

Objective Perception

I PERCEIVING OBJECTS AS SUCH

I want to distinguish and characterize what I take to be a special case among broadly sensory or perceptual phenomena. Some might maintain that this case is the only genuine case of perception properly so called; but I will sidestep that largely terminological issue, and simply specify the kind of case I mean as *objective perception*—not worrying about what to call other kinds of case. It is part of my view, though not part of what I will argue for, that objective perception, in this sense, is exclusive to people. This is not to deny that, in some other or wider sense or senses, animals perceive too; nor is it to deny that people also perceive in some or all of these other senses; nor, finally, is it to deny that some other senses of perception may also be distinctive of people (such as aesthetic perception). It is only to claim that at least some human perception is “objective” in a special sense, and no animal perception is.¹ Though I will not be defending this claim, I will have it always in mind, and will occasionally take it for granted in choosing illustrations.

The qualifier ‘objective’ is intended to suggest perception of *objects as objects*. Thus, it is very much part of the undertaking to spell out what is meant by the *object* of a perceiving. It should be clear at the outset, therefore, that ‘object’ is used in a formal sense, to identify a role or position vis-à-vis (for instance) perception; in particular, it does not carry any implications about the nature of the objects perceived—such as that they be substantial, corporeal, or otherwise thing-like. Objectivity in perception is a kind of structure that involves the perceiving, that which is perceived, and the relation between them. The aim is to delineate this distinctive structure and its presuppositions.

Several familiar problem areas in the philosophy of perception will be seen to converge in this topic, among them: (i) picking out what is perceived from among the various causal antecedents of the perceiving; (ii) the normativity of objective perception (that is, the possibility of *mis*perception); (iii) the respect in which objective perception depends on an understanding of what is perceived; (iv) the relevance of language to the possibility of objective perception; and (v) the prerequisite character of the perceiving self or subject. On the other hand, a number of other important issues connected with perception will not enter directly into the present discussion, including in particular: (i) the peculiar “of-this-ness” or indexicality of perceptual content; (ii) the richness of (apparently) ineffable yet determinate detail in perceptual content; and (iii) the special connection (if any) between perception and imagery or imagination.

2 DRETSKE ON THE OBJECT PERCEIVED

It is convenient to begin with the identification of what is perceived—the “object” in a broad sense—from among the many causal antecedents of the perceiver’s perceiving. We can take for granted that, at least in ordinary sensory perception, the perceived object is *one* (or some) of those antecedents. The question is: *Which* one? Or, rather, the question is: On what principled basis can we identify one among them as “the” object? For instance, if I look at and see a bicycle, part of the cause of that perceiving in me is the presence and properties of the bike—particularly those properties that determine in detail how it reflects light. But other important causal factors include the ambient light around me (as reflected in part by the bike), the source of the illumination of the bike in the first place, the source of the bike itself, the functioning of my visual system, and so on. We want to say that what I perceive, the object of my perception, is the bicycle, and none of these other factors. The question is: *Why?*

Dretske (1981, chapter 6, especially pages 155–168) suggests that the object of a perceptual state is whatever that state carries information about in a primary way. “Carrying information” is the operative notion. Roughly: given conditions of kind *K*, *a*’s being *F* carries the information that *b* is *G* just in case *a* could not have been *F* in such conditions had *b* not been *G*. To put it anthropomorphically, *a*’s being *F* carries the information that *b* is *G* if you can “tell” from *a*’s being *F* (and the conditions being *K*) that *b* is *G*. It carries this information in

a *primary* way if it does not carry it “via” any other fact, such as *c*’s being *H*. Intuitively, to say that *a*’s being *F* carries the information that *b* is *G* *via* *c*’s being *H* means that it carries that information *only because* it carries the information that *c* is *H*, and *c*’s being *H* carries in turn the information that *b* is *G*.²

Dretske illustrates the point by explaining why, when someone rings the doorbell, we hear the bell ringing but we don’t hear the button being pushed. Under normal conditions, our perceptual response—what Dretske calls the “perceptual experience”—carries both the information that the button is pushed and the information that the bell is ringing. This is because we wouldn’t have that experience if the button weren’t pushed, nor would we have it if the bell weren’t ringing. But the *only reason* we wouldn’t have it if the button weren’t pushed is that we wouldn’t have it if the bell weren’t ringing, *and* (under normal conditions) the bell wouldn’t ring if the button weren’t pushed. Hence, the information about the button is carried *via* the information about the bell.

To see how this works, consider the (abnormal) conditions in which the bell wires occasionally short, thereby ringing the bell without the button being pushed. Under those conditions, the same experience would still carry the information that the bell was ringing, but it would no longer carry the information that the button was pushed—precisely because the bell’s ringing would itself no longer carry that information. So, even in normal conditions, when the perceptual experience does carry the information that the button is pushed, it does so *only via* the bell’s ringing. Hence, the experience does not carry the information about the button in a primary, but only a secondary way. And this is why, according to Dretske, the button is not the object of the perception. Even though we can *tell* that the button is pushed, we do not *perceive* that it is (unless we happen also to be outside watching). What we *hear* is not the button but the bell.

The more interesting and difficult question is why the experience carries even the information about the *bell* in a primary way. It might be argued, for instance, that the experience also carries the information that the air or our eardrum is vibrating in a certain manner, and that, moreover, it carries the information that the bell is ringing *only via* this information about the air or eardrum.³ In that case, only the latter information would be carried in a primary way; and then, by the above logic, it would have to be conceded that what we *really* hear,

even under normal conditions, is the air or our eardrum, rather than the bell.

Dretske's reply to this challenge is ingenious, but ultimately, I think, unsuccessful. It has two stages. In the first stage, he points out that one thing's being the case can carry the information that another is, without carrying any information about the intervening causal processes, if (in the relevant conditions) any of several different causal processes might have occurred, and each would have had the same result. Thus, suppose that the mantle bowl only has lemon drops in it when one of the children buys some, and that the children only buy lemon drops when Grampa provides the funds. Then, if there are lemon drops in the bowl today, that fact carries the information that Grampa recently paid, but it carries no information as to *which* grandchild fetched—because they're all equally eager and able. *A fortiori*, it can't carry the former information *via* carrying the latter.

The second stage, then, is to argue that the different ways the air or eardrum might vibrate when we hear the bell ring are like the different grandchildren that might fetch the candy for Grampa. That is, the perceptual experience must have been caused by some vibration or other that could only have been caused (in the relevant conditions) by the bell ringing; but there are any number of such vibrations that would have sufficed with no difference in effect. Under those conditions, the experience would carry the information that the bell is ringing without carrying any information as to *which* of the sufficient vibrations mediated causally. Thus, the distal information (about the bell) is not carried *via* any more proximal information (about how the air or eardrum is vibrating), for the simple reason that the proximal information is not carried (by that experience) at all. And this is why it is the bell which we hear after all.

How is this argument made out? Dretske appeals to perceptual constancy effects, which, as he rightly points out, are ubiquitous. We do not perceive the table as changing shape when we walk around it, changing size when we approach, or changing color as daylight gives way to dusk and then to candles, even though the proximal stimuli—the patterns of light entering our eyes or the patterns of neuron firings in our retinas—vary dramatically. Thus, we see the table as equally square, whether our vantage is such that its retinal projection is itself square, rhomboid, or trapezoid. So, (under normal conditions) the perception carries the information that the table is square, without carrying any information as to *which* retinal projection happened, on

this occasion, to mediate the perceiving causally. Accordingly, it carries the information about the table in a primary way—not via carrying information about any proximal stimulus—and hence it is the *table* which we see.

The essential difficulty here lies in the characterization of the proximal stimulus. Note that the argument depends on the claim that the *same* perceptual experience of the *same* object can be mediated by *different* stimuli; for, if only one stimulus could mediate, then the experience would carry information about it, and hence the information about the object would not be carried in a primary way. For purposes of discussion, we can accept the suggestion, based on perceptual constancy, that in some sense the experience and the object remain the same; the question concerns the sense in which the respective proximal stimuli are different. Of course, it's *qualitative* difference—difference in kind—that's at issue; and whether two instances differ in kind depends on which kinds are being considered.

Which kinds need Dretske consider? Since the argument depends on a nonexistence claim, it must, in effect, consider *all* kinds. For if there were *any* single kind of stimulus that mediated all and only the constant perceivings (same kind of perception of the same kind of object), then the perception would carry the information that the stimulus was of that kind, and hence would carry the information about the distal object only *via* the information about the stimulus. Dretske points out that there are respects in which the stimuli differ: some are rhomboid projections, some trapezoid, and so on. But that isn't enough; he must argue that there is *no* respect in which these stimuli (and only these) are all alike.

Of course, such a kind would not be as simple as the shape of an instantaneous optical projection on the retina. And Dretske even mentions (1981, 164), as plausible explanations of constancy phenomena, the existence of “higher order” variables in the stimulus array” and “global characteristics of the entire stimulus pattern” to which our sensory systems are sensitive.⁴ Unaccountably, however, he never considers the possibility that such “higher order” and more “global” stimulus *kinds* might undermine his account of why the stimulus itself is not perceived. That is, he never considers the possibility that the perceptual response carries the information that the (proximal) stimulus is of such a (higher order, global) kind, and thus carries information about the (distal) object only via that information about the stimulus.⁵

What's worse, it seems that there *must* be such kinds, if sensory perception is to be possible at all. For if one *can* reliably recognize the squareness of the table from varying perspectives, then there must be *something*—something higher order, global, relative to context, or whatever—normally common to all and only the stimuli from such objects, on pain of rendering perception magical. To be sure, perception is not 100% reliable, in part because stimuli of the relevant kinds can be produced artificially (or accidentally) even when the corresponding objects are not present. But this is no help to Dretske. Quite the contrary: it not only suggests that there *are* the stimulus kinds that his account can't allow, but also that perceptual responses track these kinds primarily, and the object kinds only *via* them.

3 THE OBJECTS OF CHESS PERCEPTION

So far, the question has been: Why, despite being *via* (proximal) stimuli, is perception of (distal) objects. The answer, as we have seen, cannot depend on the absence of suitable kinds for the stimuli to instantiate. Thus, presumably it must depend instead on something positive about the objects themselves, and about perceiving them. It is crucial, however, not to suppose at the outset that *objects* are an unproblematic fixed point, and the only issue is how perception gets to be of *them*. In particular, we must not tacitly presuppose that being an object is tantamount to being a temporally and spatially cohesive corporeal lump. Rather, the "objecthood" of perceptual objects and the "of-ness" of perception go hand in hand, and are intelligible only in terms of one another, something like the interdependence of target and aim, or puzzle and solution. So, the deeper question is: *How and why* is such a structure—what we might call *the structure of objectivity*—imposed on the physics and physiology of sensation?

In order to avoid the prejudicial presupposition that we know in advance what "objects" are, I shall turn temporarily to a different sort of example, one in which the temptation to equate 'object' with 'body' is perhaps less compelling. In this context, it will be possible to ask what it is about the objects that lets them be objective, without the distraction of supposed "obviousness". That is, the question can really be confronted and maybe even answered. On that basis, then, we will be able to return to the case of ordinary corporeal "things" with clearer horizons.

The special case I want to consider is chess, when the game is played in a visible medium—that is, such that one can *see* positions, threats, moves, and the like. I can, for instance, see you castling, early in the midgame; a little later, I can see your knight simultaneously threatening my queen and rook, in a so-called knight fork; and, before long, I can even see the untenability of my entire king-side defense, and conclude that I might as well resign. We should not try to imagine that these perceptual abilities are all built up out of a handful of “primitive” or “atomic” abilities, such as identifying the individual pieces and squares. I can recognize a knight fork just as “immediately” as I can recognize a knight or a square; and it may not even be possible to define ‘untenable king side’ in terms of pieces and locations—suggesting that they can *only* be perceived as gestalts.

What’s involved in seeing these objects? An instructive first indication is the fact that dogs and cats can’t see them at all. Without broaching, for the moment, the question of what exactly they *do* see, we can be quite sure that no dog or cat could see even a rook as such on the k2 square, never mind the power, daring, or foolishness of that position. And the reason is not far to find: dogs and cats don’t have a clue about chess. The very possibility of such a game, and all that it entails, is utterly alien to them. We might gather, then, that some grasp or understanding of the game of chess, and maybe also of games more generally, is prerequisite to any ability to perceive the phenomena manifested in chess play. But it will be worth the trouble to proceed more slowly, and explore what such a grasp amounts to, and why it is required.

Surely no creature or system can see a given configuration as a knight fork without having some sense of what a knight fork is. To put it in a familiar but perhaps misleading terminology, nothing can *apply* a concept unless it *has* that concept.⁶ Of course, having a sense (or “concept”) of what something is cannot be merely an ability to detect it reliably, or discriminate it from other things; locks, for instance, reliably detect the keys that fit them, and it’s not out of the question that pigeons could be trained to discriminate (for a given chess set) positions containing knight forks.

In the case of chess, however, even the discrimination problem is harder than that; for chess sets can be implemented in all sorts of styles and media, from labeled slips of paper or computer-screen icons to costumed palace servants or elaborately painted helicopters. Knight

forks are equally knight forks in any of these sets; but, to see them, a perceiver would have to “abstract from” the gross and conspicuous differences of implementation, and attend only to the chess-relevant features of the configurations.

Indeed, the matter is still more complicated, because other games, quite different from chess, can be played with the board and tokens that we ordinarily call chess sets. And it could well happen that, in the course of such a game, an arrangement occurs that would be a knight fork were chess being played. But, in fact, it’s nothing like a knight fork, because, in this game, the pieces move (or are used) completely differently—the token that would typically implement a knight in chess doesn’t implement a knight in this domain at all. So, in order to see that a certain configuration is a knight fork, a perceiver must first be able to see, or otherwise appreciate, that it has occurred in the midst of a *chess* game.

These complementary considerations, that chess can be played in widely different media, and that widely different games can be played in the same media, together with the fact that knight forks can occur (and be perceived) in all of the former but none of the latter, show that the ability to perceive knight forks presupposes some grasp or understanding of the game of chess—at least enough to tell when it’s being played, regardless of medium. The same point is even more obvious for seeing that a king-side defense is untenable; and a little reflection shows that it is equally true for seeing that some piece is a knight (a *chess* knight), or anything else that occurs only in chess games. Chess games are a kind of *pattern*, and chess phenomena can only occur within this pattern, as *subpatterns* of it. The point about different media and different games is that these subpatterns would not be what they are, and hence could not be recognized, except as subpatterns of a superordinate pattern with the specific structure of chess—not a pattern of shape or color, therefore, but a pattern at what we might call “*the chess level*”. To put it more metaphysically, the game of chess is *constitutive* for chess phenomena, and *therefore* some grasp of it is prerequisite to their perceivability as such.

The pivotal phrase, “some grasp or understanding”, has so far been left dangerously vague. It is all too easy to gloss this grasp as: knowing the rules (including the starting and goal configurations) that define the game—what one might get by reading Hoyle or the instructions on the box. But there are several reasons to resist this temptation. In the first place, it’s not obvious that reading (and understanding) the

rules is the same as, or even sufficient for, understanding the game. Second, and more to the point, it's not obvious that *knowing* the rules is required for understanding the game, at least not if knowing is taken in a sense that implies discursive cognizance of the rules in some explicit formulation (such as one gets by reading). But third, and most important, chess is serving here just as an example; the hope, without which the example would not repay the effort, is to generalize its lessons to less specialized cases—in particular, to cases for which, though there are constitutive prerequisites, these prerequisites have never been (and perhaps cannot ever be) articulated in any “rules according to Hoyle”.

4 CHESS OBJECTS WITHOUT LANGUAGE

Too often, I believe, philosophers take it on faith that what *essentially* distinguishes people is language—“essentially” in the sense that all our other interesting differentia (society, morality, objectivity, self-understanding, history, science, normativity, sense of humor, ...) flow from, or at least necessarily involve, this one. It seems to me, on the contrary, that some of these—objectivity, in particular—are fundamentally independent of language, and that we misunderstand them if we overlook this independence. Accordingly, I want now to indulge in a brief thought experiment, by imagining some creatures that are, in certain respects, intermediate between known animals and people.

These creatures—“super-monkeys”, we could call them—are sub-human in that they have no linguistic facility whatever: they neither produce nor understand speech (nor any other articulate gestures or signs to the same effect). This is not to say they never cry out in pain, make warning signals, or give mating displays; but they never, for instance, tell stories, describe things, or formulate rules. On the other hand, they stand out from all other animals in that they learn and play games, including sophisticated games like chess.

Since super-monkeys are animals, and in fact a lot like us, there is no problem in ascribing to them, not only various beliefs and desires (at least in the minimal sense of Dennett's “intentional systems theory”—1981/87), but also a considerable range of affects and moods, emotions and motives. It's perfectly obvious, for instance, that they *like* strawberries (as well as the janitor who brings them), that they *enjoy* frolicking in the pool (but would sometimes rather sleep), that they get *angry* when picked on, because it *hurts*, and that's *why* they

strike back. These qualities sharply differentiate super-monkeys from any current or imminent generation of “game-playing” computer.

What should we look for in super-monkeys, if we are to understand them as playing games, even though they don’t talk about it? Well, of course, they must *at least* reliably go through the motions; that is, they must actually move the pieces, not merely in accord with the rules, but in a manner more or less suited to winning. And, in the spirit of the above remarks about emotions and motives, we can expect them to be pleased and self-satisfied when they do win (or gain ground), and dejected or frustrated at the opposite. This shows that, in one sense, they are not *just* going through the motions, for the outcomes evidently matter to them. But it does not yet show that the motions are specifically *chess* moves, for it has not been shown that they are made as and only as subpatterns of a superordinate pattern at the chess level. Such further specificity is not hard to add, however: we merely require that our super-monkeys also be able to play chess in various media, and other games in the same media. This is not at all a trivial or inconsequential requirement; but it does not seem to presuppose linguistic facility, so it remains within the parameters of the thought experiment.

Super-monkeys, as described so far, are clearly guided in their moves by sensory input from the chess board; and so, in at least that limited sense, they are certainly perceiving. When we ask, however, with Dretske, what the perception is *of*, the situation is more complicated. On the one hand, since the *responses* involved are not limited to (and may not even include) inner experiences, but are rather themselves also movings of the chess pieces, we have, so to speak, a new “angle” on the *objects*: they are *both* perceived *and* acted on. Of course, just as we can ask about perception, why it is of a distal rather than a proximal cause, we can ask about action why it is of a distal rather than a proximal effect—why it is an action of moving a knight, for instance, rather than of moving an arm and some fingers. But the fact that the piece on the chess board is the one place where these two causal trajectories intersect, and especially if there is continuous feedback, as in visually guided movement, strongly suggests that it is the proper object of both the perception and the action. It is this feedback linkage, I believe, that inclines us to identify, to the extent that we do, the “objects” of ordinary animal perception and action.

On the other hand, we should notice the restricted scope of this feedback argument. It applies to a chess piece only insofar as it is a

sensible, manipulable token, and has nothing whatever to do with its role or character in chess. In other words, the question of whether chess phenomena can be the objects of perception (or action), even for super-monkeys, has not yet been touched. And, in fact, the earlier argument brought against Dretske's information-based proposal, applies here as well. For, if chess phenomena are the phenomena they are by virtue of the way they are subpatterns within a superordinate pattern at the chess level, then, *inevitably* (under normal conditions), the proximal stimulus patterns are instances of corresponding kinds, according as they are subpatterns of higher order and more global patterns in the proximal stimuli overall. These latter patterns will have to be comparable in "level" to the chess level, in that they abstract from features that are characteristic of the various different media, and depend instead on the global relationships that define chess games as chess. But there is no reason to deny the existence of such patterns, and hence of such stimulus *kinds*.

5 LETTING THE STANDARDS GOVERN

There is, however, one more requirement that super-monkeys will have to meet if they are to be recognized as genuine chess-players; and this last requirement will prove crucial also for objectivity. It can be introduced by asking what happens in the face of an *illegal* move. We have already required, of course, that, as candidate chess players at all, super-monkeys must *reliably* make legal moves; and in undertaking to play they must, in effect, be counting on this reliability in one another. This cannot mean, however, that they have to be perfect, or that they can take perfection for granted, in the sense of being unprepared to deal with illegality.⁷ No system that would blithely ignore the illegality of an opponent's move could properly be deemed a chess player.

Hence, any genuine chess player—in fact, any game player—must, as such, occupy a peculiar position. A player must, on the one hand, be ever *vigilant against* illegal moves, yet, on the other hand, always *count on* consistent legality.⁸ These two are more intimately related than they might seem. Let's examine them in turn.

As vigilant, a player needs two quite distinct abilities: (i) to tell in general what move is made (including any relevant context),⁹ whenever a player makes a move; and (ii) to tell, for any identifiable move, whether that move would be legal.¹⁰ These abilities together (but neither of them alone) enable players reliably to detect illegal

moves—an obvious prerequisite to vigilance. The first ability, in effect, induces a field of recognizable moves (that is, identifiable phenomena that are either legal or illegal); and the second induces a partition on that field (into the legal and the illegal). Clearly, the field of identifiable moves had better include candidates on both sides of that partition, lest the vigilance be vacuous. That is: it must be possible to identify moves that would be illegal—moves, in other words, that are *ruled out* by those very rules that are constitutive for chess phenomena and their identification at all. To say that these abilities are distinct is to say that they have within them this essential possibility of *discord*.

Counting on consistent legality means more than just expecting it, but something like *insisting* on it—on pain of giving up the game. This insistence is a kind of *commitment* or *resolve* not to tolerate illegality: in case an illegal move is detected, ordinary play is breached, extraordinary measures are called for—and the game itself (that is, continued play) is at stake in the outcome. The vigilant player's insistence, therefore, is a commitment not to stand for the kind of discord (between the above two abilities), the possibility of which is prerequisite to vigilance in the first place. So the extraordinary measures, when called for, will be an attempt to eliminate this discord, by rectifying either the identification of the move or the determination of its legality. As we shall see, it is this commitment to the definitive standards constitutive for the domain as such (for instance, to the rules of chess) that transmits normative force (including the potential for correction) to the move identification ability, and hence underlies the possibility of its objectivity. Hence, the capacity for such commitment is a prerequisite character of any self or subject that is to be capable also of objective perception (or thought or action).

Suppose, for example, that Bobby Baboon is in the midst of a chess game when his "illegal-move alarm" goes off. As a genuine chess player, Bobby is antecedently disposed to take extraordinary measures in the face of such an event, and to adjust accordingly. And among those measures will be to "double check" his identification of what move was actually made. The details of this double checking are not important. But we can imagine Bobby looking over the board again, perhaps from different angles, checking for overhanging branches or bad lighting obscuring part of the board, making sure that what he took for a bishop isn't really a knight seen from the back, and so on. We can even imagine that Bobby is still getting used to an unfamiliar

chess set, and sometimes has to remind himself (or be reminded) which pieces are which. (Basically similar things could be said about Bobby “double checking” or “reconsidering” his assessment that the move was illegal—perhaps he forgot about capturing *en passant*, for example.)

It is crucial that this double checking be an effort, and that “success” not be automatic. If Bobby turned up a “misperception” every time he looked for one, then the *checking* would be a sham. Rather, we must suppose that the double-checking is an extension, elaboration, or more careful mode of the primary recognition ability, together with an ability to tell whether the result is coming up different, and to adjust if so. (Still more impressive would be a capacity to diagnose what went wrong in cases where a misperception is uncovered; but I don’t see that this is required.) The enabling assumption is that, almost always, this more careful mode is unnecessary, because it wouldn’t make any difference. But, occasionally, it would make a difference; and, in those cases (often flagged by the alarm), the result of the more careful mode is almost always *better*.

What does “better” mean? It means: more conducive to continued chess play, in the long run. Playing a game governed by constitutive standards is a nontrivial *achievement*; dogs and ordinary monkeys, for instance, are utterly incapable of it.¹¹ The mark of such an achievement is not an ability to articulate those standards as explicit rules, nor is it a disposition to go through any particular motions of play. Rather, it is the ability to play in fact, while at the same time putting the game itself at stake in insisting on the standards; that is, it is the ability to *let the standards govern* the play. Thus, the aforementioned possibility of discord—a possibility which, on the one hand, is reliably avoided, but, on the other hand, is resolutely not tolerated—is fundamental. Realizing a set of dispositions that works with this structure is the achievement; and the fine tuning effected in the double checking is best to the extent that it fosters this achievement.

6 OBJECTS DETERMINE PERCEPTUAL CORRECTNESS

Notice that, almost incidentally, we have now assembled the resources needed to answer Dretske’s question. What the perception is of is that which the constitutive standards govern: the moves, pieces, positions, and so on, in the chess game. Why? Because, when the possibility of a *misperception* arises, what the issue turns on is telling how things

stand with regard to those standards. That's what *accurate* perception is counted on to tell; it's what is checked for in the double checking; it's the way in which perceptual ability contributes to the game-playing achievement. In other words, the *norms* governing the perceptions as such, and in virtue of which they can be objective, are inseparable from the *standards* governing, and indeed constituting, the chess phenomena as such; or, to make the Kantian paraphrase even more obvious: the conditions of the *possibility of objective perception* as such are likewise the conditions of the *possibility of the objects of that perception*.¹²

To be concrete, suppose that, on a certain occasion, the actual chess position is of one kind, but (due to some quirk in the circumstances) the ambient optic array around Bobby, and hence also his perceptual response, is of the sort usually caused by chess positions of some other kind. Is his response a misperception of the chess position, or an accurate perception of the ambient array? Clearly the former—*because* what's at issue in the perceiving at all, the whole point of *looking* in the first place, is telling what the position is. And this is manifest in the fact that, if Bobby's alarm goes off and he discovers the discrepancy, he will change his response; that is, to the best of his ability, he will bring his response into line with the actual position, rather than the optic array (the entire aim and purpose of double checking). Therefore, even when there is no discrepancy, it is the position and not the array that is perceived.

This shows, I believe, that our imaginary super-monkeys are capable of objective perception—like people, and unlike any (actual) animals. The fact that these super-monkeys are completely nonlinguistic shows further that language is not prerequisite to objective perception, and hence that it is not language that fundamentally separates us from the animals with regard to the possibility of objectivity. On the other hand, the fact that putative perceptions are normatively beholden to their objects, subject to correction in light of double-checking those objects, is integral to the account. What's more, these norms are completely dependent in force and content on the constitutive standards to which the perceived objects are held, and which must therefore be counted primary. And it is, finally, the super-monkeys themselves that “hold” those objects to those standards. Such holding to standards, by simultaneously counting on it and insisting on it, is, when it succeeds, understanding. Thus, the objects of objective perception are *ipso facto* understood. *Understanding*, not language, is what separates super-monkeys (and us) from “thoughtless brutes”.

7 CHESS WITHOUT LANGUAGE (REVISITED)

That there might actually be “super-monkeys” that learn and play games like chess without benefit of language is, to be sure, practically incredible. Hence, the intended point and force of the example calls for explanation. It will help first to rehearse a few notable objections. Thus, it is extremely unlikely that the cognitive and social capacities prerequisite for chess would evolve prior to and independently of those required for language. For, clearly, there is a great deal of overlap in these capacities; and they would have far more survival value if used for language than if used only for formal games. It could even be that our own ability to play chess relies essentially on species-specific “modules” that are integral to our linguistic faculty—such as our ability to parse complex structures, or to see individual tokens as tokens of digital types, or to remember a large “vocabulary” of possible constituents.

Likewise, it is almost impossible to imagine *teaching* chess, even to a creature with all the native capacities of *homo sapiens*, without verbal instruction and admonition. How could one hope to convey the relevant alternatives and restrictions in various situations without explicit conditionals and modal adverbs? Indeed, if genuine chess presupposes not just that the moves in fact be legal, but that the players insist on this, then they had better have a way of expressing their insistence. Expressions of insistence, moreover, cannot be the same as mere expressions of displeasure, dismay, or disapproval; for those could have any number of possible grounds, whereas insistence (as used here) can be grounded only in issues of constitutive legality. But that means that expressions of insistence must be somehow marked as such, marked as concerned with issues of legality, so as to communicate specifically that concern. And wouldn't such specifically marked, communicative expression be tantamount to language?¹³

Objections based on actual or plausible evolutionary history, however, are beside the point. Thus, it might also be that chess-like games could not have emerged until after prehensile hands; but, even if so, that wouldn't say much of interest about the games as such. The arguments that matter, therefore, are those to the effect that chess itself presupposes language, either for learning it or for playing it. But to those, I think, a simple reply is decisive. It is certainly no harder to learn and play chess than it is to learn and speak a natural language. Quite the contrary: games are clearly much less demanding than

languages on all counts. In particular, languages are just as constituted by standards, hence just as dependent on speakers' *insistence*, as any game. Yet, it must be possible to learn and speak a language without benefit of (any other or prior) language, on pain of regress. So, in principle, it's possible for games as well.¹⁴

There remains the idea that chess itself might *be* a language. It goes without saying that languages and games are similar in many ways—to the extent that it has become a philosophical cliché to refer to languages as “games”. Nevertheless, there are also many differences, including the following, which is both easy to see and crucial. Chess, like all formal games, is *self-contained* in the sense that what moves would be proper (legal) on any given occasion depends only on factors that are *internal* to the game—previous moves, the locations of other pieces, and the like. By contrast, for any genuine language, what it is proper to say on a given occasion—what is true, authorized, materially valid, and so on—depends in general on more than the specifically *linguistic* context. Therefore, no formal game is (in and of itself) a language.

Why belabor the point? I certainly do not deny that language is characteristic of people, and centrally important to what we are. But I want to resist what strikes me as an all too prevalent tendency to make of language a magic fulcrum on which every human quality turns—as if language were the whole essence of what distinguishes us from machines on the one hand, and beasts on the other. That, I believe, is a grave distortion, for many reasons, most of which go beyond the present discussion. Even in the context of a theory of perception, however, it is important not to overestimate the role of language (if only because *correct* perception is prerequisite to dialog itself).

Playing chess, like speaking, involves interacting with items and structures in ways that depend—not just causally but normatively—on their types (roles), their reidentification as individuals (histories), and their relations to other items and structures (“contexts”). The ability to engage in such interaction is at least the greater part of what is meant by a *conceptual* grasp. Granted, a person who can talk and theorize might have more and richer concepts of chess phenomena than could a nonlinguistic super-monkey. Yet the basic concept of, say, a rook is determined by its role in the game; and any creature that can play must have mastered that. Concepts are in general articulated by their relations to one another, as expressed in the detailed contin-

gencies of acceptable practice. Often these contingencies are primarily linguistic, and, in particular, inferential. But, as the chess example shows, they need not be: proper chess play itself, without verbal accompaniment, is sufficiently structured to institute a conceptual articulation of the corresponding phenomena.

8 SCIENTIFIC COMMITMENTS AND OBJECTIVITY

It must be acknowledged that chess is not a typical example of a perceptual domain. The question therefore arises whether the account of objectivity, in terms of constitutive standards, insistence, and achievement, might be limited to games and their ilk. We ask, that is, whether the account depends essentially on any features peculiar to such domains, or whether it depends only on features characteristic of objective perception more generally. The following stand out as distinctive of chess and the like—in contrast, for instance, to scientific observation and everyday perception of sticks and stones.

- 1 Chess is defined by constitutive standards (rules) that are arbitrary human inventions; they do not have the character of empirical discoveries, liable to refutation. (Even so, the rules of a playable game must be consistent, complete, and followable—an achievement that is not at all automatic.)
- 2 When chess players insist upon legal moves, this is understood primarily in terms of rule-compliance by other players—agents who can be held to account for their behavior. (Less often an issue, though just as important, is insistence that the board and pieces function properly—by not melting, wandering around, or such like.)
- 3 It is almost always easy to tell what move has been made in a chess game, and whether that move is legal—because the game is digital, explicitly defined, and relatively simple. (On the other hand, further perceptual skills, such as recognizing strategic weaknesses and opportunities, are not easy and may well not be reducible to these basic discriminations.)

Science, by contrast, seems anything but arbitrary an invention: its discoveries are a paragon of the empirical and refutable. Moreover, scientists never hold observable phenomena to account for their behavior: if some observed phenomenon fails to accord with scientific

expectations, it is the observation or expectations that are at fault, not the phenomenon itself. Finally, scientific investigation is difficult: it takes years of training plus painstaking effort, both to perform reliable experiments and to tell with confidence what results are acceptable.

This, however, is not to deny that the objects of scientific study are held to standards. Donald Davidson (1970/80, 211 and 219–222), for instance, more or less defines the *physical* as that which can be picked out in terms drawn exclusively from a closed, comprehensive theory with strong constitutive elements and precise, explicit, exception-free laws.¹⁵ This is as much as to say that being subsumable under such a theory—*strict subsumability*, we might call it—is a constitutive standard for the physical: to *be* physical is to be strictly subsumable.

Speaking in a similar vein, but of the history of chemistry, Thomas Kuhn writes:

Changes in the standards governing permissible problems, concepts, and explanations can transform a science. In the next section I shall even suggest a sense in which they transform the world. (1962/70, 106)

And, in the next section, he continues:

For Dalton, any reaction in which the ingredients did not enter in fixed proportion was *ipso facto* not a purely chemical process. A law that experiment could not have established before Dalton's work, became, once that work was accepted, a constitutive principle that no single set of chemical measurements could have upset. ... As a result, chemists came to live in a world where reactions behaved quite differently from the way they had before. (133–134)

In other words, the principle of fixed proportions became a constitutive standard for what it is to *be* a chemical reaction—and thereby also for being an element or a compound. Kuhn, however, unlike Davidson, is concerned not merely with the notion of standards for *entities*, but with the implications of this view for the conduct of science.

Finally, at a still higher level, there is another set of commitments without which no man is a scientist. The scientist must, for example, be concerned to understand the world and to extend the precision and scope with which it has been ordered. That commitment must, in turn, lead him to scrutinize, either

for himself or through colleagues, some aspect of nature in great empirical detail. And if that scrutiny displays pockets of apparent disorder, then these must challenge him to a new refinement of his observational techniques or to a further articulation of his theories. (42)

Scientists, that is to say, are *scientists* by virtue of their *commitments*—in particular, commitments that require what we have earlier called *vigilance*. “Pockets of apparent disorder” are nothing other than apparent breaches of the relevant constitutive standards, the exact analog of apparent illegal moves; and scientists must simultaneously be on the lookout for them, and resolved not to tolerate them. Scientists *insist* that the scientifically observed world be orderly (strictly subsumable, for instance). Moreover, their alternatives in the face of an apparent breach are essentially the same as those available to a chess player: refinement of observational technique, further articulation of the theory,¹⁶ or giving up the game (scientific breakdown and/or revolution).

The constitutive standards for the objects of scientific research—whether local to particular disciplines, like combining in fixed proportions, or global throughout science as such, like displaying an order with precision and scope—are not arbitrary inventions. On the other hand, they are not exactly empirical discoveries either; for, as Kuhn is at pains to show, accurate observations and discoveries presuppose them, and they are not readily dislodged. They are somehow *both* empirical *and* invented—“synthetic a priori”, Davidson says (1970/80, 221). Standards for the constitution of objects are worked out by fits and starts over many years, such that, in accord with them, ever more objects can be scrutinized in great empirical detail and ordered with precision and scope. “It is hard”, Kuhn notes (135), “to make nature fit a paradigm”; and a fundamental component in that difficulty is coming up with paradigms that nature can be made to fit. Scientists may invent the recipe; but experiment is the proof of the pudding. Like chess—only far more so—science is an *achievement*.

It is an old problem in the philosophy of science to say just what is measured by a scientific instrument. In particular, in what sense can it be said that so-called “theoretical” entities and properties are measured, as opposed to “phenomenal” or “observational” properties or states of the instrument itself? Notice that this question has essentially the same structure as Dretske’s question about how we can say that we hear the bell, and not the vibrations in the air or our

eardrums. And the solution, it seems to me, is essentially the same too. What the measurement is of is that which the constitutive standards govern: the entities, properties, and relations in terms of which the theory is expressed. Why? Because, when the possibility of a *mismeasurement* or experimental *error* arises, what the issue turns on is telling how things stand with regard to those standards. That's what *accurate* instrumentation is counted on to measure; it's what is checked for in double checking; it's the way in which experimental ability contributes to the scientific achievement. Successfully holding those entities to those standards, in the face of ever more precise and detailed experimental testing, is scientific *understanding*—understanding what the entities *are*.

Philosophers of science speak more often of explanation than of understanding, but these come to the same: to *explain* is to render intelligible—to show that and how something can be understood. Thus it is that explanations are of two sorts. An entire domain of phenomena can be constituted and explained when standards to which they can be held are grasped and successfully insisted upon. (Such a success is a *paradigm*.) And, within a domain, particular phenomena, or particular classes of phenomena, are explained by showing how they in particular can be held to the constitutive standards. In case the pertinent standards include strict subsumability, those internal explanations will be, in part, deductive nomological. In case the domain-constituting standards are otherwise—for instance: the Davidson/Dennett notion of rationality, the integrated operation of functional systems and organisms, or the historical/institutional dynamics of cultures and subcultures—then the explanations and intelligibility will be of different sorts. What they have in common is the structure of *insistence*: practitioners have the ability to recognize phenomena that are ruled out by the standards that are constitutive of that recognizability, and will refuse to accept them.

9 EVERYDAY OBJECTIVITY

Can the same account be extended also to the objectivity of everyday perception, say of sticks and stones? It would be here, presumably, if anywhere, that dogs and cats would have perceptual abilities comparable to our own. No doubt, we share with higher animals various innate “object-constancy” and “object-tracking” mechanisms that automatically “lock onto” medium sized lumps—especially ones that are

moving and/or staring at us. The question is whether, for us in contrast to animals, there is any more to the objectivity of perceivable “things” than that they trigger such mechanisms—in particular, whether there are constitutive standards to which we, as perceivers, hold them.¹⁷ It seems to me that there are, though they are somewhat vague and difficult to formulate. The essential tenet is something like: *things* are integral bearers of multiple properties. Integrity is the way the properties belong together in space and time, in (at least) two respects: cohesively and compatibly.

A thing is a cohesive spatio-temporal unit. At a time, a thing occupies—that is, its properties are instantiated at—exactly one place, a place which has a definite size and shape, and which is occupied by no other thing. Through time, the place of a thing can vary, as can the properties collected in it there. But mostly these variations can be only gradual—that is, relatively slow and continuous. For otherwise, it would cease to be well-defined which properties belong together as properties of which things; their staying identifiably together in one continuing thing is what fixes their belonging together at all.

The properties of a single thing must always be mutually compatible, and they can be interdependent; that is, some combinations of properties and/or lacks are not permissible in one thing. But no properties or lacks in distinct things are incompatible; that is, any individually possible things are compossible with any others, or none at all (Leibniz notwithstanding). This is to say that things are what they are independently of one another, that their properties are intrinsic to them. Properties *as such* are “proper to” the things that have them.¹⁸

Can we make sense of the suggestion that human perceivers (but not animals) *hold* things to some such standard—*insist* upon it? Consider first (before turning to things) how the members of a family are perceivable (on a corporeal level): each has his or her own characteristic visual appearance, sound of voice, odor, way of moving, and so on; and, of course, their various parts stay attached in the same way. But suppose, one day, all these aspects started permuting: what looks like Sister sounds like Father, moves like Grandma, and smells like Kid Brother. Even the parts could mix up: Mother’s head (but Father’s hair) on Uncle’s torso with Baby’s limbs—or just two heads with no limbs or torso at all (sounding like a truck and smelling like a watermelon). And moments later, they switch again, with new divisions and new participants. What would you say?

Surely something like: “Egad! Am I going crazy? Am I being tricked or drugged? I can’t *really* be seeing this—it’s *impossible*”. That is, you would *reject* what you seemed to perceive, you would not accept them as *objects*. Now suppose that, instead of you, it were the family dog who came home to this. We can’t ask what it would say, because dogs can’t talk; and, of course, any estimate of its reaction at all is bound to be largely conjecture and prejudice. But, by way of counterpoint to sharpen the main point, I’ll express my own prejudice: I think the dog would *bark*. I expect it would be disoriented and distressed, maybe even frightened. But I can’t imagine any part of a dog’s reaction amounting to a rejection of the scene, a discounting of its reality, on the grounds that it’s impossible. Though Fido can tell Sister from Brother, and humans from cats, I don’t think he can distinguish in any sense between possible and impossible. And this, I believe, is the same as to say that he holds no objects to constitutive standards, and therefore understands nothing.

The integrity of family members—people—is certainly a different matter from that of things. Yet analogous permutations of the sensible properties of rocks and blossoms, comets and waterfalls would be equally fantastic. We might occasionally accept, even relish, such disintegrated phantasmagoria in dreams or hallucinations; but no one who understood what it meant could accept them as objective things. That is, the experiences, whatever else they might be, could not be objective perception. To perceive things as objects is to insist upon their coherent integrity—the constitutive standard for thinghood—just like insisting upon legality in chess, rationality in interpretation, and ordering with precision and scope in empirical science.

NOTES:

- 1 This is not intended as a conceptual or definitional point, but as a factual claim. It could turn out, I suppose, that dolphins or extra-terrestrials are capable of objective perception; but (at least in the case of the former) I’m highly dubious.
- 2 Abbreviate “*a*’s being *F* carries the information that *b* is *G*” with “ $Fa \Rightarrow Gb$ ”. Clearly, for a fixed specification of the conditions, this relation is transitive: If $Fa \Rightarrow Hc$ and $Hc \Rightarrow Gb$ in the same conditions, then (in those conditions) $Fa \Rightarrow Gb$. More, however, is required for $Fa \Rightarrow Gb$ to be *via Hc* than that $Fa \Rightarrow Hc$ and $Hc \Rightarrow Gb$ in

the same conditions—because, intuitively, it might be that $Fa \Rightarrow Gb$ independently of Hc . Suppose, however, that there are some conditions in which $Fa \Rightarrow Hc$ but not $Hc \Rightarrow Gb$, and some in which $Hc \Rightarrow Gb$ but not $Fa \Rightarrow Hc$, but the only conditions in which $Fa \Rightarrow Gb$ are those in which both $Fa \Rightarrow Hc$ and $Hc \Rightarrow Gb$. Then, plausibly (and according to Dretske), it is only *via* Hc that $Fa \Rightarrow Gb$.

- 3 It would not be difficult to contrive “abnormal” conditions in which the air or eardrum were vibrated in such a manner as to cause the perceptual experience, even without the bell ringing—for example, using an audio recorder.
- 4 The terminology is reminiscent of Gibson, whom Dretske in fact cites, in a slightly different context, on the following page.
- 5 In terms of the earlier illustration, this amounts to suggesting that, although the presence of the lemon drops does not carry any information about *which* child fetched, it does carry the information that *some* child fetched (so, ‘child’ = the higher-order kind); and, moreover, it carries the information that Grampa paid only *via* the information that some child fetched. (Note, by the way, that ‘some child’ is not in general equivalent to any disjunction: ‘child *a*, or child *b*, or ...’.)
- 6 This way of putting it may be misleading because, as I will argue in sections 4 and 5, it is possible in principle for nonlinguistic creatures to play chess; the reader, therefore, should not suppose that I intend this mention of conceptual ability to entail linguistic ability. (I will suggest in section 7 that conceptual understanding is possible without language.)
- 7 An existence proof for illegal moves is not required: Murphy’s Law is the best confirmed generalization in all of empirical metaphysics.
- 8 This counting-on and vigilance are to be understood as implicit in persistent behavioral dispositions, and not (at least, not necessarily) as conscious conviction, or deliberate being on the lookout for.
- 9 For nonlinguistic players (like our super-monkeys), *telling* what moves are made will be exhibited, in the first instance, in their

own legal move making (in the light of that move, so to speak). But it could also take the form of producing the equivalent moves (and positions) in various alternative media—in effect, “transliterating” them.

- 10 In chess, of course, both these abilities are already required in normal play, and are exhibited in a player’s ability to make legal moves of its own. Note that these two abilities “to tell” amount to abilities to *recognize*, respectively, the subordinate and superordinate patterns mentioned in section 3.
- 11 If I’m wrong about ordinary monkeys (or even dogs), then so much the better for them; what’s important is that there’s a line to be drawn, not where.
- 12 “...the conditions of the *possibility of experience* in general are likewise conditions of the *possibility of the objects of experience...*” (1781/1929, A 158/B 197).
- 13 This paragraph is based closely on a conversation with Jim Conant.
- 14 It could be argued that a first language is learnable only by bootstrapping. That is, only a minimal initial “core” is learnable without benefit of language; the remainder of the language is not so learnable, because learning (and/or speaking) it relies on the core. Then, if it were further maintained that chess is essentially more difficult than this core, the argument in the text would fail. What I find highly dubious in this line of thought, however, is the suggestion that anything *essentially* less sophisticated than chess (and similar formal games) could be in any proper sense a language. Yet, without that, the challenge collapses.
- 15 I say “more or less” because Davidson, in fact, nowhere offers an explicit definition of the physical, nor do the various remarks on which my attribution is based quite add up to one.
- 16 For a scientist, of course, “further articulation” can include more than just a better ability to tell what is and isn’t permissible; it can include modest modifications (“friendly amendments”) to the theory itself. (There is no analog to this for chess players.)
- 17 By ‘things’ here I mean mere things, *realia*—like sticks and stones. “Things” in the sense of paraphernalia or equipment, things

constituted by their roles and appropriate uses, are also held to constitutive standards, but different ones. Accordingly, the most proper perception of them as objects, and what is insisted upon in such perception, is also different

- 18 A fuller discussion would: disengage locatedness at a particular place from the togetherness (collocation) of cohesion; add a requirement for concreteness (complete determinacy) to that for compatibility; connect locatedness and concreteness with actuality (as opposed to mere possibility), and therefore with each other (and with particularity); and so on. But exactly how is another question.