

Intellectualism as Cognitive Science

Martin Roth

Robert Cummins

In *Knowledge and Representation: Recent Developments*. Albert Newen, ed. Mentis, 2011.

I. Introduction

A well entrenched view in epistemology, philosophy of mind, and cognitive science is that there is an important distinction between knowing *how* to do something and knowing *that* something is the case. According to this view, while knowing-that requires standing in a relation to a proposition, knowing-how does not; instead, knowing-how is a matter of possessing abilities, capacities, and skills, and thus constitutes a form of “non-propositional” knowledge. However, the integrity of this distinction has been challenged, and some philosophers argue that *intellectualism*—the thesis that knowing how to do something is a matter of possessing propositional knowledge—is true. Pertinent to how this issue will be resolved, of course, are certain assumptions about the appropriate methodology for resolving it, and in this paper we argue that a certain method for establishing intellectualism is deeply suspicious. Furthermore, once we get clear about how intellectualism is understood as an *hypothesis in cognitive science*, it will become apparent that there are serious grounds for doubting the truth of intellectualism.

II. Against Poaching

Minimally, intellectualism is the thesis that, for any ψ , knowing how to ψ requires propositional knowledge of a way of ψ -ing. For example, according to intellectualism, if Mary knows how to tie shoes, then Mary stands in a knowledge-that relation to a

proposition concerning a way of tying shoes. The most famous attack on intellectualism is due to Gilbert Ryle (1949), who argued that intellectualism faces insuperable difficulties. Ryle's arguments are well known, and we will discuss some of them in the next section. Ryle's alternative to intellectualism is also well known: knowing-how is a matter of possessing abilities (multi-track dispositions to behave), not knowing that something is the case. For Ryle, such abilities include both "physical" abilities (e.g., knowing how to swim) and "mental" abilities (e.g., knowing how to multiply). Recently, however, there have been attempts to revive intellectualism, and in this paper we focus on an attempt due to Jason Stanley and Timothy Williamson (2001).

Stanley and Williamson offer both a criticism of the view that knowledge-how is an ability and a positive argument that knowledge-how is a species of knowledge-that. The criticism is quite simple: there are correct knowledge-how ascriptions to people who lack the abilities purportedly specified by the ascriptions. For instance, they offer the examples of a ski instructor who, though he himself cannot perform a complex stunt, may know how to perform such a stunt, and of a master piano player who loses both of her arms in a car accident (thereby losing the ability to play the piano) but is still correctly described as knowing how to play (p. 416). Since the thesis that knowing-how is an ability implies that these situations are impossible, and since it is intuitively plausible that there could be such an instructor and piano player, the thesis that knowledge-how is an ability is false.¹

¹ How conclusive are these examples? The thought underlying both cases is that if a person is presently unable to do ψ , then the person lacks the ability to ψ . Though this account of the connection between what a person is able to do and her abilities is questionable (Millikan, 2000; Noë, 2005), we will concede the examples for the purposes of this paper.

Turning to their positive argument, Stanley and Williamson's claim that knowledge-how is a species of knowledge-that relies on syntactic and semantic analyses of knowledge-how ascriptions. If we assume generally that the specification of truth conditions for knowledge-that ascriptions involve reference to propositions, we can clarify Stanley and Williamson's conclusion as follows: the correct specification of the truth conditions of knowledge-how ascriptions involves reference to a proposition. Furthermore, it is in virtue of the appeal to propositions in specifying the truth conditions of knowledge-how ascriptions that knowledge-how is a species of knowledge-that. According to Stanley and Williamson, we should analyze sentences like

(1) Hannah knows how to ride a bicycle

in the same way we analyze other sentences containing embedded questions, such as

(2) Hannah knows what to do in the case of an earthquake

(3) Hannah knows whom to call for help

After describing commonly accepted syntactic and semantic analyses for sentences involving embedded questions, they offer the following statement of the truth conditions for sentences like (1):

(1) is true iff, for some contextually relevant way w which is a way for Hannah

to ride a bicycle, Hannah knows that w is a way for her to ride a bicycle, under a practical mode of presentation (p. 430).

Stanley and Williamson claim that if Hannah knows how to ride a bicycle, then Hannah stands in a knowledge-that relation to a Russellian proposition that contains a way of riding a bicycle as a constituent (p. 427). Furthermore, Stanley and Williamson claim that the contextually relevant way w that is a constituent of a Russellian proposition can be picked out demonstratively. For example, the way which Hannah knows how to ride a bicycle may be introduced by a friend pointing to a man riding a bicycle and saying, “That is a way to ride a bicycle” (p. 428). Thus, Hannah’s knowing how to ride a bike does not require that she be able to describe non-indexically a way for her to ride a bike. Because Stanley and Williamson show that it is possible to give an analysis of knowledge-how ascriptions using a standard account of knowledge-that ascriptions, they conclude that knowledge-how is a species of knowledge-that:

We take our view of ascriptions of knowledge-how to be the default position. From a linguistic perspective, very little is special about ascriptions of knowledge-how. It is hard to motivate singling them out for special treatment from the rest of a family of related constructions. Our view of ascriptions of knowledge-how is the analysis reached on full consideration of these constructions by theorists unencumbered by relevant philosophical prejudices (p. 431).

On the face of it, these remarks are puzzling. As Alva Noë (2005) points out, Ryle’s distinction is not intended as a thesis about the *sentences* used to attribute knowledge, but rather about the nature of knowledge itself. So why do Stanley and Williamson take

themselves to have established intellectualism on the basis of linguistic considerations of knowledge-how ascriptions?

There is a tempting line of argument running from the semantic analysis of sentences involving “mental” terminology to conclusions about the structure of the mind. One starts, for example, with a truth-conditional analysis of belief sentences, and argues, let’s say, that ‘believes’ is a three place relation between a believer, a proposition and a sentence in the language of thought (LOT) that expresses it (e.g., Fodor 1981). One then notes that some belief attributions are true. Since the truth condition for belief attributions requires a belief relation between a believer, a proposition and an expression in LOT, it seems that anything to which one can truly attribute beliefs must harbor psychological states with precisely that structure. And, on the assumption that the brain realizes psychological states, we get a conclusion about the structure of the brain *without having to do a single experiment!*

This is way too much for way too little. The lure of armchair neuroscience and psychology is surely too good to be true. It is what we call *epistemological poaching* (Cummins, 1991). A poacher is someone who hunts (or fishes) without a valid license. No sane epistemology of science will grant a neuroscience license, or a cognitive psychology license, on the basis of truth-conditional semantics and the truism that people have beliefs.

We think Stanley and Williamson are poaching when they apply their results to the debate over intellectualism. At face value, the take home message here seems to be that something like LOT and the classical program must be on the right track scientifically. The scientific evidence against LOT is not just outweighed, but trumped by

some linguistically bolstered truth-conditional semantics. But the truth-conditional semantics of knowledge-how attributions, however convincingly bolstered by linguistic evidence, cannot possibly license conclusions about the viability of intellectualism in cognitive science, or about the shortcomings of its opponents. If the Stanley and Williamson analysis of knowledge ascriptions is correct (which we doubt), then we must concede the cognitive sciences may determine that those ascriptions are false. You cannot have it both ways: you cannot assume the truth of knowledge-how ascriptions, while analyzing their truth conditions in a way that prejudices substantive scientific issues.

As we will show in section IV, the debate in cognitive science between advocates of intellectualism and their opponents is evidently an empirically grounded debate about what sort of functional analysis of the brain is required to explain cognitive, perceptual and motor capacities. More particularly, it is about the LOT hypothesis central to intellectualism. As that hypothesis is understood, by both its advocates and its critics, it is not a hypothesis that has implications for the proper truth-conditional semantic analysis of knowledge-how attributions. Nothing in the connectionist inspired treatment of know-how (e.g., Churchland, 1989(a); Bechtel and Abrahamsen, 1991; Haugeland, 1998), for example, has any implications for the truth-conditional semantics of know-how attributions. Nor do semantic analyses of knowledge-how attributions have implications for the debate over LOT. These issues are, or should be, orthogonal to each other. If one is inclined to think that there *is* a genuine inference from a premise about the truth-conditional semantics of psychological attribution sentences to conclusions about the scientific viability of intellectualism, one should become skeptical of truth-conditional semantics. No program that licenses poaching can be right.

So, just what exactly *is* wrong with the simple move from the truth of some know-how attributions, plus the semantics of such attributions, to conclusions about intellectualism? Notice first that there obviously *is* something wrong with such arguments: they license poaching. To repeat, no sane epistemology can allow that the semantics of ordinary talk can substitute for serious neuroscience or psychology. We do not need a persuasive deconstruction of poaching arguments to know they are no good.

With that said, it seems to us that there are two possible lines of deconstruction available. The first and most radical is that truth-conditional semantics is just wrong as an account of meaning. While we have some sympathy with this idea, it doesn't really address the problem. Whether or not you believe that truth-conditional semantics is the right approach to meaning, it is hard to deny that statements have truth-conditions. Anywhere there is an overlap of vocabulary between ordinary talk on the one hand, and science on the other, truth-conditional semantics will provide a link between the substantive claims of the science and ordinary speech. That link—Davidson's Demon, we might call it, since he released it in "Truth and Meaning," (1967)—will force us either to constrain science to honor ordinary talk—to license poaching—or to constrain ordinary talk to honor science. The latter alternative, on reflection, seems obviously right. Truth-conditional semantics tells us that meanings are truth-conditions. It does not, and should not, tell us the truth-values of contingent statements. We can grant Stanley and Williamson's analysis of know-how statements, but deny that it has any implications at all for cognitive science. The truth-values of our ordinary statements *should* be held hostage to the progress of science.

There is another possibility: truth-conditions, if they are to be meanings accessible from the armchair, must not go more than skin deep. They must be little more than meta-language clones of the object language expressions under analysis. In the case of belief, for example, they should be indifferent to the differences between Dennett, Churchland, and Fodor about what beliefs actually are, if there are any. This would make semantics much less fun, of course, but it might make it more honest.² In the case of know-how, the lesson is that the kinds of arguments advanced by Stanley and Williamson should interest the cognitive scientist only if such arguments could bear on the debate between, e.g., connectionists and classicists; but they can't, so they shouldn't.

We do not propose to settle the issue of how semanticists should think of what they are doing so as not to engage in poaching. That, of course, is their problem. But it is a problem they must face. Semantics cannot license poaching.

III. Ryle on Intellectualism

Some people know how to ski, and some people who know how to ski occasionally display their knowledge on the slopes. How should we understand the relationship between knowing how to ski and particular instances of skiing? In chapter II of *The Concept of Mind*, Gilbert Ryle characterizes and critiques an “intellectualist” answer to this question. According to intellectualism, as Ryle describes it, knowing how to ski requires propositional knowledge (knowledge-*that*) of ways to ski. When people act on the intention to ski, they first consult their knowledge-that and formulate a plan for action. The action itself is the output of this planning process. If we think of

² A fair chunk of analytic metaphysics would pretty much evaporate, as well.

propositional knowledge of ways to ski as constituting a “mini-theory” of how to ski, then we can say that actions that manifest know-how are the causal consequences of a theory driven planning process. As Ryle puts it, we first do a bit of theory and then do a bit of practice (p. 29). Famously, Ryle thought intellectualism was deeply flawed. For example, he argues that intellectualism leads to an unacceptable regress: since theorizing *itself* is something we know how to do, explaining any particular instance of theorizing would require positing still *further* acts of theorizing, ad infinitum (p. 30). Hence, there has to be some knowing-how that cannot be explained in terms of knowing-that.

If intellectualism is a non-starter, what is the alternative? According to Ryle, knowledge-how is not to be found in causes, but in “capacities, skills, habits, liabilities, and bents” (p. 45). Exercises of know-how are simply manifestations of capacities or abilities, where capacities and abilities are understood by Ryle to be dispositions “the exercises of which are indefinitely heterogeneous” (p. 44). For example, although Tom and Jerry may each succeed in producing the correct answer to a particular multiplication problem, if Tom is a skilled mathematician but Jerry is merely answering by rote, then Tom’s behavior counts as an exercise of know-how, and Jerry’s does not. The etiology of their respective behaviors is simply irrelevant. Knowing-how is a disposition, not a cause.

Ryle’s argument for identifying know-how with abilities is based, in part, on how Ryle thinks we use the phrase ‘know-how.’ For example, consider what is involved in knowing how to play chess. Ryle writes, “It should be noticed that the boy is not said to know how to play, if all that he can do is to recite the rules accurately. He must be able to make the required moves. But he is said to know how to play if, although he cannot

recite the rules, he normally does make the permitted moves, avoid the forbidden moves and protest if his opponent makes forbidden moves” (p. 41). However, it would be a mistake to conclude from this that Ryle identified know-how with dispositions to perform publicly observable behavior, for while Ryle was certainly skeptical of ‘mechanistic’ psychology, he did not deny the existence of operations that occur ‘in the head.’ For example, Ryle wasn’t claiming that knowing how to multiply requires performing bits of publicly observable behavior, e.g., writing marks on paper with pencil. Rather, he emphasized that the difference between solving a problem ‘in the head’ and solving a problem with pencil and paper does not mark a theoretically interesting distinction, as far as the issue of *intellectualism* is concerned. To put the point another way, since the abilities that constitute know-how can be exercised ‘in the head,’ the intellectualist’s distinction between theory and performance does not simply reduce to some sort of “private/public” distinction. In addition, Ryle did not deny the existence of intellectual operations or their importance to know-how. For example, Ryle writes “the learning of all but the most unsophisticated tasks requires some intellectual capacity,” which in turn requires understanding and “propositional competence” (p. 49). Ryle was not hostile to granting the existence of intellectual abilities (know-how), as long as this was not understood as requiring that the exercise of such abilities be preceded by some prior act of theorizing.

IV. Intellectualism as Cognitive Science

For better or worse, in the mouths of philosophers and psychologists, the phrase ‘know-how’ was more or less used to refer to what Ryle said it did—abilities, capacities, and

skills. Whether using ‘know-how’ in this way reflected actual linguistic practice simply became irrelevant—there were substantive questions about the nature and explanation of such capacities and skills, regardless of what we *called* them. Along with the cognitive revolution in psychology, however, came the revival of intellectualism. Recall that, for Ryle, intellectualism is the thesis that knowing how to ψ consists in propositional knowledge of ways to ψ , knowledge that is deployed in the course of generating ψ -behavior, and when Chomsky (1959) showed that Skinner’s behaviorism was inadequate to explain language acquisition, intellectualism became a compelling *empirical hypothesis*—children possess innate knowledge of rules of grammar (knowledge-*that*), and this knowledge allows children to formulate, test, and confirm a *tacit theory* of English, a theory whose application (in part) explains their linguistic performance.

Since it was central to the cognitive revolution that much of human and animal behavior *is* best explained in terms of the deployment of propositional knowledge, if cognitive psychology had any chance at succeeding, Ryle needed to be wrong. An early attempt to refute Ryle can be found in Jerry Fodor’s 1968 paper “The Appeal to Tacit Knowledge in Psychological Explanation.” In this paper, Fodor defends intellectualism as an account of mental competences, which Fodor identified with abilities (e.g., the ability to play chess, the ability to speak Latin). To give a feel for his version of intellectualism, Fodor begins his paper by inviting the reader to take seriously the following explanation of how we tie our shoes:

There is a little man who lives in one's head. The little man keeps a library.

When one acts upon the intention to tie one's shoes, the little man fetches down a volume entitled *Tying One's Shoes*. The volume says such things as: “Take the

left free end of the shoelace in the left hand. Cross the left free end of the shoelace over the right free end of the shoelace...,” etc. When the little man reads the instruction ‘take the left free end of the shoelace in the left hand,’ he pushes a button on a control panel. The button is marked ‘take the left free end of a shoelace in the left hand.’ When depressed, it activates a series of wheels, cogs, levers, and hydraulic mechanisms. As a causal consequence of the functioning of these mechanisms, one's left hand comes to seize the appropriate end of the shoelace. Similarly, *mutatis mutandis*, for the rest of the instructions. The instructions end with the word ‘end.’ When the little man reads the word ‘end,’ he returns the book of instructions to his library. That is the way we tie our shoes (p. 63).

According to this explanation, shoe tying behavior is mediated by internal processes that employ “propositions, maxims, or instructions” (p. 76) regarding shoe tying, and Fodor urges us to take seriously the hypothesis that a whole host of motor, perceptual, and cognitive abilities are best explained in this way. Moreover, Fodor argues that this hypothesis yields an intellectualist account of know-how: knowing how to do something consists in having an ability whose exercise has the right kind of *etiology*: “If an organism knows how to X, then nothing is a simulation of the behavior of the organism which fails to provide an answer to the question ‘How do you X?’” (p. 75). In other words, the ability to ψ constitutes knowing how to ψ only if the causal processes that generate ψ -ing behavior represent a way to ψ . In the shoe tying example, the organism *knows how* to tie shoes since a representation of a way to tie shoes is part of the process that generates shoe tying behavior.

As Fodor was quick to acknowledge, the sorts of representations that are involved in his story about shoe tying are not usually, if ever, consciously available to the person tying her shoes. Rather, such knowledge is *tacit*, and ‘tacit knowledge’ is first and foremost “a theoretical term in psychology” (p. 76). Thus, if the kind of intellectualism that Fodor has in mind is vindicated, it will be vindicated by psychology, not philosophy or common sense. Furthermore, as his discussion makes clear, Fodor was not really interested in whether tacit knowledge is knowledge in the justified-true-belief sense that epistemologists have in mind—what he really cared about was the cognitive resources needed to explain the host of abilities and capacities that concern psychologists. According to Fodor, those resources include LOT and computationalism: the scheme of mental representation consists in primitive symbols that can combine to form more complex representations, and mental processes are computational processes defined over the syntactic structure of LOT expressions. On this view concepts turn out to be symbols of LOT, symbols that can combine with other symbols to form complex representations with propositional content (thoughts). Since the content of a complex representation is determined by the content of its constituents and their manner of combination in the way familiar from Tarski, LOT has a compositional semantics.³ Successfully explaining knowing-how in terms of knowing-that thus consists in showing how the manifestation of capacities is the result of processing expressions in LOT, expressions that have propositional contents.

The explanatory strategy that Fodor develops in “Tacit Knowledge” (and the intellectualist account of know-how that comes with it) has had a tremendous impact on

³ Fodor and Zenon Pylyshyn (1988) give a thorough exposition and defense of this position, which they call the “classical” account of concepts.

theorizing about the mind; indeed, we can think of the history of what John Haugeland (1985) calls Good Old-Fashioned Artificial Intelligence (GOFAI) as the history of defending the very intellectualism that Ryle discarded. In order for such a program to succeed, however, Ryle's argument against intellectualism needed refuting. So where did Ryle go wrong? The key is in appreciating that whatever know-how is required for theorizing can eventually be explained in terms of something that is itself not an exercise of the agent's know-how. In this way, Fodor's account is an instance of *homuncular functionalism* (Dennett, 1978; Lycan, 1987), which is a particular instance of functional analysis (Cummins, 1975). The complex capacities of a person are analyzed into simpler capacities of sub-personal systems. The capacities of the sub-personal systems are in turn analyzed into even simpler capacities until we reach capacities that require no know-how to be exercised. Fodor (1987) describes the strategy as follows:

According to this story, intelligent behavior typically exploits a 'cognitive architecture' constituted of hierarchies of symbol processors. At the top of such a hierarchy might be a quite complex capacity: solving a problem, making a plan, uttering a sentence. At the bottom, however, are only the sorts of unintelligent operations that Turing machines can perform: deleting symbols, storing symbols, copying symbols, and the rest. Filling in the middle levels is tantamount to reducing—analyzing—an intelligent capacity into a complex of dumb ones; hence to a kind of explanation of the former...At the very top are states which may correspond to propositional attitudes that common sense is prepared to acknowledge...But at the bottom and middle levels there are bound to be lots of

symbol-processing operations that correspond to nothing that people—as opposed to their nervous systems—ever do (pp. 23-24).

Returning to the example of language, why did Chomsky hypothesize that children possess a tacit *theory* of language? As Ryle pointed out, our know-how consists in our open-ended ability to satisfy criteria (e.g. our ability to produce an indefinite number of novel yet grammatically correct sentences), and Chomsky argued that the best explanation of this ability is that we possess general knowledge of the principles of language—“laws” of language, so to speak—and we apply this general knowledge when we manifest our capacities. According to the computational account sketched above, theories are sets of sentences in LOT, sentences which encode the laws or principles of the domain being cognized, and applying a theory consists in inferring conclusions about particular cases from those sentences. “At the bottom,” however, the symbol manipulating capacities of the hardware will constitute neither know-how nor know-that, since everything cognitive is ultimately analyzed into something non-cognitive.

By the mid 1980’s, however, many researchers in cognitive science and artificial intelligence (as well as philosophers) grew skeptical that LOT accurately describes the structure of mental representations. The emergence of connectionism and computational neuroscience brought an alternative to LOT into play, and soon it became clear that many kinds of representations, including scale models, pictures, graphs, diagrams, maps, synaptic weight configurations and partitioned activation spaces and activation vectors could be evaluated for accuracy along many, often simultaneous and competing, dimensions, but not for truth. For example, consider Paul Churchland’s characterization of the motive behind eliminative materialism:

The real motive behind eliminative materialism is the worry that the ‘propositional’ kinematics and ‘logical’ dynamics of folk psychology constitute a radically false account of the cognitive activity of humans, and of the higher animals generally. The worry is that our folk conception of how cognitive creatures represent the world (by propositional attitudes) and perform computations over those representations (by drawing sundry forms of inference from one to another) is a thorough-going misrepresentation of what really takes place inside of us (1992, p. 40).

For Churchland, however, the abandonment of the propositional attitudes does not signal the abandonment of cognition, representation, and knowledge. Rather, representation and cognition need to be understood in new terms:

The basic idea is that the brain represents the world by means of very high-dimensional activation vectors, that is, by a pattern of activation levels across a very large population of neurons. And the brain performs computations on those representations by effecting various complex vector-to-vector transformations from one neural population to another (p. 41).

How do we tie our shoes, according to Churchland? There is a massively recurrent neural network that resides in your motor cortex. Given the appropriate inputs, the weight matrix of this network induces a trajectory of activation within the hidden activation space of the network, a trajectory that represents a sequence of motor outputs appropriate for shoe-tying. Because these are motor neurons, the trajectory distinctive of shoe-tying will result in orchestrated muscle movements that constitute the act of tying one’s shoes.

Put in the context of cognitive science generally, the debate between those in the GOFAI camp and those in the connectionist camp is a debate over the best explanation of the target phenomena of psychology: capacities such as speaking a language, finding one's way home, seeing depth, tying one's shoes, etc. Both camps are committed to explaining capacities and abilities in terms of intentionally characterized internal states and processes; the real disagreement is over what the form and content of internal representations must be like in order to do their explanatory work.

So what, if anything, does this have to do with intellectualism, understood as the thesis that knowing-how can be reduced to, or is a species of, knowing-that? According to both Fodor and Churchland, the debate within cognitive science does have implications for intellectualism. As we've seen, Fodor claims that his proffered explanation for how we tie our shoes supports intellectualism, since the causal explanation of shoe-tying behavior involves representations of instructions for tying shoes, instructions that provide an answer to the question "How do you tie your shoes?" For Fodor, it is in virtue of shoe-tying behavior having this sort of etiology that it is correct to say that someone 'knows-how' to tie shoes. On this point, Churchland may have no disagreement with Fodor. For example, Churchland (1989(a)) argues that the detailed motor representations of the limb movements involved in a golf swing, the representations that subserve the ability to play golf, constitute the golfer's know-how. Insofar as we can think of the processes that generate golf-playing behavior as answering the question "How do you play golf?" there appears to be no substantive difference between Fodor and Churchland. What, then, divides them? Clearly, the disagreement is over the structure and content of the representations that produce behavior. Fodor thinks

that the relevant representations are sentences in LOT, sentences that have structured propositions as their contents. Churchland thinks that the relevant representations are partitioned activation spaces, spaces that represent in the way that maps do and that are not candidates for truth-conditional analysis. Thus, if by intellectualism we mean merely that the explanation of ψ -ing involves representations that provide an answer to the question “How does one ψ ?”, then both Fodor and Churchland are intellectualists. If by intellectualism we mean that the relevant representations have propositional contents, then Fodor, but not Churchland, is an intellectualist.

Why do Fodor and Churchland bother discussing know-how, and how much of the cognitive science hangs on the discussion? Perhaps—though it is not likely—Fodor and Churchland think that their favored theory about the architecture of cognition in some way supports, or finds support from, the folk understanding of know-how and the truth conditional analyses of know-how ascriptions such understanding supports. If so, then we say a pox on both their houses; just as physics need not support, or find support from, the folk understanding of motion, cognitive science need not support, or find support from, the folk understanding of know-how. If we bracket this motivation for speaking in terms of ‘know-how,’ however, then talk of ‘know-how’ may be rather harmless. As far as the disputes in cognitive science are concerned, everyone pretty much understands what is being asserted when someone claims that knowing-how cannot be reduced to knowing-that: the correct explanation of certain abilities will not involve processing representations with propositional contents.

Let’s take stock. Practically every party to the dispute between GOFAI and connectionism accepts that knowing-how involves the possession of abilities.

Furthermore, accepting this does not depend on accepting some controversial premises about the concept of know-how; using ‘know-how’ to refer to abilities of a certain sort is as much a matter of stipulation as anything else. Thus, in the context of cognitive science, *intellectualism is compatible with identifying know-how with certain abilities, capacities, and skills*. The real controversy is over the explanation of skills and capacities: are they best explained in terms of representations with propositional content, or are they best explained in terms of, e.g., partitioned activations spaces and prototype points, representations that lack propositional content? The reason connectionists are often aligned with Ryle (Moffett and Bengson, 2007) is that they agree that we will not account for many abilities in terms of representations with propositional content (knowledge-*that*); but in so agreeing, connectionists are not embracing Ryle’s identification of know-how with *mere* abilities, for connectionism accounts for abilities in terms of something “cognitive” as much as GOFAI does. For example, consider how Bechtel and Abrahamsen (1991) contrast their position with Ryle’s:

Ryle’s approach to analyzing knowing how was behavioristic; he treated this knowledge as manifest in our actions and so not to be understood in terms of something hidden and internal... Knowing how to perform an action consisted in a disposition to perform that action when appropriate circumstances arose. For cognitivists, however, merely saying that someone has a disposition to behave in a certain way does not suffice; what is sought is an explanation of the internal mechanisms involved. One possibility is that intelligent dispositions may not be explainable in cognitive terms; it may have to do simply with the physiological conditions accruing in one’s body. But insofar as we speak of knowing how to do

these things, and of these being things we learn how to do, it seems plausible to attempt to give a cognitive account of what such knowledge consists in, that is, and account in terms of one's mental activities (pp. 152-153).

What is particularly striking about this passage is that, although Bechtel and Abrahamsen are open to the possibility that capacities and skills (know-how) may be given a non-cognitive explanation, calling these capacities and skills *knowledge* suggests a cognitive one; and *this* suggests that if know-how were not explained in cognitive terms, then perhaps calling these capacities and skills 'know-how' would be inappropriate. Since this is a *very* un-Rylean thing to say, why are Bechtel and Abrahamsen aligned with Ryle? Because in the context of the dispute over intellectualism in cognitive science, Bechtel and Abrahamsen are anti-intellectualists. But their anti-intellectualism is not anti-*cognitivism*, but anti-*propositionalism*. Bechtel and Abrahamsen accept that the abilities they are calling 'know-how' have a basis in mental representations and processes; they just don't think the mental representations relevant to explaining skills and capacities will look like LOT. If by intellectualism we mean simply a rejection of behaviorism, then almost everyone in cognitive science is an intellectualist.

As we've seen, however, using 'know-how' and 'know-that' to formulate claims in cognitive science may provide an occasion to poach. As such, if there are independent reasons for thinking that the intellectualist account of know-how offered by Stanley and Williamson has nothing really to do with intellectualism in cognitive science, then it might be tempting to drop the talk of 'know-how' in cognitive science—it would simply do no theoretically important work. Of course, there are such reasons. The first concerns Fodor's use of 'knowledge.' As we pointed out above, Fodor's commitment to tacit

knowledge is not a commitment to knowledge in the epistemologist's sense; what he really cared about was psychological explanation and the LOT he thought psychology needed. A second—and perhaps more important—reason for thinking that Stanley and Williamson's intellectualism is irrelevant to intellectualism in cognitive science concerns the resources required for intellectualism to provide a plausible explanation of behavior. Recall Fodor's explanation of how we tie our shoes: the little man in your head fetches a *book* about tying shoes. Why a book? Because according to intellectualism, the sentence in LOT that would express the relevant Russellian proposition (the kind of proposition that Stanley and Williamson would think sufficient for know-how) has to have the same form or structure as the proposition expressed, *and nobody—not even Fodor—thinks that tokening such a sentence in LOT would come anywhere close to being sufficient for an intellectualist explanation of shoe tying*. The problem, in short, is that the science requires that the structure of the representations (e.g., LOT sentences) mirror the structure of the explanatorily relevant content (e.g., a way to tie shoes), and Stanley and Williamson's account does not deliver this. Having a sentence in LOT that contains a demonstrative (e.g., 'this') that picks out a way to tie shoes will be of no help when it comes to actually getting shoes tied since, from the perspective of the science, explanatorily relevant content cannot outstrip computational form. This, by the way, is part of the reason why getting GOFAI to succeed has proven so *hard*.

The tactic of divorcing the kind of intellectualism that Stanley and Williamson defend from intellectualism in cognitive science has its costs, however. Much of the interest in intellectualism presumably rests on the assumption that establishing intellectualism would have profound consequences for our understanding of the cognitive

architecture of humans. Thus, to the extent that one insulates intellectualism from the science, a vindication of intellectualism no longer looks interesting or significant.

Without such insulation, however, the arguments that Stanley and Williamson adduce in support of intellectualism are irrelevant. This is no surprise: it is simply incredible that we could establish the truth of intellectualism, understood as an hypothesis in cognitive science, on the basis of the sorts of considerations they offer.

V. Conclusion

Intellectualism is (or should be) a substantive hypothesis about the nature of human cognition, and anti-intellectualists are betting that the best explanation of know-how will not involve LOT and propositions (GOFAI). Furthermore, though anti-intellectualism may turn out to be false, it will not be refuted by the kinds of considerations offered by Stanley and Williamson. Looking back at Ryle through the eyes of contemporary cognitive science, we can see that perhaps Ryle's most important contribution to the philosophy of mind was his insight that there is an intimate connection between *knowing* and *doing*, and that locating the knowing in an "inner place" while locating the doing in an "outer place" could not do justice to the intimacy of that connection. Ryle's solution, of course, was to say that the knowing *is* the doing (or the tendency to do). But there is nothing wrong with positing inner mental processes to explain capacities and skills; as the failures of GOFAI have shown, what may be wrong is (1) thinking that the kind of knowledge that explains our capacities and skills is propositional, and (2) thinking that understanding abilities, capacities, and skills is peripheral to understanding cognition. Understood this way, the problem with conceiving of knowing and doing in terms of

“inner” and “outer” has nothing really to do with traditional Cartesian dualism *per se*, since the problem arises for any account of the mind that would distinguish having the knowledge relevant to abilities from applying that knowledge. Part of the appeal of connectionism is that it offers an account of how and why knowing and doing are so intimately connected, since the knowledge relevant to know-how is in the very connection weights that sub-serve the exercise of abilities, capacities, and skills.

References

Bechtel, W. and A. Abrahamsen. 1991. *Connectionism and the Mind*. Cambridge (MA): Basil Blackwell.

Chomsky, N. 1959. Review of Skinner’s *Verbal Behavior*. *Language* 35: 26-58.

Churchland, P.M. 1989(a). “Knowing qualia: A reply to Jackson.” In *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge (MA): MIT Press. Reprinted in *On the Contrary*. Cambridge (MA): MIT Press.

--1989(b). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge (MA): MIT Press.

--1992. “Activation Vectors vs. Propositional Attitudes: How the brain represents reality.” *Philosophy and Phenomenological Research* 52. Reprinted in *On the Contrary*. Cambridge (MA): MIT Press.

--1998. *On the Contrary*. Cambridge (MA): MIT Press.

- Cummins, R. 1975. "Functional Analysis." *Journal of Philosophy*, 72: 741-765.
- 1991. "Methodological Reflections on Belief." In *Mind and Common Sense*. R. Bogdan (ed.). New York: Cambridge University Press.
- Davidson, D. 1967. "Truth and Meaning." *Synthese* 17: 304-323.
- Dennett, D. 1978. *Brainstorms: Philosophical Essays on Mind and Psychology*. Cambridge (MA): MIT Press.
- Fodor, J. 1968. "The Appeal to Tacit Knowledge in Psychological Explanation." *The Journal of Philosophy* 65: 627-640. Reprinted in *Representations*. 1981. Cambridge (MA): MIT Press.
- 1981. *Representations*. Cambridge (MA): MIT Press.
- 1987. *Psychosemantics*. Cambridge (MA): MIT Press.
- Fodor, J., and Z. Pylyshyn. 1988. "Connectionism and Cognitive Architecture: a Critical Analysis." *Cognition* 28: 3-71
- Haugeland, J. 1998. *Having Thought*. Cambridge (MA): Harvard University Press.
- Lycan, W. 1987. *Consciousness*. Cambridge (MA): MIT Press.

Millikan, R. 2000. *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.

Moffett, M. and J. Bengson. 2007. "Know-how and concept possession." *Philosophical Studies* 136: 31-57.

Noë, A. 2005. "Against Intellectualism." *Analysis* 65: 278-290.

Ryle, G. 1949. *The Concept of Mind*. Chicago: University of Chicago Press.

Stanley, J. and T. Williamson. 2001. "Knowing How." *Journal of Philosophy* 98: 8: 411-444.