinction Nelson Goodman (1976) was getting at in his contrast between a picture of a black horse ("black horse" here specifying the object the picture is a picture of) and a black-horse picture ("black-horse" here specifying what the picture depicts the object as). Some pictures of black horses do not represent the black horse as a black horse. They are not black-horse pictures. Imagine taking a picture of a black horse at a great distance. The picture depicts the black horse as, say, a spot in the distance. It could be anything. Or imagine taking a picture of a black horse disguised to look like a brown camel. This picture of a black horse is a brown-camel picture. Others (e.g., Lloyd 1989, p. 14) make the same distinction by contrasting explicit with extensional content. The extensional content is what object it is that is being represented. The explicit content is the way this object is being represensted—as a so-and-so. The object a representation is a representation of is not determined by the properties that object is represented as having. Nothing guarantees that black-horse pictures are pictures of black horses.

Experiences are like pictures in this regard; they are *de re* modes of representation (Burge 1977; Bach 1986, 1987; Recanati 1993). Although token experiences may be individuated by being experiences of a particular object, *k*, tokens of the same experience type can occur with a different *k* or with no object at all (hallucinations, dreams, imagination). Although my present black-horse experience is an experience of a black horse (I am, that is, seeing a black horse), I can nonetheless have exactly the same type of experience, a black-horse experience, without its being an experience of a black horse or, indeed, an experience of any object at all.[20] What determines the reference for a *de re* mode of representation (the object it is a representation of) is not *how* it is represented, but a certain external causal or contextual relation I will designate as C.[21] There is nothing in the content of the representation, nothing the representation *says*, which

makes it about this object rather than that object or no object at all. *De re* modes of representation have their reference determined contextually, by the relation I am here calling C. Since the veridicality of an experience depends on its reference—on what object (if any) it is an experience of—the veridicality of experience is determined, in part, by context (C). C makes it the case that *k* is the object S represents as blue. C thereby helps to determine whether the representation (that *k* is blue) is veridical or not. Nonetheless, the fact that it is *k* (rather than some other object or no object at all) that stands in relation C to the representation is not what the representation represents. Representations do not (indeed, cannot) represent context. They represent *k* as being blue, but they do so without representing it to be *k* that they represent to be blue.

The speedometer in my car is connected to the axle of my car, not your car. It therefore represents (or misrepresents, as the case may be) the speed of my car, not your car. It is in virtue of this special relation, C, to my car that it can do what other speedometers cannot do—viz., say something (whether truly or falsely) about my car. Other speedometers cannot even say something false about my car. C is the relation such that, when a representational system S is functioning properly, and *k* stands in C to S, then S will indicate the F of *k*. If *k* is the object whose F is represented at time t, then we can say that, at time t, S represents *k*. At other times S may be deployed differently—thus, at different times, representing different objects. Changing the way a system is deployed, changing the object that stands in relation C to S, changes what S represents (what *object* it represents), but not necessarily what S says (represents) about it.

As I am using terms, then, the fact that S represents *k* is not—at least not a pure—representational fact about S. It is a fact about a representation—an important fact for many practical purposes—but it is not a fact that has to do exclusively with what the system has the function of indicating. S has the function of indicating the F of those objects which stand in C to it, but it does not have the job of indicating—does not therefore represent—which objects—or even whether there is an object—that stands in C to it. When we describe S (whose function it is to indicate F) as representing *k*, this implies that, for some F, S represents (possibly misrepresents) the F of *k*. That S represents *k*, therefore, implies a representational fact—that, for some F, S represents the F of *k*. But it also implies something that is not a representational fact—viz., that *k* stands in relation C to S. So facts about the object of representation are *hybrid* facts—part representational, part not. This will be important when we discuss (in chapter 2) introspective knowledge. For what is known by introspection are mental—hence, (according to the Representational Thesis) representational—facts. Hybrid representational facts (that S represents *k*) are mixtures of representational facts and facts about representations. This is why one cannot know, at least not by introspection, what object (or whether there is an object) one is experiencing.[22]

Clearly, then, my use of the word "something" in the description of S as representing something as being blue is not an existential quantifier. It may turn out that S is misrepresenting something to be F when there is something (in the next room, say) that *is* F. The fact that there is something in the world that is going 25 mph does not mean that a speedometer registering "25 mph" says something true about the world. For what the speedometer "says" is not that there is something in the world going 25 mph, but that *this* (whatever it is that stands in C to the instrument) is going 25 mph. If representational system S says anything at

all when it represents color—and, thus, represents something as being (say) blue—it is that *this* is blue where *this* is whatever object (if there is one) to which S stands in the C-relation. If there is no object of representation, then, S represents (i.e., misrepresents) *this* as being blue when there is no *this*. If we take this to mean that, owing to a failure of presupposition, what S represents to be so (the content of the representation) is neither true nor false, this merely shows something about the nature of sense experience that we knew all along: viz., that misrepresentation takes two forms. An experience can misrepresent by (1) saying something false, by saying that *this* is blue when *it* (the object of representation) is not blue, and (2) by saying what is neither true nor false—that *this* is blue when there is no "this" that is not blue. When there is no object of representation, nothing that S represents to be blue, I shall continue to describe S as representing something to be blue. It should be understood, however, that when there is no object, "something" stands in for a failed indexical.

We thus have two different kinds of misrepresentation:

S represents color
S represents something to be (say) blue

Object of representation
There is an object S represents to be blue

- *Veridical*
  The object S represents to be blue is blue
- *Misrepresentation*
  The object S represents to be blue is not blue

No object of representation
There is no object S represents to be blue

- *Misrepresentation*
  S represents something to be blue, but there is no object that S represents to be blue

1 French Poodles and French Wine

Susan, a child of normal eyesight and intelligence, has never seen a dog. She doesn't know what dogs are. The first one she sees is a French poodle. Does it look like a poodle to her? Susan will not say or think it is (or looks like) a poodle. She will not see it *as*, or see that it is, a poodle. This is not the way she will conceptualize what she sees. Will it, nonetheless, look like a poodle to her?

Arthur is a young toad of normal toad-eyesight and toad-intelligence. He has never seen a dog. The first one he sees happens to be a poodle—the same one Susan sees. Does it look like a poodle? Like Susan, Arthur will not say or think it is (or looks like) a poodle. He will not see it as, or see that it is, a poodle. This is not the way he will conceptualize what he sees. Will it, nonetheless, look like a poodle to him?

Given what we know about the eyesight of human beings and toads, it seems clear that the poodle looks different to Susan than it does to Arthur. The dog looks to Susan the way poodles look to you and me—different from the way bulldogs, terriers, and sheepdogs look. You and I can describe the way the dog looks ("like a poodle") and Susan cannot, but this, surely, makes no difference to how the dog looks to her. Arthur is different. Though the dog looks to Arthur the way poodles look to normal toads, and Arthur is a normal toad, the dog does not look like a poodle to Arthur. Poodles do not look like poodles to half-blind people despite the fact that they look the way poodles normally look to half-blind people. Poodles look like blurry spots to half-blind people—the same almost everything else this size looks. Toads have the visual acuity of half-blind people. So poodles look to toads the way poodles look to half-blind people.

In saying that the dog looks like a poodle to Susan I am not saying that Susan's visual experience of the poodle represents it as a poodle and, therefore, *mis*represents it in some way if it is not a real poodle. No, we are here speaking of sensory, not conceptual, representations of the poodle. Susan's experience of the dog represents the dog as a poodle in the sense that it represents the dog as having what McGinn (1982, p. 40; see also Millar 1991, p. 42) describes as the manifest properties of poodles, those properties that make poodles look so much different from other dogs (not to mention bicycles, etc.). A variety of nonpoodles also have the manifest properties of poodles: *trompe l'oeil*, pictures of poodles, disguised terriers, poodle robots, and so on. These objects also look like poodles. They, too, produce, poodle qualia in Susan. Although she would be making a mistake in believing (i.e., conceptually representing) a good fake to be a poodle, Susan is not misperceiving (sensuously misrepresenting) it when she visually represents a fake poodle in the same way she visually represents a real poodle. Good fakes are supposed to cause the same kind of experiences as the originals.

Some will feel it wrong to describe anything as looking like a poodle to someone who does not have the concept of a poodle. Millar (1991, p. 32), for instance, says that a person could have pumpkin experiences of exactly the same type I have when I see a pumpkin, but if that person did not have the concept of a pumpkin, did not know what a pumpkin was, it would not "seem" to that person that there was a pumpkin there. I do not want to quarrel about words. There may be a sense of "look" and other "appear" words (especially when used with a factive clause) in which Millar is right. To keep things straight I will call these uses the *doxastic* (= belief) sense of "look" ("appear," "seem," etc.) and will subscript it accordingly (Jackson [1977, p. 30], calls it the *epistemic* use of "look"). To say that a dog $looks_d$ like a poodle to S is to say that, in the absence of countervailing considerations, this is what S would take the dog to be, what S's perception of the dog would (normally) prompt S to believe (see Millar 1991, p. 19). Describing the dog as $looking_d$ like a poodle to S implies that S has the concept POODLE, understands what a poodle is, and classifies or identifies (or *would* do so in the absence of countervailing considerations) what she sees in this way. In the doxastic sense, the dog Susan sees does not $look_d$ like a poodle to her. That isn't what her perception of the dog causes her to believe.

There is, though, another sense of these words, a sense in which if the dog looks the same to Susan as it looks to me, and it looks to me like a poodle, then it must look to Susan like a poodle whether or not she understands what a poodle is, whether or not she has the concept of a poodle. Following a long tradition, I will call this the *phenomenal* sense of "look" ($look_p$). To say that the dog looks—phenomenally—like a poodle to S (= $looks_p$) is to say two things: (1) that the dog looks to S the way poodles normally look to S; and (2) the dog looks different to S from other dogs (bulldogs, terriers, etc.).[1] The second clause—call it the *discriminatory* clause—is necessary to prevent a dog looking like a poodle to a toad (or a half blind person) to whom all dogs look the same. If you are color blind, if you cannot discriminate one color from another, an object does not $look_p$ red to you simply because it looks the same as red objects normally look to you. You do not satisfy the discrimination condition. The discrimination condition is here left deliberately vague to reflect the context-sensitivity and circumstantial relativity of these judgments (*k* can look like a poodle to you in circumstances C but not in C′). Even if there is some rare breed of dog that you cannot distinguish from a French poodle, a dog might nonetheless "$look_p$ like a French poodle to you" as long as the implied set of contrasts did not include the rare breed. This is why a dog can $look_p$ like a French poodle to one even though one cannot discriminate it from a very good fake. Fake poodles are not—not normally, anyway—in the contrast-class.[2] Hence, despite the discriminatory clause, one does not have to be able to discriminate French poodles from very good fakes in order for a dog to $look_p$ to you like a French poodle.[3]

It should be noted that simply because an object $looks_p$ F to two people does not mean it $looks_p$ the same to them. It may turn out S_1 and S_2 are equally good at discriminating Fs from other objects and that a given object, *k*, looks to S_1 the way Fs usually look to S_1 and looks to S_2 the way Fs usually look to S_2. None of this implies that the object looks to S_1 the way it looks to S_2. Maybe it does, maybe it doesn't. The fact that a dog $looks_p$ like a poodle to Susan *and* (say) to another poodle does not mean it looks the same to them. The possibility of inverted qualia is not foreclosed by these definitions. More of this in a moment.

With these terminological stipulations, we can say that the dog $looks_p$ like a poodle to Susan but not to Arthur (it $looks_d$ like a poodle to neither, of course). A toad's visual system was not designed to discriminate the forms of middle-sized objects in the way the visual system of humans was. The proof of this lies in the inability of toads to perceptually discriminate the shapes of middle-sized objects. Just what poodles look like to a toad depends, in part, on the toad's powers of discrimination. If, as I have been assuming (on the basis of neurophysiological and behavioral data[4]) middle-sized animals that are moving (toads don't seem able to see the stationary ones) look the same to toads, then toads cannot see differences that normal human beings see, differences in form and detail that make poodles look different from bulldogs to us. A toad's visual system does not represent$_s$ poodles the way our visual system represents$_s$ them. So poodles, I infer, must look$_p$ to toads the way poodles look to us when we see them in poor light and without our glasses—pretty much the way bulldogs and terriers look.

We are also in a better position to describe certain (otherwise) puzzling situations. Suppose all red wines taste the same to me. Does this mean that the red Burgundy I am presently tasting tastes$_p$ like a red Burgundy to me? It *does*, after all, taste the way red Burgundy normally tastes to me. No, not if I cannot discriminate it from a Chianti.[5] If Burgundies taste the same to me as Chianti, chances are neither Burgundies nor Chiantis taste like Burgundy to me. In all likelihood, both taste like an ordinary red "jug" wine. Maybe not. It depends on whether I can distinguish "jug" wine from Coca-Cola. If I stand (gustatorily) to wines the way toads stand (visually) to dogs, Burgundy and Chianti are to me what poodles and bulldogs are to toads—the gustatory equivalent of blurry spots.

Application of the same criteria for telling how things "look" or "seem" to a system is evident in the way we think and talk about instruments. If a measuring instrument is not designed to discriminate between 7.00 and 7.01, if numbers beyond the decimal point are not (as they say) significant digits for it, then the instrument cannot represent anything as 7.00. Nothing can "seem" like 7.00 to an instrument that has no internal state with the function of indicating a value of between 6.99 and 7.01. This, basically, is the reason nothing can "seem" like a poodle to Arthur. Arthur's visual system does not have the requisite "resolving" power. Nothing in Arthur has the function of indicating that yonder object has *this* (a poodle) shape rather than *that* (a bulldog) shape.

Susan's visual system is a more sensitive instrument.[6] She, as it were, experiences 7.00 shapes and Arthur does not.

We must, however, be careful in using discriminatory data in making inferences about appearances$_p$—about the quality of someone's experience. Such inferences move from premises about what animals can and cannot discriminate in a given sensory mode to conclusions about how something appears$_p$ to them in that mode. That this is fallacious—at least suspicious—can be seen by considering the wine example. Even though all red wines taste the same to me, maybe they all taste$_p$ like exquisite Burgundy and not, as I was assuming, like ordinary red table wine. If this were so, then you and I might be the same in our discriminatory powers (we both embarrass ourselves at wine tasting parties) even though all wines taste different to us. All wines taste to me the way fine Burgundies taste to a connoisseur. All wines taste to you the way cheap Chianti tastes to a connoisseur. This gustatory anomaly may be hard (impossible?) to detect. Every time you, sipping an expensive Burgundy, say that the wine tastes like *that* (gesturing toward a cheap Chianti), I agree. That is exactly the way it tastes to me.

If the wine example seems far-fetched, Clark (1993, p. 167) gives a more compelling example:

> If you close one eye, stare at some bright colour for 30 seconds, and then blink your eyes successively, you will note shifts in the chromatic appearance of most things. The adaptation difference between the two eyes vanishes quickly, so the effects will soon disappear. While they last, however, these adaptation effects are similar in several ways to the purported differences between people. Discriminations made with the adapted eye will match those made with the unadapted eye. Any two items that present different hues to the unadapted eye will present different hues to the adapted eye. Matching items will continue to match as well. But the apparent hue of everything shifts.

Clark concludes by saying that the possibility of interpersonal differences of this sort seems fatal to the project of defining sensory qualities in terms of what can be discriminated. Shoemaker (1975) makes the same point.

It seems, then, that we might be indistinguishable in our discriminatory behavior and, yet, different in the way things taste to us—different in the way we sensuously represent the objects we perceive. This, of course, is the inverted spectrum problem.[7] The "problem" is a problem for those—e.g., behaviorists and functionalists[8]—who think the quality of experience must, somehow, be defined in behavioral or functional terms.

The Representational Thesis is a naturalistic theory that avoids this problem.[9] The qualitative character of perceptual experience, it concedes, is not functionally definable. It is, however, physically definable. By identifying qualia with the properties that the experience represents$_s$ things as having, a representational approach to the mind does two things: (1) It respects the widely shared (even by functionalists[10]) intuition that qualitative aspects of experience are subjective or private: they do not necessarily express themselves in the behavior (or behavioral dispositions) of the system in which they exist. (2) It provides an account of sense experience which makes the qualitative aspects of experience objectively determinable. In identifying qualia with experienced properties, experienced properties with properties represented$_s$, and the latter with those properties the senses have the natural function of providing information about, a representational approach to experience makes qualia as objectively determinable as are the biological functions of bodily organs. It may be hard—sometimes (from a practical standpoint) impossible—to discover what the function$_s$ of a certain state is, but there is nothing essentially private or exclusively first-person about functions.

Eduardo Bisiach (1992), a neuropsychologists, despairs of finding an objective way to study qualia. "There is," he says, "no way in which a natural science of consciousness could have anything to do with qualia" (p. 115). What I hope to do in the remainder of this chapter is to show that such pessimism can be avoided. A representational account of experience not only makes room for qualia, it provides an objective way of studying them.

2 Qualia as Represented Properties

It is hard to find a description of qualia with which two (let alone all) philosophers would agree, but it seems safe enough to begin by saying that the qualia in sense modality M (for S) are the ways objects phenomenally appear or seem to S in M. In accordance with the Representational Thesis, I continue to identify qualia with phenomenal properties—those properties that (according to the thesis) an object is sensuously represented (represented$_s$) as having. This means that questions about another person's (or animal's) qualia are questions about the representational$_s$ states of the person (or animal), questions about what properties these states have the function$_s$ of indicating.

As we saw in chapter 2, §2, the way an experience represents$_s$ an object is the way that object would be if the representational system were working right, the way it is supposed to work, the way it was designed to work. In the case of conventional representations, we look to the purposes and intentions of the designers and builders. If you want to know how things "seem" to a ringing doorbell, how the bell is currently representing the front door, ask what would be the case if the device was working the way it is supposed to work. *That* will be how things "seem" to the doorbell. Using this test, when the bell rings, it must "seem" to the system that there is someone at the door. That is what the ring means whether or not there is someone at the door. It does not mean that a bill collector is at the door even if, in point of fact, that is who is there. Given the operation of this system, the bell, even under optimal conditions, cannot distinguish between bill collectors and visiting relatives.[11] That is why we say that the bell means that someone is at the door—not a bill collector, not a visiting relative, but *someone*. A person hearing the bell might represent (i.e., believe) it to be a bill collector, but that is not what the bell means, not what it says or represents. The state of a representational system means whatever it has the function of indicating. That is why doorbells do not "lie" about who is at the door. They never say who is there.

Arthur's visual system cannot represent something to be a poodle for the same reason a doorbell system cannot represent someone at the door as a bill collector. The reason it cannot "seem" to a doorbell that a bill collector is at the door is the reason it cannot seem to Arthur as though there is a poodle nearby. This is not—I emphasize, *not*—because Arthur cannot distinguish poodles from bulldogs. For we can imagine Susan's eyesight so bad that she, too, cannot distinguish one from the other. Yet, though Susan could not, thus impaired, distinguish poodles from bulldogs, we can still imagine her having poodle qualia. We can imagine Susan dreaming of poodles.[12] We can imagine her hallucinating poodles. We can even imagine Susan, with vision so impaired she cannot distinguish poodles from bulldogs, hallucinating poodles when she looks at dogs—thus having poodle qualia when she sees poodles. This is what the wine example shows. We can imagine someone with virtually no ability to discriminate wines experiencing an exquisite

Burgundy taste when he or she drinks wine—any wine. Another person always experiences the taste of cheap Chianti. No, the key difference between Susan and Arthur is not the difference in their discriminatory powers. We can imagine their discriminatory powers being the same and, yet, their having quite different experiences. The key difference is not in what information their visual system provides (this might be the same), but in what information their visual systems have the function of providing.[13] Susan occupies perceptual states that, whether or not they can any longer perform their function, whether or not they enable her to distinguish poodles from bulldogs, nonetheless have the function (i.e., the $function_s$) of distinguishing middle-sized objects (like poodles and bulldogs) from one another. Arthur occupies no such state. That is why he can no more experience poodles as poodles than standard issue doorbells can represent bill collectors as bill collectors.

Imagine two instruments, J and K. J is a precision device manufactured to measure speed in hundredths of miles per hour. When it registers 78.00, that (as the digits after the decimal point suggest) means not 77.99 (and below) and not 78.01 (and above). J has a discriminating speed "palate." Instrument K is a less expensive device, designed to provide rough information about speeds. Its registration of 78 means not 77 (and below) and not 79 (and above).[14] There is no state of K (as there is of J) that has (by design) the job of indicating that the speed is between 77.99 and 78.01 speeds. Since no state has this function, no state of K represents the speed as being 78.00. K's speed "palate" is much less discriminating than J. J is a speed connoisseur. He can tell a 77.92 speed from a 77.98 speed. All these speeds "feel" the same to K. J has speed "quale" that K never "experiences."

J is the instrumental analog of Susan, K of Arthur. In representing the speed as 78.00 mph and 78 mph respectively, J and K are responding to the same objective condition just as both Susan and Arthur are responding to the same poodle. But Susan represents this object in a more discriminating way. Small changes in shape make a difference to Susan. Just as J's registration would change (from "78.00" to "78.05") if the speed changed from 78.00 to 78.05, Susan's representational "needle" would move if the objective shape changed from that of a poodle to that of a bulldog. Just as J's mechanism is sensitive to, and was designed to be sensitive to, these differences in speed, Susan's visual system is sensitive to, and (we are assuming) was designed to be sensitive to, these differences in shape. This is not true of K and Arthur. Neither of them would register, nothing in them was designed to register, these differences. Arthur represents poodles the same way he represents bulldogs—as a blob—just as K represents a speed of 78.00 the same way it represents a speed of 78.05—as 78.

We can imagine our two instruments "experiencing different qualia" (i.e., representing speed differently) while being indistinguishable in discriminatory behavior and capacity. Suppose speedometer J, via damage or wear, loses sensitivity. It becomes as insensitive to speed as K. Speeds that, before damage, "seemed" (i.e., were represented as) different to J now seem the same. After damage to J's information delivery systems, the two instruments provide the same information about speed. As the speed varies between 77.84 and 78.23 J's needle (just like K's needle) doesn't budge (it may help to think of these as digital devices). When the speed rises from 77 to 78, J (like K) sluggishly responds with a shift from "77.00" to "78.00" at approximately 77.50. The two instruments thus become functionally

indistinguishable. Despite this equivalence, a registration of "78.00" on J means something different from a registration of "78" on K. The fact that J no longer delivers the information it is its function to provide does not mean that it loses the function of providing it. As long as it retains this function, a registration of "78.00" on J means something different from a registration of "78" on K even if the two responses carry the same information. After damage (injury, old age, whatever) things still "seem" like 78.00 mph to J. That is what a reading of "78.00" has (and retains) the function of indicating. J can no longer (veridically) "perceive" a speed of 78.00 mph, but J can still "hallucinate" or "dream" about this speed. K cannot. K never experienced 78.00 mph qualia and still doesn't. He cannot "dream" about 78.00 mph speeds. There is no state of K that has this representational content. The representational difference between J and K lies, not in what they do, not even what they can do, but in what their various states have the function of doing.

The difference between J and K is that J does, while K does not, have an internal state that was designed to indicate a speed of 78.00, an internal state that, therefore, has this function. J, therefore, can represent the speed to be 78.00 mph. K cannot. The fact that when things are not working properly the internal states of the two devices indicate the same thing does not imply that they *mean* the same thing. Damage changes what information a state carries (what it correlates with) but not what it means (what it has the function of correlating with). This is why we can imagine Susan, with impaired vision, seeing every dog as a poodle, but not Arthur. She has sensory states that mean this even if they no longer carry this information.

I agree, therefore, with Shoemaker (1991, p. 508), who agrees with Ned Block and Jerry Fodor (1972), that qualia