

5

Kinds of Things—Towards a Bestiary of the Manifest Image

Daniel C. Dennett

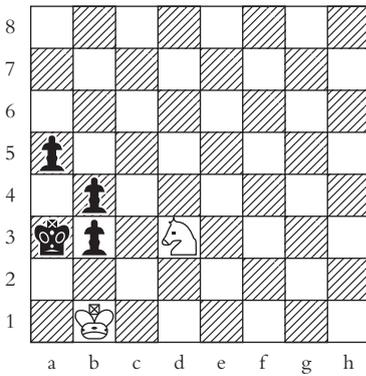


Figure 5.1 White to play (hint: checkmate in two moves).

WHITE TO PLAY
Hint: Checkmate in 2 moves.

Consider this chess puzzle. White to checkmate in two. It appeared recently in the *Boston Globe*, and what startled me about it was that I had thought it had been proven that you can't checkmate with a lone knight (and a king, of course). I was wrong; as David Misialowski (personal correspondence) has pointed out to me, what had been proven is just that you cannot checkmate your opponent's king when only his king and your king and knight are on the board. The fact that the proposition *you can never checkmate with a lone knight and king* is not a truth of chess is a *higher-order* truth of chess. Now let's consider chmess (I made up the term); it's the game that you get by allowing the king to move two spaces, not one, in any direction.

I made it up and I don't know if anybody has ever played it. I have no idea whether it's worth playing. Probably it isn't, but it doesn't matter; it's not even worth our attention long enough to figure out. But a moment's reflection reveals that there are exactly as many higher-order a priori truths of *chmess* as there are of chess, namely, an infinity of them. And no doubt they would be roughly as difficult to discover and to prove as the higher-order truths of chess. There are people who make a living working out the truths of chess and certainly it's been a big avocation for many other people. But I doubt if anybody yet has spent more than five minutes trying to work out the a priori truths, and the higher-order truths, of *chmess*.

I use this example (complete with my mistake, now corrected by Misialowski) in a paper, 'Higher-order truths about *Chmess*' (Dennett, 2006), my own salvo against analytic metaphysics and other dubious battles in philosophy. It is, more broadly, my advice to graduate students in philosophy on how to avoid the trap of getting stuck in a career devoted to an artifactual 'problematic.' As one of my heroes, Donald Hebb once said, 'If it isn't worth doing, it isn't worth doing well.' I think we have to admit that there are quite a few philosophers who do very well what they do, executing their self-assigned tasks with cleverness, technical proficiency, and rigor, but engaged in an enterprise that's just not worth doing.

So I applaud the first chapter of Ladyman and Ross (2007), which articulates the grounds for my conviction much better than I could ever have done. They comment briefly on my discussion (in Dennett, 2005) of Patrick Hayes, the artificial intelligence researcher, who set out on a project to axiomatize the naïve (or folk) physics of liquids. I'm going to expand a bit on this point, to clarify the message we both derive from this ingenious enterprise. The idea was that a robot who interacted with everyday folk would need to understand folk physics—the manifest image of liquids, solid objects, and so on—so, in classic GOFAI¹ style Hayes set out provide such a robot with the propositions it would need to use as its core beliefs. It proved to be more challenging than he had anticipated, and he wrote an interesting paper about the project, 'The Naïve Physics Manifesto' (Hayes, 1978). In the naïve physics of liquids, everything that strikes naïve folks as counter-intuitive is, of course, ruled out: siphons are 'impossible' and so are pipettes, but you can mop up liquid

¹ Good Old-Fashioned Artificial Intelligence.

with a fluffy towel, and pull water out of a well with a suction pump. A robot equipped with such a store of ‘knowledge’ would be as surprised by a siphon as most of us were when first we saw one in action. Hayes’ project was what I would call *sophisticated* naïve physics, because he was under no illusions; he knew the theory he was trying to axiomatize was false, however useful in daily life. This was an exercise in what might be called axiomatic anthropology: you treat what the folk say—and agree about—as your axioms or theorems, and try to render the data set consistent. And, of course, he didn’t bother rounding up any actual informants; he figured that he knew the naïve physics of liquids as well as any normal person did, so he used himself as his sole informant: axiomatic *auto*-anthropology.²

Now compare Hayes’ project with the projects of analytic metaphysics, which often strike me as *naïve* naïve auto-anthropology since the participants in this research seem to be convinced that their program actually gets at something true, not just believed-true by a particular subclass of human beings (Anglophone philosophers of the analytic metaphysics persuasion). Otherwise, the programs seem identical: you gather your shared intuitions (testing and provoking them by engaging in mutual intuition-pumping) and then try to massage the resulting data set into a consistent ‘theory,’ based on ‘received’ principles that count, ideally, as axioms. I’ve asked a number of analytic metaphysicians whether they can distinguish their enterprise from naïve auto-anthropology of their clan, and have not yet received any compelling answers.

The alternative is *sophisticated* naïve anthropology (both auto- and hetero-)—the anthropology that reserves judgment about whether any of the theorems adduced by the exercise deserves to be trusted—and this is a feasible and frequently valuable project. In this essay I am going to propose that this is the enterprise to which analytic metaphysicians should turn, since it requires rather minimal adjustments to their methods and only one major revision of their *raison d’être*: they must roll back their pretensions and acknowledge that their research is best seen as a

² Hayes’ work inspired others. See the anthology edited by Daniel Bobrow, *Qualitative Reasoning about Physical Systems* (1985). Authors of one of the essays therein remark: ‘Naïve physics is in itself an ambiguous term. Is it just bad physics? Is it psychology? Artificial intelligence? Physics?’ (de Kleer and Brown, p. 13) The answer, I submit, is that naïve physics is the attempt to make the physics part of our manifest image rigorous enough to support automated deductive reasoning.

preparatory reconnaissance of the terrain of the manifest image, conducted between brackets, as it were, like the Husserlian *epoché*: let's pretend for the nonce that the natives are right, and see what falls out. Since at least a large part of philosophy's task, in my vision of the discipline, consists in negotiating the traffic back and forth between the manifest and scientific images, it is a good idea for philosophers to analyze what they are up against in the way of folk assumptions before launching into their theory-building and theory-criticizing. Philosophical work on the perennially hot-button topic of free will, for instance, certainly must be guided by an appreciation of what non-philosophers think free will is or might be—and why it matters so much to them.

One of the hallmarks of sophisticated naïve anthropology is its openness to making counterintuitive discoveries. As long as you're doing naïve anthropology, counterintuitiveness (to the natives) counts against your reconstruction; when you shift gears and begin asking which aspects of the naïve 'theory' are true, counterintuitiveness loses its dispositive force and even becomes, on occasion, a sign of significant progress. In science in general counterintuitive results are prized, after all.

One of the weaknesses of auto-anthropology is that one's own intuitions are apt to be distorted by one's theoretical predilections. Linguists have known for a long time that they get so wrapped up in their theories that they are no longer reliable sources of linguistic intuition. Their raw, untutored intuitions have been sullied by too much theory, so they recognize that they must go out and ask non-linguists for their linguistic intuitions. Philosophers have recently begun to appreciate this point, in the new enthusiasm for so-called experimental philosophy. It is early days still, and some of the pioneer efforts are unimpressive, but at least philosophers are getting used to the idea that they can no longer just declare various propositions to be obviously true on the sole grounds that they seem smashingly obvious *to them*. (In a similar vein Hayes *might* have surprised himself about the chief tenets of folk physics if he had gone to the trouble of interviewing a random sample of folk instead of just treating himself as exemplary.)

So there is a project, a particular sort of sophisticated naïve anthropology, that philosophers should seriously consider as a propaedeutic inquiry before launching into their theories of knowledge, justice, beauty, truth, goodness, time, causation, and so on, to make sure that they actually aim their analyses and arguments at targets that are relevant to the rest of the

world, both lay concerns and scientific concerns. What they would get from such a systematic inquiry is something like a *catalogue raisonné* of the unreformed conceptual terrain that sets the problems for the theorist, the metaphysics of the manifest image, if you like. This is where we as philosophers have to start in our attempts to negotiate back and forth between the latest innovations in the scientific image, and it wouldn't hurt to have a careful map of this folk terrain instead of just eyeballing it. (This is the other half, one might say, of the reform that turned philosophy of science from an armchair fantasy field into a serious partnership with actual science when philosophers of science decided that they really had to know a lot of current science from the inside.) Once one thinks about our philosophical tasks with this image in mind, one can see that a great deal of the informal trudging around, backing and filling, counter-example-mongering and intuition-busting that fills the pages of philosophy journals is—at best—an attempt to organize a mutually acceptable consensus about this territory. Maybe a Mercator of the Manifest Image will emerge from experimental philosophy.

But that won't be me. I am instead going to fall back on the good old-fashioned informal methods of ordinary language philosophy—the 'discipline' I was 'trained' to do—and draw attention to some of the unanswered questions that one might want to take seriously in sketching the ontology of the manifest image. The only innovation in my method is that I encourage us all to be *sophisticated* naïve (auto-) anthropologists, perhaps *savoring* our intuitions about cases, but not *endorsing* them—yet.

The main goal in the rest of this paper is simply to draw attention to the riotous assembly of candidate *things* that we find in the manifest image. Wilfrid Sellars (1963) 'discovered' the manifest image much the way Columbus discovered the western hemisphere, but he, like Columbus, was more the explorer than the systematic cartographer, and, as just noted, our Mercator has yet to appear. But what about a Linnaeus, a systematic taxonomist of all the things we find in this new, if all too familiar, world? Ontology, it seems to me, has been resolutely reductive and minimalist in its attempts to come up with an exhaustive list of kingdoms, classes, genera and species of things. No doubt the motivation has been the right one: to prepare the disheveled cornucopia for scientific accounting, with everything put ultimately in terms of atoms and the void, space-time points, or (by somewhat different reductive trajectories) substances and universals, events and properties and relations. As is well known, these procrustean

beds provide metaphysicians with no end of difficult challenges, trying to fit all the candidates into one austere collection of pigeonholes or another. Where do we put numbers, or melodies, or ideas, for instance? Do numbers even exist, and are there different kinds of existence? Instead of attempting any answers of my own to these oft-asked questions of systematic taxonomy, I am going to push, informally, in the opposite direction, like the bestiarists of yore, recounting with as much ontological tolerance as I can, a few of the neglected candidates for thing-hood that everyday folks have no trouble countenancing in their casual ontologies. Metaphysics as folklore studies, you might say.

Quine did some pioneering on this score in *Word and Object* (1960). Are there *sakes*? Are there *miles*? He used these homely examples to loosen up his readers for his theory-driven proposals, his quest for a 'desert landscape' clean-shaven by Occam's Razor, a minimalist ontology for science (and for everything worth talking and thinking about—he was Carnap's student, after all). Since his overall aim was eliminative, he no more had to address the question of whether sakes and miles belonged to different kinds of thing than Linnaeus had to worry about whether mermaids were mammals and dragons were reptiles. Nor were the Oxbridge ordinary language philosophers interested in constructing a systematic taxonomy of the ordinary. Their informal but clever sallies did unearth some useful examples, however, and I want to add some more for the consideration of sophisticated naïve (auto- and hetero-) anthropologists.

Let's consider some of Quine's examples in a little more detail. In addition to Paris and London are there the *miles* that separate them? Are there *sakes*, as suggested by such familiar idioms as 'for Pete's sake' and 'for the sake of — '? What are the identity conditions for sakes? How many sakes might there be? How do they come into existence? Can one still do something for FDR's sake, or did it expire with him decades ago? Art for art's sake shows that not only living things can have sakes. Quine gave us reasons to see the ontological candidacy of sakes as something of an accident of ordinary language (more specifically, of English, since there aren't clear synonyms in all other languages), a sort of potential cognitive illusion frozen in a few idioms, waiting to be vivified by an unwary and profligate semanticist. The restriction to a few idioms ('for the sake of,' 'on behalf of,' and the like) which might rather gracefully be 'fused' into unanalyzed prepositions does suggest, as Quine proposed, that this noun is *defective*, and its putative referents need not be taken seriously from an

ontological point of view. He suggested that we dismiss all the nouns of measure—miles, ounces, minutes, and their ilk—by the same principle, a more demanding exercise of self-denial, since talk about gallons and kilograms and minutes is less easily paraphrased away, and plays a much richer role in daily life.

What about *dints*? Again, the noun, descended from a now extinct term for a blow or strike, has lost its combinatorial power, almost as inert as the *board* in *cupboard*. Not a compelling candidate—on the surface, anyway—for thing-hood, but we shouldn't jump to conclusions. It is noteworthy that computer scientists in the last few decades have found the task of knowledge representation for software agents or bots, where the demands of automated inference put a severe strain on data structures, is often clarified and expedited by the devising of ingenious ontologies that carve up what might be called the manifest image of their bots in novel ways. Dints could conceivably prove to be just the things to represent when reasoning about actions and effects of certain sorts.

Then there are *cahoots* and *smithereens*, both particularly unpromising. These are probably just what Quine said about sakes, degenerate cases that do not require patient elaboration. It is bad enough having graduate students writing whole dissertations on the relationship of constitution between the stone and the statue; we need not create a microdiscipline, smithereen theory, to work on the sub-problem of what to say about all the bits when the wrecking ball hits the statue.

David and Stephanie Lewis (1970) made famous the difficulties of either banishing *holes* altogether or fitting them into any of the austere schemes favored by Quine and others. As they amusingly showed in the dialogue between Argle and Bargle, getting through the day without holes in one's ontology is no easy matter. They never considered the tempting hypothesis that our reliance on the concept of holes has a deep biological or ecological source, since holes are Gibsonian affordances *par excellence*, put to all manner of uses in the course of staying alive and well in a hostile world. Ubiquitous and familiar though they are, it is surprisingly hard to say what holes are made of, if anything, what their identity conditions are, whether they are concrete or abstract, how to count them, and so forth. In the end Bargle says: 'I for one have more trust in common opinion than I do in any philosophical reasoning whatsoever.' To which Argle replies: 'My theories can earn credence by their clarity and economy; and if they disagree a little with common opinion, then common opinion may be

corrected even by a philosopher' (Lewis and Lewis, 1970: 211–12). The sophisticated naïve anthropologist can calmly note the impasse from a position of neutrality and conclude (tentatively) that holes are virtually indispensable in the ontology of the manifest image, but probably not portable, without strikingly counterintuitive revision, to the austere scientific image. But not so fast: holes may play a potent role in organizing the patterns studied in the special sciences—consider the importance of membranes and their multiple varieties of portals in cell biology, for instance—and what about the slit that figures so prominently in quantum physics?

Not so fundamental or indispensable, no doubt, but still not easily dismissed are the *voices* I discussed (Dennett, 1969) at some length at about the same time. My project then, rather like the Lewises' project with holes, was to coax philosophers out of their self-imposed disjunctions about *minds*—are they brains *or* spirits, concrete *or* abstract, *or* are they, strictly speaking, as non-existent as witches and mermaids? I tried to do this by showing that voices, ostensibly nowhere near as mysterious as minds (there is no voice-throat problem to set alongside the mind-body problem, thank goodness), were just as uncategorizable, and none the worse for that. Why are voices more robust candidates for thing-hood than sakes? In part because synonyms or near-synonyms for 'voice' occur in many languages, and the contexts in which these nouns appear are remarkably diverse and varied. You can hear a voice, lose a voice, strain a voice, record a voice, recognize a voice, disguise a voice, and use technology to see through an acoustic disguise to identify the underlying voice, for instance. What *kind* of thing is a voice? It can be strained, like a larynx or knee, but it is not an organ, since it can be lost and recovered. Is it a disposition or prowess or idiosyncratic competence or method, something like a *style*? In a way, but it can also be recorded, like a sound or a melody. Maybe a voice is the sole exemplar of an ontological kind all of its own. Maybe bass voices, soprano voices, childish voices, raspy voices and nasal voices are all kinds of voices but voices aren't kinds of anything at all except the voice kind. (There could be other such kinds, of course. In some cultures that use bells a lot, there might be the recognizable *ding* of this or that bell.) Here is an interesting question in the empirical theory of perception and knowledge as it merges with semantics (see Jackendoff, 2002, 2007): why is it so natural for human minds speaking many different languages to hit upon voices as a valuable addition to their ontology? But remember: we

are bracketing for the time being the question of which elements of the ontology of the manifest image we take ourselves to be committed to.

Haircuts (and hairdos more generally) are another variety that is strangely resistant to pigeonholing. A haircut is not really an event or a property or a relation or a material object (made of hair, presumably). How long haircuts last, and just how to apply the type-token distinction to them (think of the Beatles haircuts all around the world), are not readily answered questions.

The wonderful world of economics provides us with further socially constructed things, such as dollars and euros. Do dollars exist? It seems obvious that they do, and—unlike some of our other problematic things, they seem eminently countable, designed to be countable in fact. But we should note, before moving on, that there has been a historical progression in the direction of abstraction. No doubt a large source of the confidence with which we countenance the existence of dollars as things is that today's dollars are the descendants, in effect, of paradigmatic things, made out of metal and having a shape and weight, like dimes and nickels and silver dollars. These were followed by first paper dollar *bills*, five-pound *notes*, and their kin, and then entirely abstract dollars, by the trillions. It's a good thing we don't have to make them all out of metal; they would bury the world. We gain a tremendous efficiency from this transition, and all we lose is the standard way of keeping track of individual concrete things. It makes sense to ask whether this particular dime or silver dollar once passed through the hands of Jacqueline Kennedy, but not whether the hundred dollars I just sent to my PayPal account ever did.

The concept of abstract dollars is undeniably useful in our everyday affairs, but is utility tantamount to reality? This question sits happily within the perspective of sophisticated naïve anthropology but is out of place in traditional metaphysics. Consider the reality of centers of gravity or mass. Are they real or not? Some philosophers have said yes, and others no, and both sides have used their conviction to disagree with my use of them as a useful analogy for the soul, as the 'center of narrative gravity' (Dennett, 1991). Compare the concept of the center of gravity of, say, an automobile with the concept of Dennett's lost sock center, defined as the center of the smallest sphere that can be inscribed around all the socks I have ever lost in my life. The former is an undeniably useful concept, and so *palpably* useful that you can almost feel a center of gravity once you know what to feel for. The toddler's toys, called Weebles, 'wobble but they don't fall down';

watching a child discover this marvelous property by exploratory manipulation, it is hard to resist casting the task of *finding the center of gravity* alongside the task of *finding the hidden coin* as different instances of the same activity. But however perceptible or tangible a center of gravity is, it is (one would suppose, as a metaphysician) in exactly the same basic ontological genre as Dennett's lost sock center, and kazillions of other equally pointless abstractions we could invent. Must we fill our world with all that pseudo-Platonic garbage? Scientific utility, as Quine never tired of reminding us, is as good a touchstone of reality as any, but why shouldn't utility *within the manifest image* count as well? Is there anything dangerously relativistic in acknowledging that the two images may have their own 'best' ontologies, which cannot be put into graceful registration with each other? (This suggestion is a close kin, and descendant, of Quine's ontological relativity (1969), and his claims about the indeterminacy of radical translation, but I cannot do justice to exploring the relations in the limited space of this essay.)

In the introduction to *Brainstorms* (1978), I imagined a society that speaks a language just like English except for

one curious family of idioms. When they are tired they talk of being beset by *fatigues*, of having mental fatigues, muscular fatigues, fatigues in the eyes and fatigues of the spirit . . . When we encounter them and tell them of our science, they want to know *what fatigues are* . . . what should we tell them? (pp. xix–xx)

Our task, I suggested then, and still maintain, is more a matter of *diplomacy* than philosophy. It is not that there is or must be—there *might* be—a univocal, all-in, metaphysical *truth* about what there is, and whether fatigues can be identified as anything in that ultimate ontology, but just better and worse ways of helping people move between different ontological frameworks, appreciating at least the main complexities of the failures of registration that are encountered. For anyone interested in such diplomatic projects, a vigorous anthropological scouting is good preparation.

Metaphysicians and philosophers of mathematics have devoted years of work to the question of whether numbers exist, while taking the existence of *numerals* more or less for granted, but why? Numerals, as tokens, are no doubt unproblematic trails of ink or chalk or some other physical particular, but as types, they are more problematic. They are like words in this regard. Words are rather like voices: they seem to be a set with no obvious

superset. Nouns and verbs are kinds of words, but what are words a kind of? Symbols? Sounds? I have suggested that words are a kind of *meme*; words are memes that can be pronounced (Dennett, 2009). They are also, as Jackendoff (2002) puts it, ‘autonomous, semi-independent information structures, with multiple roles in cognition.’ In other words, they are a kind of software that runs on the brain.

And how about software? Does it exist? I have been fascinated to learn recently that there are philosophers who are reluctant to countenance software (and all its subspecies, from data structures and subroutines to apps and widgets) as being among the things that truly exist. This is a descendant, apparently, of a similar line of ontological skepticism or mistrust about information in general, as something—some thing—that can be stored, moved, processed.

Once again, the perspective I would recommend is that of the diplomatic anthropologist, not the metaphysician intent on limning the ultimate structure of reality. The ontology of everyday life is now teeming with items that, like fatigues, sit rather awkwardly in the world of atoms and molecules. If we can understand how this population explosion came about, and why it is so valuable to us as agents in the world, we can perhaps discharge our philosophical obligations without ever answering the ultimate ontological question. To me it looks more and more like professional make-work, an artifact of our reasonable but ultimately optional desire for systematicity rather than a deeper mystery in need of solving.

References

- Bobrow, D. (1985). *Qualitative Reasoning about Physical Systems*, Cambridge, MA: MIT Press.
- Dennett, D. (1969). *Content and Consciousness*. London: Routledge & Kegan Paul.
- (1978). *Brainstorms: Philosophical Essays on Mind and Psychology*. Montgomerly, VT: Bradford Books.
- (1991). Real Patterns. *Journal of Philosophy*, LXXXVIII: 27–51.
- (2005). *Sweet Dreams: Philosophical Obstacles to a Science of Consciousness*. Cambridge, MA: MIT Press.
- (2006). Higher-order truths about chmess. *Topoi*, 25: 1–2.

- (2009). The cultural evolution of words and other thinking tools. *Cold Spring Harbor Symposia on Quantitative Biology*, 74: 435–41 [online] <<http://ase.tufts.edu/cogstud/papers/coldspring.pdf>> accessed 11 May 2012.
- Hayes, P. (1978). The naïve physics manifesto. In D. Michie (ed.) *Expert Systems in the Microelectronic Age*. Edinburgh: Edinburgh University Press.
- Jackendoff, R. (2002). *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
- (2007). *Language, Consciousness, Culture: Essays on Mental Structure*. Cambridge, MA: MIT Press.
- Ladyman, J. and Ross, D. (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.
- Lewis, D. and Lewis, S. (1970). Holes. *Australasian Journal of Philosophy*, 48, 2: 206–12.
- Quine, W. (1960). *Word and Object*. Cambridge, MA: MIT Press.
- (1969). *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Sellars, W. (1963). *Science, Perception and Reality*. London: Routledge & Kegan Paul.