CONSCIOUSNESS, CONTEXT, AND KNOW HOW

1.) INTRODUCTION

Most philosophers since Ryle agree that we attribute knowledge how to individuals and that attributions of knowledge how are as explanatorily relevant to our understanding of mentality as attributions of beliefs and desires. Even a cursory glance at the literature in cognitive science reveals the pervasiveness and perceived centrality of knowledge how to contemporary cognitive science. Work in cognitive science focuses, not merely upon representations of the world which may be true or false, like beliefs, but on human abilities. Computer scientists like Herbert Simon more often model the human ability to play chess than the capacity to have beliefs about chess or to represent current states of play. (de Groot, A. & Gobet, F., 1996; Gobet, F., and Simon, H. A. 1996, 1998; Waters, Gobet, and Leyden 2002) Research into human abilities produces programs like ACT, SOAR, and IMPROV that are best understood as models of knowledge how, and not models of beliefs. (Anderson, J., Kline, P, & Beasley, C., 1979; Rosenbloom, P., Laird, J., McDermott, J., Newell, A., & Orciuch, E., 1985; Laird and Pearson, forthcoming) These programs are sets of

condition action rules, bits of practical reasoning, not to be confused with representations of the world that are true or false. These programs are, nevertheless, evaluated epistemically. A given chess move is not true or false, but better or worse relative to the goal of winning, more or less justified given the circumstances. The program that generated the move is not true or false, but more or less optimal. Some connectionist researchers depart even more strongly from the staples of belief-desire psychology by claiming that in their connectionist models of abilities, as in the brain itself, "...almost all knowledge is *implicit* in the structure of the device that carries out the task rather than *explicit* in the states of the units themselves." (Rumelhart, 1989, p.136) A similar acknowledgment of the importance of knowledge how is evinced in cognitive psychology. Cohen and Squire as well as Stillings equate the important declarative/procedural knowledge distinction in psychology to the knowing that/knowing how distinction in philosophy. (Cohen and Squire 1980; Stillings, et al 1987, pp.18-21)

Despite the seeming acceptance of the existence of know how and its seeming importance to the scientific study of human cognition, philosophers devote little attention to knowledge how. Few accept Ryle's dispositional account. Fewer still offer theories themselves. Without a serious and informed theory of knowledge how philosophers lack insight into one of the more important explanatory constructs of contemporary cognitive science. Equally lacking is a perspective from which to understand and evaluate the claims of cognitive scientists regarding knowledge how. One might well describe the need for a theory of know how as paramount in both the philosophy of mind/cognitive science and in epistemology.

Philosophy is not entirely bereft of perspectives on know how. One pervasive reaction to Ryle's 1949 conception of knowing how accepts the existence of knowing how while denying Ryle's contention that it was fundamentally dispositional in nature.[1] These theorists re-describe knowing how in terms of practical reason, i.e., in terms of conscious, general reasoning about practical matters as opposed to theoretical ones. Such a position is nicely summarized as follows:

> Inter alia, knowing how in the strong sense to play football is knowing the rules of the game, but a statement of the rules of the game is not a theoretical statement but a description of a set of rules of *practice*, and mastery of the rules brings with it an understanding of an activity rather than a theory. Statements of the rules of a game are essentially of relations between *prescriptions* rather than descriptions requiring a grasp of practical rather than theoretical discourse. (Carr 1981, pp.60-1)

Essential to this view is the idea that knowing how, like knowing that, is really a subset of represented knowledge, where each knowledge set has its own characteristic logic. Knowing that, or theoretical knowledge, is propositional in nature. Practical knowledge, though also consciously represented, is a knowledge of procedures. Knowing what we might express as, "combustion engines generate force by igniting compressed gases within closed, expandable cylinders," counts a theoretical knowledge. Knowing what we might express as, "to start standard U.S. production cars, turn the key clockwise," is practical knowledge. As David Carr (1981) speculates:

> Learning to perform a given complex task or coming to know how to do it is essentially a matter of learning to reason practically...a matter of instructing him in practical reasoning by means of practical directives;... (Carr 1981, p. 60)

Indeed, David Carr (1979, 1981, 1981a) represents both one of the clearest criticisms of Ryle and the most significant recent example of the practical-knowledge analysis of

knowledge how. Like Ryle and myself, Carr holds that ascriptions of knowledge how play a role in explanations of human cognitive ability, telling readers that

> it seems to me that our common ascriptions of knowing how and ability to agents, enjoy the same sort of explanatory role in human affairs as has been claimed for reports of belief and other mental acts. (1979, p.395)

In this paper I acknowledge the crucial importance of knowledge how in our contemporary understanding of the mind, and offer a treatment of knowledge how capable of providing insight into the use of knowledge how in contemporary cognitive science. Specifically, I look at Carr's analysis of the necessary conditions for knowledge how. I claim that Carr's conditions fail to capture the distinction he himself draws between ability and knowing how. Moreover, Carr ties knowing how to conscious intent, and to an explicit knowledge of procedural rules. I argue that both moves are mistakes. Together these mistakes render Carr's theory an inadequate account both of common ascriptions of knowledge how and of widely accepted ascriptions of knowledge how to human beings within explanations in cognitive science. I note that Carr's conditions fail to capture intuitions (that he shares) regarding the ascription of know how to persons who lack ability. Finally, I consider the position advocated by Stanley and Williamson (2001), which seems avoid Carr's commitments to conscious intent and explicit knowledge while still maintaining that"knowledge-how is simply a species of knowledge-that." (Stanley and Williamson 2001, p. 411) I argue that Stanley and Williamson's attempt to frame a reductionist view that avoids conscious beliefs and explicit knowledge of procedural rules is both empirically implausible and explanatorily vacuous.

In presenting a case against Carr's theory as well as Stanley and Williamson's more

contemporary alternative, I challenge the presuppositions of the most pervasive responses to Ryle in the philosophic literature, what might be described as "the received view" on the nature of knowledge how. I also establish several facts about knowledge how. First, given the theory of cognition current in cognitive science, neither conscious intent nor explicit representation (much less conscious representation) of procedural rules are necessary for knowledge how. I argue that Carr's and similar analyses fail to capture the necessary conditions for knowledge how because knowledge how requires the instantiation of an ability and of the capacities necessary for exploiting an ability--not conscious awareness of purpose or explicit knowledge of rules. Second, contemporary reformulations of reductionist accounts neither reduce knowledge how to propositional knowledge nor do they solve difficulties of traditional reductionist analyses. Third, one must understand knowledge how as task-specific, i.e., as presupposing certain underlying conditions. Conceiving of know how as task-specific allows one to understand ascriptions of know how in the absence of ability as counterfactual ascriptions based upon underlying competence.

2.) CARR ON KNOWLEDGE HOW, INTENTION, AND RULES

Carr proposes the following necessary conditions for knowing how:

> A knows how to ϕ only if:
> (CD1) A may entertain ϕing as a purpose,
> (CD2) A is acquainted with a set of practical procedures necessary for successful ϕing,
> (CD3) A exhibits recognizable success at ϕing. (1981, p.58)

For Carr "entertaining ϕing as a purpose," consists in consciously intending to ϕ, understood as the act of ϕing. Carr stipulates that acquaintance with a set of practical procedures consists exclusively in possessing conscious directives (rules) about practice: Carr has a

long, essentially correct argument that an agent can know how to ϕ only if they can entertain ϕing as a purpose. Carr's argument is as follows:

> (C1) Attributions of ability can always be substituted *salva veritate* for equivalent ability descriptions.
>
> (C2) Attributions of *knowing how* cannot always be substituted *salva veritate* for equivalent *know how* descriptions.
>
> (C3) So, *knowing how* to ϕ must differ from having the ability to ϕ.
>
> (C4) The difference in such cases is always the agent's ability to entertain ϕing as a purpose.
>
> (C5) Hence, being able to entertain ϕing as a purpose must be necessary for knowing how to ϕ.

Carr offers several examples to illustrate his distinction between ability and knowledge how in (C1) and (C2) as well as his assertion in (C4).[2] In one such example, a dancer performs "Improvisation #15," a dance piece of his own creation. One would say that the dancer both knows how to perform "Improvisation #15" and has the ability to perform "Improvisation #15." Unbeknownst to the dancer, who has no knowledge of semaphore whatsoever, the dance matches a semaphore rendition of Gray's "Elegy" move for move. Hence, the dancer would seem to have the ability to perform a semaphore version Gray's "Elegy" without knowing how to perform it. That is, he would meet one's requests for a semaphore recital of Gray's "Elegy" with puzzled looks, statements to the effect, “Sorry, I don’t know that one,” etc.. Yet, such a performance is well within the dancer’s abilities.

Carr suggests that examples like the dancer indicate that the ability to consciously intend to ϕ differentiates having the ability to ϕ and knowing how to ϕ [Hence, (C4)]. In fact, Carr argues for the correctness of his necessary conditions, in part, because they are

sufficient to distinguish know how from ability in such cases. (1981, pp.58-9) It is instructive, however, to note that the above example satisfies Carr's necessary conditions for knowledge how with respect to Grey's ‘Elegy’. In the above example, the dancer may entertain performing a semaphore recital Gray's "Elegy" as a purpose. All he need do is merely intend to perform a semaphore recital of Gray's ‘Elegy’. He may think, "For my next piece, I want to perform a semaphore recital of Gray's 'Elegy'," though falsely supposing that a different dance is actually a semaphore recital of Gray's ‘Elegy’. He thus satisfies clause (CD1). Though the dancer does not realize it, he is acquainted with a set of practical procedures necessary for successfully performing a semaphore version of Gray's ‘Elegy’. Hence, he satisfies (CD2). He also exhibits recognizable success at performing a semaphore version of Gray's ‘Elegy’, satisfying (CD3). The dancer can satisfy the latter two conditions without an awareness that they are satisfied. So, the dancer satisfies all of Carr's conditions in a case where there is ability, but no knowledge how. As a result, the conditions fail to capture Carr's intended distinction between knowing how and ability.

3.) DIAGNOSIS

Why does Carr fail to capture his own distinction between know how and ability? I claim that Carr misdiagnoses the difference in the above example. Carr correctly supposes that knowledge how ascriptions are tied essentially to intentional characterizations. Hence, as in the case of belief and other propositional attitudes, knowledge how ascriptions cannot be substituted *salva veritate* for alternative knowledge how ascriptions--even when the substituted know how ascriptions pick out the same sequence of actions. Though Carr makes this important observation, his supposition as to the proper explanation for the substitution

failure, viz., that it lies in conscious intent, proves somewhat less laudable. Though awareness or consciousness of purpose has played a large role in the discussion so far, I claim, in contrast to Carr, that awareness or consciousness is an inessential property of the examples. One might express my criticism by noting that the dancer has the ability to perform a semaphore recital of Grey's 'Elegy', but lacks the functional connection between the ability to dance a semaphore recital of Grey's 'Elegy' and the dance so described/understood. It is this lack of a functional connection, and not an inability to intend to perform a semaphore recital that prevents knowledge how in the case of the dancer.

Knowledge how, on my understanding, is the successful and reliable linkage between ability, represented knowledge, and task-specific conditions. The cognizer, in other words, must instantiate a process or set of processes that generates actions constrained in ways appropriate to the task at hand in the appropriate circumstances. Thus, the dancer must not only have the ability to generate the relevant movements, he or she must have a reliable mechanism which underlies their generation of those movements as a response those situations in which an appropriate response to the situation is to dance a semaphore recital of Grey's 'Elegy'. It is this underlying reliable mechanism integrating ability, knowledge, and sensitivity to appropriate circumstances to which I refer when claiming the dancer lacks a functional connection between the ability to dance a semaphore recital of Grey's 'Elegy' and the dance so described. Carr perceptively notes that conscious intent mediates many such functional connections. However, functional connections between an ability, knowledge, and task-specific conditions are often, *but not always*, mediated by conscious intent.

To further illustrate this point, consider the case of a woman who irrevocably loses

her sight. On a literal reading of Carr's conditions, she no longer knows how to read because she violates (CD3).[3] Her loss of knowledge how, however, does not result from a loss of the ability to entertain reading as a purpose, nor has she lost her acquaintance with the practical procedures necessary for reading. Indeed, though she must learn to read all over again, she does not need to relearn spellings, sentence structures, and so on. She must learn how to exploit her knowledge of spelling, sentence structure, and so on by learning a new, non-visual code for letters (i.e., she must learn braille) and integrating it with her existing knowledge. She no longer knows how to read because she is no longer able to exploit visual patterns as inputs, i.e., she lacks a functional connection between her ability and conditions underlying the performance of the task.

To illustrate my point in a different context, consider two games, "Puzzle Master" and "Sort the Letters." Puzzle Master is a hand held puzzle like those one finds in toy stores.[4] To play Puzzle Master one moves the squares into the blank area so as to unscramble the message:

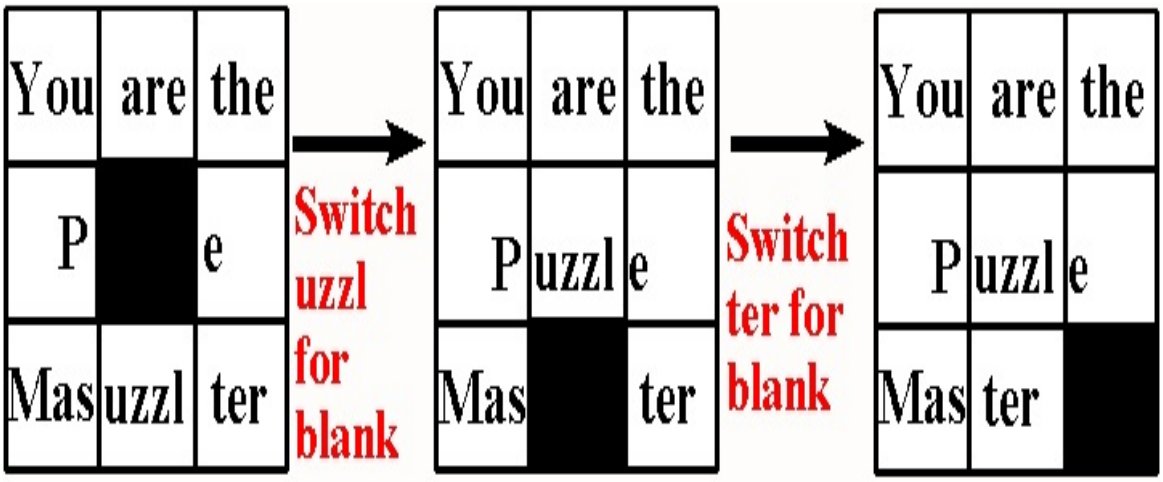

Sort the Letters is a pen and pencil game. One is given a list like the ones below.

One then tries to put the letters in the second half of the right column in alphabetical order. One sorts the letters by switching 'I', *and only* 'I', for other letters in the second half of the right column. In addition, one can only switch 'I' with letters, the number of which is given in the left hand column of the same row in which one has placed 'I'. So, the following is a series of legal move from the position given to a win:

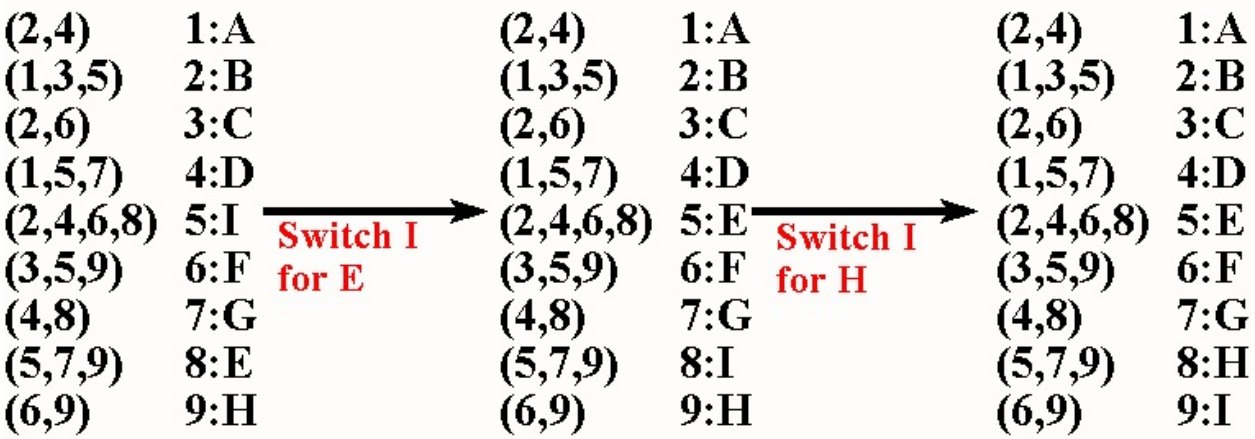

Now, suppose that I initially start out having only the bare prerequisite abilities and knowledge of the rules of both Sort the Letters and Puzzle Master. One might claim that I know how play both games, since I can make legal moves, but it is clear that I do not know how to play to win at either game. That is, I don't know how to sequence moves in such a way as to make it more likely that I achieve the winning state. However, if I learn an algorithm to find a set of moves that successfully sorts the letters, I seem to thereby come know how to play to win at Sort the Letters, though I am still terrible at Puzzle Master. That is, I embody a process that reliably generates a winning set of moves for Sort the Letters.

However, as one might suspect, Puzzle Master and Sort the Letters are formally equivalent games. There is a one-to-one interpretation mapping states of Puzzle Master to

states of Sort the Letters that preserves the formal properties of each game. So, my algorithm could also generate a winning set of moves for Puzzle Master, if I enter initial configurations for Puzzle Master using the following interpretation:

1 A	2 B	3 C
4 D	5 E	6 F
7 G	8 H	9 I

As with the case of the blind woman, the difference between my knowing how to generate a winning series of moves for Sort the Letters, and my knowing how to generate the winning moves for Puzzle Master is the ability to exploit my knowledge in a specified set of circumstances, in this case, my ability to exploit the algorithm via an appropriate interpretation. If I use the Puzzle Master interpretation, it allows me to exploit the algorithm so that I know how to generate a winning series of moves for Puzzle Master. If I use the Sort the Letters interpretation, I know how to generate a series of winning moves for Sort the Letters. If I use neither interpretation or a third interpretation, then I may not know how to play to win at either game.

4.) CONSCIOUS INTENT AND KNOWLEDGE HOW

The cases considered so far are ones where one might plausibly suggest that conscious intent together with conscious reasoning using practical directives provides the necessary functional connection between an ability and the context in which the person can

exercise the ability. Carr's necessary conditions clearly require that a person be conscious in order to manifest knowledge how. In fact, Carr's rationale for his (CD1) is that "our familiar understanding of exercises of know how is that they are expressions of deliberate purposes,...." (1981, pp.58-9) However, a wide variety of cases provide concrete *prima facie* evidence that one can manifest know how in ways requiring neither conscious intent nor conscious reasoning using practical directives. These examples undermine the relation between consciousness and know how that Carr intends to capture using (CD1). I start with examples to which philosophers might feel more sympathetic and work my way to, in my opinion, equally valid cases that philosophers might challenge. I note at the onset, that acceptance of only one of the examples is sufficient to prove my general point, that conscious intent is not necessary for the possession and/or exercise of knowledge how.

For example, experts are certainly candidates for ascriptions of knowledge how with regard to their field of expertise. Physicists know how to solve physics problems. Grand masters know how to play chess. If know how consists in conscious or consciously accessible beliefs, and the exercise of know how consists of conscious, general reasoning about practical matters, then one would expect these experts to be able to easily articulate the practical directives guiding their actions (know how). However, it is well-known in psychology and artificial intelligence that experts have often have great difficulty in consciously accessing and in expressing their knowledge. The hypothesis that knowledge how can be implicit in the operations of a system is often employed to explain why it can prove difficult for these experts, even experts in abstract, contemplative areas such as physics, to generate an explicit set of rules which govern their expert behavior. Indeed, elicitation of such "tacit knowledge" from

experts is an important step in developing expert systems.[5] Similar difficulties emerge when trying to illicit the practical directives of commonsense that seem to inform our ordinary actions. (Reber and Lewis, 1977; Bliss and Ogborn 1994; Sun 2005; Ellis 2005)

Likewise, if knowledge how consists in conscious or consciously accessible representations of practical directives, and the exercise of knowledge how consists of conscious, general reasoning about practical matters, then one would expect a reasoner's introspective reports to prove reliable regarding their thought processes and practical directives during such tasks as problem solving. In stark contrast, psychologists have long known that subjects are notoriously unreliable in their introspective reports of their thought processes during problem solving. (Nisbett and Wilson, 1977) For example, Mandler tells readers that "there are many systems that cannot be brought into consciousness, and probably most systems that analyze the environment in the first place have that characteristic. In most of these cases, only the products of cognitive and mental processes are available." (Mandler 1975, p. 245) Nisbett and Wilson (1977) summarize their extensive review of the psychological literature by saying that with regard to high level reasoning (such as that involved in problem solving, making evaluations, making judgements, and initiating actions) subjects; (1) often had no awareness of environmental factors or information shown by analysis to have importantly influenced their responses, (2) often remained unaware of their responses, and (3) were unaware of any connection in their thought between such environmental/problem factors and their responses. (Nisbett and Wilson, 1977 p. 231) Moreover, Nisbett and Wilson further conclude that the best explanation of accurate reports regarding cognitive processes is that the subjects construct an explanation based upon salient environmental or problem factors and

accounts of reasoning held by the subjects to be plausible cognitive processes. For example, if one presents subjects with a tray of identical objects and asks them to choose one, the factor most predictive of subject choice is their handedness. However, subjects rarely cite handedness as a basis for their choice. Instead subjects are much more likely to offer accounts of their mental processes during the choice which relate the choice to some (false) reason for preferring one of the objects to the others, such as, "the one on the right looked newer."

The constant and undeniable repetition of such experiments as noted by Nisbett and Wilson across a wide range of tasks leads Bargh and Chartrand to reflect 22 years later that,

> To consciously and willfully regulate one's own behavior, evaluations, decisions, and emotional states requires considerable effort and is relatively slow. Moreover, it appears to require a limited resource that is quickly used up, so conscious self-regulatory acts can only occur sparingly and for a short time. On the other hand, the nonconscious or automatic processes we've described here are unintended, effortless, very fast, and many of them can operate at any given time. Most important, they are effortless, continually in gear guiding the individual safely through the day. Automatic self-regulation is, if you will, thought lite—"one third less effort than regular thinking" (Gilbert, 1989, p. 193). (Bargh and Chartrand 1999, p. 476)

So, we act unconsciously when we are awake. We also act occasionally when we are unconscious. Somnambulism (sleepwalking) is a well-known disorder that occurs most often during periods of brain activity associated with deep, dreamless sleep. Hence, one finds both physiological (EEG) and behavioral evidence that a sleepwalker is not conscious.

> The essential features of this disorder are repeated episodes of a sequence of complex behaviors that progress to leaving the bed and walking about, without the person's being conscious of the episode or later remembering it. (American Psychiatric Association 1987, p.311)

Though such episodes are relatively infrequent, sleepwalkers have driven cars and even committed murder. (Luce and Segal 1966, pp.134-37) During sleepwalking episodes, sleepwalkers manifest knowledge how (for example, knowledge of how to drive a car). Yet, as in above-discussed cases, conscious intent plays no role in the manifestation of this

knowledge how. Again, Carr's requirement that the system have the ability to consciously entertain ϕing as a purpose looks incorrect. The system need only be capable of exploiting the ability that allows it to reliably perform the task.

Furthermore, Schacter (1989), in reviewing psychological literature on procedural knowledge, notes that;

> ...it has been demonstrated repeatedly that even profoundly amnesic patients, such as the well-known case H.M., can show normal or near-normal learning of various perceptual and motor skills without any conscious memory for the experiences of learning. ... The dissociations discussed so far indicate clearly that memory for recent experiences can be revealed in performance without any conscious experience of remembering, and also indicate that various kinds of knowledge can be expressed in the absence of conscious experiences of perceiving, identifying, or knowing. (pp. 357-362)

Indeed, persons with bilateral hippocampal damage, like H.M., show a normal or near-normal learning curve both on tasks involving problem solving and tasks involving motor and/or perceptual skills, like the Tower of Hanoi puzzle and mirror drawing. They display this learned knowledge even though they exhibit no indications of possessing a conscious knowledge of, for example, particular moves or their relative merit or demerit. (Cohen and Corkin 1981) Strong dissociations between conscious intent and knowledge how clearly show that one can acquire and exercise knowledge how without either conscious intent or conscious reasoning using practical directives.

Similarly, many animals exhibit a knowledge of how to perform acts requiring relatively complex and adaptable interactions with the environment. These animals range from animals, like monkeys, that might well possess the ability to consciously represent knowledge or consciously entertain purposes, to animals that seem unlikely to possess the ability to consciously represent knowledge or consciously entertain purposes. The caddis fly

larvae, for example, though having only a few hundred neurons builds characteristic cases which can protect them from predators, increase water flow over gills, screen water, and even capture food. They exhibit definite preferences in selecting materials for use, and will rebuild parts of cases that have been altered in accordance with their species-specific blueprint, even when such alterations damage function in unobvious ways. (Hansell 1984, 1972; Griffen 1992). Of course, one might dismiss such cases as mere anthropomorphizing (Stanley and Williams 2001) or claim that such attributions never rise above ability to achieve knowledge how (Carr 1981). However, given the widespread use and success of studying human cognition through the study of animal cognition, the complexity of the tasks, the fact that the animals often acquire the abilities through learning and problem solving, and they often teach others, such moves strike me as quite implausible.[6]

Finally, a similar case can be made against Carr using explanations from cognitive science of clearly nonconscious cognitive processes such as early visual processing. I must preface my remarks here by acknowledging certain differences between Carr's project and my own. Though Carr clearly wishes to capture a non-trivial sense of knowledge how that underlies psychological explanations, he also wishes to capture what he perceives to be the normal usage of the term in ordinary language. To this end, Carr claims that,

> the logic of knowing how differs from the logic of ability with respect to simple tasks or actions. Although we may say a man is able (has the physical power) to perform a simple or basic action that he performs, it is odd to describe many basic actions as exercises of knowing how. ... "He knows how to raise his arm," then, transgresses good sense (except in exceptional circumstances largely irrelevant to present concerns) in a way that "he is able to raise his arm" does not. (1981, p.53)

Carr might, in light of the above claim, argue that the sorts of abilities studied by vision

researchers are simple, and hence, not appropriately explained by reference to know how. I think that such an assertion is false. Research in vision has shown that the sorts of visual tasks I discuss are, in fact, remarkably difficult, requiring that a system solve several computationally complex problems (some in parallel, some in proper sequence) and integrate the individual solutions into the final solution. Upon studying vision it becomes unsurprising that large portions of our brain seem to be devoted exclusively to such "simple" tasks. So, Carr's simplicity criterion fails to exclude the sorts of tasks I consider on an informed notion of the actual level of difficulty of cognitive tasks. Indeed, once one understands the complex nature of visual processing, it does seem "to involve the sort of distinctive techniques and procedures" Carr associates with knowledge how. (Carr 1981, p.53) Second, I would claim that even on Carr's own construal, his simplicity criterion is false. Carr informs us that it "transgresses good sense" to talk of knowing how to walk or raise one's arm. However, any child psychologist will attest that children spend a significant portion of their first years of life learning how (acquiring know how) to perform "simple" tasks like raising their arms and walking. (Thalen 1995) These abilities are, of course, so widely mastered that we often forget that a difficult cognitive effort was required for their acquisition. Nevertheless, abilities like walking are cases of cognition and represent a significant cognitive accomplishment, one which is difficult to replicate in artificial intelligence (or for that matter when information from either of the vestibule's two sensory areas, the utricle and the saccule are cut off or altered). In fact, I was often asked if my nine month old child "knows how to walk yet." So, it is false that such tasks are simple in any way that would preclude the necessity of explanation by appeal to knowledge how. It is likewise false that using knowledge how to describe the sorts of tasks

classified as simple by Carr transgresses good sense. For these reasons I think that any insistence by Carr that one evaluate his theory only for what *he* classifies as non-simple acts is unmotivated and renders Carr's theory irrelevant to much of the psychological explanation he seeks to understand.

To summarize, visual detection of object edges seems a *prima facie* strong candidate for knowledge how despite Carr's (and probably many philosopher's) assertion that such ascriptions of knowledge how "transgress good sense," i.e., seem radically at odds with one's ordinary understanding and use of the phrase "know how." Knowledge how ascriptions are, *de facto*, ordinary ascriptions within vision science. Moreover, one cannot simply say, for instance, that we "can" or have the "ability" to detect edges. Ascriptions of knowledge how with regard to edge detection are neither simple nor are they bereft of epistemic evaluation. Rather, such ascriptions constitute the ascription of reliable mechanisms which serve to make complex inferences in response to environmental circumstances. These inference mechanisms are used to explain the difference between a consistently reliable ability to detect object edges–indeed to see at all--and an inability to perform and/or an unreliable ability.

Returning to my example, suppose that basic outlines of contemporary theories of early vision, for instance Marr's (1981), are approximately correct in supposing that humans are able to detect object boundaries by computing edges, depth information, etc. from retinal arrays. Such explanations of vision in cognitive science attribute to humans the knowledge how to detect object boundaries. Nevertheless, detecting object boundaries by computing edges and depth information from retinal stimulation is not something that most humans consciously entertain as a purpose. One can, in fact, exercise *only* limited indirect conscious control over

their object boundary detection.[7] Moreover, until one studies theories of vision, one lacks the theoretical concepts necessary to form an intention to detect object boundaries by computing edges and depth information. Must one learn theories of early vision in order to know how to detect object boundaries? Carr and others like him seem forced to answer affirmatively, but such an answer seems implausible, since learning such theories does not effect performance.

One objection to the above example runs as follows: This example commits the fallacy of composition. The fact that a subcomponent of a system has a capacity hardly entails that the system as a whole has that capacity. The mere fact that some part of their brain knows how to detect edges does not imply that the ***person*** has such know how. This point is usually followed by some analogy: The fact that the treasurer knew that the college was about to go bankrupt does not entail that the college knew it, UNLESS he reported THAT information upstairs. I find this objection ill-conceived. First, and most glaringly, if we as persons did not know how to detect edges the world would appear to us as an undifferentiated two-dimensional surface--making it hard, indeed, to do things like grab coffee cups, walk through doors, or any other mundane object-oriented activity. In short, edge-detection is constitutive of conscious visual ability in ways that the beliefs of treasurers are not constitutive of the financial awareness of an institution. Second, by describing the fallacy as one of composition one implies that subcomponents of a system are appropriately described as possessing knowledge how, which grants my conclusion in an effort to criticize my argument. To deny the relevance of cases like edge detection, the objector must deny that *any* component of the human nervous system is appropriately described as knowing how to detect edges unless it is also appropriately described as consciously intending when exercising that know how. Such a denial not only

runs strongly contrary to current, well-supported theories in cognitive science, it violates the general, shared framework for understanding ascriptions of knowledge how as explaining the difference between a consistently reliable ability to perform some cognitive task and an inability to perform and/or an unreliable ability.

Even should one reject the notion that there is a meaningful sense in which we know how to detect edges, the fact remains that one can garner a wide variety of well-documented examples to support the claim that the possession and the exercise of knowledge how does not require conscious intent. To summarize, a great deal of *prima facie* know how involving common sense as well as expert knowledge utilized such diverse cognitive tasks as complex as well as simple problem solving, evaluations, judgements, motor skills, and perception are not plausibly understood as resulting from conscious intent to exercise knowledge in the form of inferences involving practical directives for the following reasons: **(1) Cognizers often lack any conscious intent to perform the task. (2) Cognizers often perform the task without any conscious experiences of remembering that they know how to perform the task, perceiving that they are performing the task, identifying the task they are performing, or knowing that they can perform the task. In fact, as with H.M. they may be physically unable to exhibit these properties. (3) Cognizers often remain unaware of their responses, and hence exhibit no conscious intent to perform the task. (4) Cognizers often cannot articulate the beliefs which would constitute the knowledge (practical directives) which they supposedly are consciously intending to employ in performing the task. (5) Cognizers often exhibit no conscious awareness of environmental factors or information shown by analysis to have importantly influenced their responses, and hence no conscious intent to**

utilize that information. (6) Cognizers are often unaware of any connection in their thought between such environmental/problem factors and their responses, and hence exhibit no conscious intent to employ practical directives regarding such connections. (7) Cognizers often offer demonstrably false or highly implausible accounts of the reasoning underlying their performance, and hence cannot be supposed to conscious intend to employ the practical directives regarding the task that they likely employed.

What about Carr's view that the person must be acquainted with a set of practical instructions for performing ϕ? It is clear that Carr intends this clause to be interpreted as asserting that the person explicitly represents a set of practical instructions to themselves, hence Carr's characterization of knowledge how in terms of reasoning from premises, and his comparison between knowledge how and knowledge that in the following quote:

> The main thrust of our argument has been that knowing how involves not theoretical reasoning but *practical reasoning*, which we agree with a number of recent writers, differs from its theoretical counterpart in several crucial respects--those pertaining to its validating principle, its rules of inference and to the status of its conclusions. Inter alia, knowing how in the strong sense to play football is knowing the rules of the game, but a statement of the rules of the game is not a theoretical statement but a description of a set of constitutive rules of a *practice*, and mastery of the rules brings with it an understanding of the activity rather than a theory. (1981, p.60)

Carr has erred here as well. The data noted by Nisbett and Wilson and Schacter's data as well as data regarding many cases of animal know how, like the case of the Caddis Fly, speak against Carr. Specifically, **(1) Cognizers often cannot articulate the practical directives which would constitute the knowledge with which they supposedly are familiar. (2) Cognizers often exhibit no conscious awareness of environmental factors or information shown by analysis to have importantly influenced their responses, and hence exhibit no basis for supposing that they understand that information relevant to their performance is,**

in fact, relevant. (3) Cognizers are often unaware of any connection in their thought between such environmental/problem factors and their responses, and hence exhibit no basis for supposing that they possess or employ practical directives regarding such connections. (4) Cognizers often indicate that their reasoning was based on practical directives that they likely do not employ in performing the task. Likewise, in the case of early visual processing, no one expects that human brains actually represent an algorithm for computing object boundaries from retinal stimulations. Instead, cognitive scientists suppose that brains instantiate such algorithms. That is, the processing centers in the brain behave in a way which the algorithm would dictate. Remember that there are two ways to compute; a system can instantiate a simple algorithm and a complex set of instructions, or a system can instantiate a complex algorithm. A Turing machine that computes the successor function instantiates a somewhat complex algorithm for adding one to inputs. Its state transitions are disciplined by the physics so that they obey the set of instructions as given in the flow chart (not given in the machine). On one's computer, however, one can actually write a set of instructions (program) which the computer represents and follows to compute the successor function.

Likewise, persons who regularly, successfully ride their bicycles know how to ride bicycles, and one would not dispute that ascription simply because they do not or cannot express their ability in terms of a set of condition-action rules. One need not have an explicit understanding of the dynamics of bicycle riding to know how to ride a bicycle. Rather, one's actions need to be coordinated in a way disciplined by the dynamics of bicycle riding. Indeed, knowledge of how to perform physical skills often literally moves from

areas of the brain associated with conscious executive function to areas of the brain strongly disassociated with conscious executive functioning. (Pascual-Leone, A, Grafman, J. and Hallett, M, 1994) As noted earlier, some connectionist cognitive scientists suggest that their models invite just such a nonrepresentational understanding of cognitive functioning when they claim that "...almost all knowledge is *implicit* in the structure of the device that carries out the task rather than *explicit* in the states of the units themselves." (Rumelhart 1989, p.136)

5.) NEO-REDUCTIONIST VIEWS

Before continuing it is worthwhile noting that the criticisms offered above prove equally telling against the currently popular neo-reductionist view of Jason Stanley and Timothy Williamson. (Stanley and Williamson 2001) S&W have two arguments for their view. The first argument consists in denying that there exists a significant syntactic differences between sentences ascribing knowledge how and sentences ascribing propositional knowledge (knowledge that). The second argument utilizes Lauri Karttunen's theory regarding the semantics of the two respective classes of attributions to provide a propositional interpretation of statements attributing knowledge how to individuals. I ignore both the details of the development of S&W's semantic and syntactic framework and their claims that it can easily be adapted to other semantic and syntactic frameworks to yield the same results. S&W's penultimate conclusions, then, are expressed in the following set of three interlocking assertions:

> Then (29) {Hannah$_i$ knows [how PRO$_i$ to ride a bicycle].} is true relative to a context *c* if and only if there is some contextually relevant way *w* such that Hannah stands in the knowledge-that relation to the Russellian proposition that *w* is a way for Hannah to ride

> a bicycle, and Hannah entertains this proposition under a practical mode of presentation. (Stanley and Williamson 2001, p.430)
>
> ...employments of knowledge-that are often unaccompanied by distinct acts of contemplating propositions. (Stanley and Williamson 2001, p.415)
>
> ...thinking of a way under a practical mode of presentation undoubtedly entails the possession of certain complex dispositions. (Stanley and Williamson 2001, p.429)

Simply put, then, S&W's analysis of sentences ascribing to an individual the knowledge of how to X requires; (1) that the person to whom the knowledge how is attributed has a belief, (2) the employment of the belief quite possibly never involves the belief being consciously occurrent, (3) the content of the belief is that certain of the person's own complex dispositions result in successful Xing as Xing is described in the ascription sentence.

Possibly the most striking fact regarding S&W's position is that it is not, as they describe it, a view in which "knowledge-how is simply a species of knowledge-that." (Stanley and Williamson 2001, p. 411) On their view knowledge how requires at least one belief--at least one element of propositional knowledge. So, propositional knowledge is necessary for knowledge how. However, as the above tripartite assertions indicate, propositional knowledge is not sufficient for knowledge how on S&W's account. One still needs to possess a complex set of dispositions, the execution of which in the appropriate context reliably results in Xing. This last point, has been perceptively noticed by Koethe (2002), though I press the point harder here that Koethe did his brief but insightful critical comment. I call Stanley and Williamson's view neo-reductionist because on their view, unlike Carr or other reductionists since Ryle, knowledge how is not reduced to propositional or quasi-propositional knowledge in the form of beliefs or directives. Instead, S&W require only that knowledge how involves propositional

knowledge, a weaker claim. In criticizing S&W's view, I do not argue that knowledge how never involves propositional knowledge. I merely criticize S&W's claim that ALL knowledge how requires a least this one bit of propositional knowledge.

S&W themselves acknowledge that such second-order beliefs about what I would term the epistemic status (or neutrally, the reliability of) complex sets of dispositions are not sufficient for knowledge how. (Stanley and Williamson 2001, p.416) But, they do assert the necessity of such beliefs. Koethe notes the former point, but does not go on to challenge the latter. There are, in fact, three reasons to suppose that Stanley and Williamson's account is clearly and demonstrably false. First, individuals regularly utilize what S&W would classify as contextually reliable complexes of dispositions to successfully perform tasks without the sorts of beliefs S&W hold to be necessary. Indeed, they will often explicitly deny that they have such beliefs. (For example, Bargh and Chartrand 1999, Cermak *et al.* 1973, Cohen and Squire 1980, Cole and Rotella 2002, Corkin 1968, Jacoby and Witherspoon 1982, Nisbett and Wilson 1977, Schacter 1980, Schmand *et al.*1992, Schröder *et al.* 1996, van Gorp *et al.* 1999, Vicari 2001) Thus, Stanley and Williamson's claim that such beliefs are necessary for knowledge how is clearly empirically false. As noted above: **(1) Cognizers often perform the task without and conscious experiences of remembering that they know how to perform the task, perceiving that they are performing the task, identifying the task they are performing, or knowing that they can perform the task. (2) Cognizers often cannot articulate the information or indicate the complexes of dispositions which would constitute the content of their supposed beliefs when explicitly questioned about such information and dispositions. (3) Cognizers often exhibit no conscious awareness of environmental factors or information**

shown by analysis to have importantly influenced their responses, and hence exhibit no basis for supposing that they believe that the information is relevant to their performance. (4) Cognizers often offer demonstrably false or highly implausible accounts of the reasoning underlying their performance inconsistent with any beliefs regarding how they actually perform the task, and hence cannot be supposed to have true beliefs regarding how they, in fact, perform a task.

The second reason to suppose that S&W's account is clearly and demonstrably false consists in the very nature of the beliefs S&W require for knowledge how. While some knowledge how plausibly involves complexes of dispositions which one might be supposed to grasp as a single belief, knowledge how often involves complexes of dispositions so elaborate and so diverse in their temporal and environmental contexts as to make the notion that individuals have single beliefs articulating these "ways," Russellian or otherwise, implausible in the extreme. For example, knowing how to repair cars at an expert-level, knowing how to diagnose cancer from pathological slides, knowing how to make good wine from grape to bottle, etc.. All these instances of knowledge how stretch across time periods far exceeding any practical period for inner or outer ostentation to the relevant complexes of dispositions. These complexes of dispositions are, likewise, so elaborate that their propositional expression requires multiple volumes of hundreds of pages. Individuals who suppose the sorts of beliefs to which Stanley and Williamson appeal are ubiquitous, acquired in many cases without explicit instruction, and often exercised unaccompanied by distinct acts of contemplating propositions are encouraged to attempt such comparatively simple feats as never again writing down directions or forming a single conjunctive belief that for each and every person listed on

a page of a phone book, that conjunction of people lives in their city.

The third reason to suppose that S&W's account is clearly and demonstrably false consists in the fact that the brain areas operant in the elicitation and generation of such contextually reliable complexes of dispositions are strongly dissociable from areas of the brain responsible for propositional knowledge. (See for instance, Doya 2000, Mogenson, Jones, and Yim 1980, Pascual-Leone, Grafman, and Hallett, 1994, Schacter 1989, Smeets, Marin, and Gonzalez 2000.) To summarize these three points, S&W's hypothesized beliefs are not necessary for knowledge how because: (1) Knowledge how occurs in their absence–even when the person's beliefs actually explicitly contradict S&W's hypothesized beliefs. (2) S&W's beliefs are implausible posits on their own merit in many cases of knowledge how. (3) Neuroscience indicates that S&W's beliefs are not causally operant in many of the behaviors for which knowledge how is an explanation. These points are very significant in my mind since, as stated earlier, the claim that knowledge how reduces to propositional knowledge has always, and rightly, been understood as requiring that: **(1) The knowledge itself be propositional knowledge or at least encoded explicitly as particular linguistic or quasi-linguistic expressions that were (2) causally operant in the manifestation of the knowledge.** These are the two fundamental claims denied by Ryle (1949) and reasserted by the likes of D.M. Armstrong (1968), D.G. Brown (1970), David Carr (1979 and 1981), P.T. Geach (1957 and 1966), Jaakko Hintikka (1974), A.J. Kenny (1966), Charlotte Katzoff (1984), and J. Mackie (1974). These are likewise the claims of relevance to epistemologists and to such programs in the philosophy of mind as belief-desire psychology.

I conclude that neither the reductionist strategy nor the neo reductionist strategy prove

viable general accounts of knowledge how. One cannot completely reduce all knowledge how to consciously accessible, explicitly represented propositional or quasi-propositional knowledge. Nor can one require of all knowledge how that it involve at least one piece of propositional knowledge that might never actually be conscious.

6.) CARR ON MANIFESTING AN ABILITY

Before offering my own theory of knowledge how, I must note one more difficulty with Carr's account. Though Carr rejected Ryle's dispositional analysis, he did follow Ryle's thinking in that he linked knowledge how to actual ability. I noted (footnote 2) that Carr must qualify his necessary conditions to avoid inconsistency in that his (CD3) requires current ability, while in other places Carr explicitly allows past ability to count towards knowledge how. In defending his intuition that past ability supports ascriptions of know how, Carr asserts that past performance together with a grasp of practical procedures can underwrite knowledge how. He distinguishes a strong from a weak sense of know how. Strong know how requires current ability, whereas weak know how requires only that the person have had the ability at some time. Though Carr's realization that knowledge how is intimately linked to action seems sound, Carr's suggestion for understanding knowledge how ascriptions in the absence of ability is far from adequate. Carr's distinction between senses of knowledge how is unmotivated by factors other than our noticing that we sometimes ascribe knowledge how in the absence of ability. Moreover, Carr's weak sense of know how applies to cases where ascriptions of knowledge how seem inappropriate. Consider the case of John, a famous, but elderly logician who first proved an important result, P1. John's aging has brought with it a decline in his mental abilities such that he still knows the language and inference rules of logic, but can no

longer reconstruct the proof of P1. Clearly, John has had the ability to prove P1 in the past. He is also acquainted with the practical procedures necessary for proving P1, and can entertain proving P1 as a purpose. But, it seems obvious that we would say that, sadly, John no longer knows how to prove P1. John no longer knows how to prove P1, since he can no longer construct the proof using the rules of inference. That is, his native cognitive abilities are now deficient to utilize the basic rules of logic to construct a proof.

Forgetting the particulars of some instance of knowledge how seems to present a similar situation of unwarranted knowledge how ascription by Carr. For example, when one forgets how to get to a friend's house, one is clearly still acquainted with a set of practical procedures necessary for getting to the house. One still knows how to go left, go right, go straight for four blocks, read street signs, etc.. Moreover, one has had success in the past. Carr, therefore, seems committed to an ascription of knowledge how even though one clearly no longer knows how to get to the house.

So, past ability together with acquaintance with practical procedures does not always support knowledge how. Moreover, though Carr does not explicitly consider such cases, knowledge how ascriptions also occur in cases where a subject has always lacked an ability or possesses an unsystematic ability. Consider the case of Bill, who consistently bakes perfect cakes. Bill uses normal baking methods, which he learned in school. We would want to say that Bill knows how to bake a cake. We would further not deny Bill his cake baking know how simply because he does not know how to bake cakes in a microwave oven. Here Bill has knowledge how even though he has always lacked ability to bake a cake in certain situations.

A theory of knowledge how that ties knowledge how to ability must therefore provide a

mechanism to explain and understand ascriptions of knowledge how in the absence of ability or in cases of unsystematic ability. Such a mechanism ought not be a mere *ad hoc* addition to a theory, but follow directly from the theory.

7.) Theory

With the above in mind, I define *knowledge how* as follows:

> Ψ *knows how* to Φ iff
> (1) Ψ is a cognizer,
> (2) Φ is a cognitive task,
> (3) Ψ instantiates and exploits a tractable set of processes tokening outputs that reliably satisfy the satisfaction conditions associated with the idealized target function characteristic of the Φ across the task domain, and
> (4) the process reliability in those cases where the input given the system satisfies the satisfaction conditions associated with the target function is greater than or equal to the general level of reliability of the process across the task domain.

In the above definition, the notion of a system is considered a primitive. The notion of a cognitive task is the notion of an inferentially characterizable capacity. Specifically, one can understand a given Φ as a cognitive task only to the extent that one can generate a well-specified task description of Φ. Within cognitive science, one explains a system's performance in terms of a well-specified cognitive task, that is, by showing how the system's interactions (between the system and the environment as well as between the system's states) typify a rationale for using information either to generate more information or to control behavior. This point is often expressed by saying "...that cognitive capacities are *inferentially characterizable capacities*...." (Cummins 1983, p. 33) A well-specified task description (**WSTD**) includes three elements. (1) A **WSTD** includes a specification of input and output types in terms of an idealized target function that defines success conditions for performing the task by mapping input sets to sets of acceptable outputs. For instance, planning involves having a goal state and

initial state descriptions as input and the planner computes plans as output. A system computing the successor function has representations of numbers for which one wants the successor as input and representations of successors as output. The dancer has requests for dances, etc., as inputs and sequences of movements as outputs. (2) A **WSTD** states the nomic correlations (including statistical correlations) that underlie the system's performance of a task. For instance, David Marr discusses what he calls the "underlying physical assumptions" of human vision. Marr's assumptions amount to nomic generalizations about the task domain, and include the existence of surfaces, spatial continuity, and so on. (Marr 1981, pp. 44-51) The dancer's performance likewise relies upon underlying assumptions including such things as the resistance and stability of the floor, the spatial medium through which movement occurs, leg and arm mass, muscle stretch and spring dynamics, etc.. (See for example, Thalen 1995) (3) Finally, a **WSTD** includes a specification of the relevant process(es) by reference to the system's laws of operation within the domain, viewing these operations as a strategy or set of strategies for generating outputs from inputs and/or generating new strategies relying upon certain nomic correlations.

There are two important things to note regarding the above definition. First, knowledge how is a counterfactual supporting attribution that continues to apply even when the abilities, knowledge and environmental conditions for the exercise are not currently actual. For example, people would normally continue to say that Alan Alborn (US Olympic Ski Jump Team member) knows how to ski jump even during the summer when conditions are such that Alborn cannot exercise his ski jumping know how (no snow). Alborn still knows how to ski jump in snowless climes and times because the world is such that he can regularly inhabit

climes and times when he can exercise his knowledge of how to ski jump. Likewise, people would normally continue to say that Alborn knows how to ski jump even during his recovery from knee surgery. Though Alborn might not be physically able to ski jump during his recovery period, his ski jumping motor skills and knowledge will not degenerate too dramatically during his recovery period and he will regain his ability to ski jump during the normal course of events. The counterfactual elements of knowledge how ascriptions, though false at times or in some circumstances, are relevant and appropriate precisely because they are quite likely to become actual for the individual as a part of the range of their normal circumstances, abilities, and knowledge. Thus, one can account for cases like Alborn and our cake baker Bill if one allows that ascriptions of know how are task-specific. That is, ascriptions of know how presuppose certain underlying conditions which are part of the task one is thought to know how to perform. These conditions are often implicit in one's thinking about and ascriptions of knowledge how, and can lead one both to attribute knowledge how to individuals even when conditions exist in which they would lack ability, and to attribute knowledge how to individuals in a counterfactual sense. Bill knows how to bake a cake, because implicit in one's ascription is the presupposition that Bill can use a conventional oven in performing the cake-baking task. Similarly, when we say that someone knows how to ski, we presume implicitly that this task presupposes access to skis, snow on the ground, etc..

The second important thing to note about the above definition of knowledge how is really an extension of the first point. People often make counterfactual ascriptions of knowledge how where the presupposed counterfactual connections between abilities, circumstances, and knowledge may never become actual. Such counterfactual ascriptions,

though strictly speaking false, are necessary and important both because an individual's knowledge how often extends beyond their current range of normal circumstances, abilities, and knowledge and because individuals can often easily augment or extend their knowledge how by adaptive extensions of their current connections between ability, knowledge, and circumstance. Part of understanding and predicting one's own behaviors and the behavior of others is anticipating that ways in which their behavior might be altered by adaptions in the form of novel or renewed connections between ability, knowledge, and circumstances. The key to understanding such ascriptions is understanding that they are, in fact, counterfactual ascriptions. Alborn would no longer know how to ski jump were he permanently incapacitated anymore than Christopher Reeves continued to know how jump horses once he became a paraplegic and his motor skills and knowledge degenerated through normal brain reorganization. Thus, such false ascriptions are relevant and appropriate speculations to the degree to which, the counterfactual elements of the ascriptions are likely to become actual. An important way in which counterfactual elements of know how ascriptions can become actual is through an individual's creating connections between ability, knowledge, and circumstances in others. I'll discuss this point in more detail in the next section.

Finally, in allowing that ascriptions of know how are task-specific, one does not thereby offer a dispositional analysis of knowledge how on the lines of Ryle. A dispositional analysis holds that one is disposed to behave in a certain way in the specified conditions. Task-specificity, in contrast, states only that certain conditions underlie the manifestation of knowledge how, regardless of whether they are sufficient to elicit its manifestation. So, task-specificity asserts that Bill requires access to conventional oven in order to manifest his cake-

baking know how. It does not claim that access to a conventional oven will elicit Bill's cake-baking behavior.

One can illustrate the relationship between closeness of possibility and relevance of knowledge how ascription as follows:

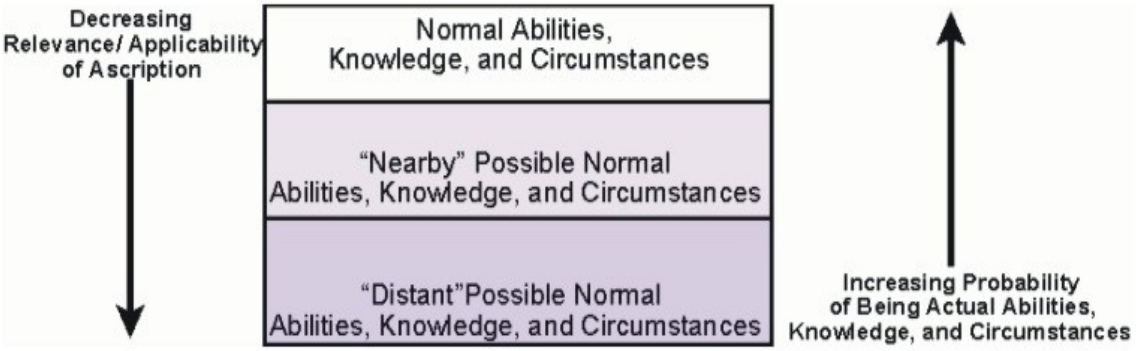

8.) Ordinary Attributions, Vagary, and Conflicting Intuitions

Ordinary attributions of knowledge how are, needless to say, far from achieving the sort of rigor one might expect in a fully developed explanation in Cognitive Science. Nevertheless, I claim that ordinary ascriptions of knowledge how have the same general structure of positing reliable processes operating within a context. Ordinary ascriptions are merely much more vague in their understanding of both processes and context. While the vagary of ordinary ascriptions allows people ascribe knowledge how in, from a scientific perspective, relative ignorance, it likewise muddies intuitions regarding the aptness of such descriptions in important and predictable ways. Indeed, intuitions regarding various thought experiments discussed in this and other papers are sometimes diverse. One of the more profound vagaries

that results from ordinary ascriptions are those that stem from the applicability of counterfactual ascriptions of knowledge how. As I note, people often make ascriptions of knowledge how based upon counterfactual abilities, knowledge, and circumstances. These ascriptions are relevant or appropriate only to the degree to which, the counterfactual elements of the ascriptions are likely to become actual for the individual. Since in ordinary attributions circumstances and abilities are more often than not quite vaguely understood, the accuracy, relevance and appropriateness of such attributions can be muddied so as to generate conflicting intuitions.

I'll illustrate the above point by considering an alternative construal of the dancer example which occasionally arises and which one might suppose undermines my analysis: One could claim that the dancer does know how to dance a semaphore rendition of Grey's "Elegy," he merely does not *know that* he *knows how*. I have two objections to this treatment. First, ascriptions of knowledge how are made in explaining the difference between an inability to perform and/or an unreliable ability to perform some cognitive task in appropriate circumstances and a consistently reliable ability in such circumstances. In other words, lacking an ability to reliably perform some cognitive task, one has no motivation to ascribe knowledge how. Cognizer's are thought to know how to do something in that they tend to perform correctly in circumstances where it is plausible to suppose that they could get it wrong. The above treatment of the case, however, prescribes just such an unmotivated ascription. Second, the missing knowledge (that he knows how to perform a semaphore rendition of Grey's "Elegy") does not explain the dancer's failure. Imagine that an internationally acclaimed semaphore expert, introduced to the dancer as such, tells the dancer that one of his dances is a

semaphore rendition of Grey's "Elegy." Unfortunately, the conversation ends before the expert reveals which dance is the semaphore rendition. It seems that on the above account the dancer then knows that he knows how to perform a semaphore rendition of Grey's "Elegy," yet still fails to perform the dance so described. One might claim that the required knowledge is a knowledge that "Improvisation #15" is a semaphore rendition of Grey's "Elegy." However, I would suggest that this merely grants my point, as this knowledge would provide a functional connection between ability and the dance so described.[8]

Such intuitions regarding the dancer case are relatively uncommon, but perfectly explicable in a manner which also undermines their veracity as counterexamples to my theory. Presumably the dancer is thought to know how to dance a semaphore rendition of Grey's "Elegy" because he knows a sequence of movements that correspond to a semaphore rendition of Grey "Elegy". However, the correspondence between movements known under a different set of circumstances and knowledge, and those movements under an envisioned set of circumstances and knowledge remains, nevertheless, counterfactual. The closeness of the counterfactual situation to the actual situation together with the vagueness of ordinary ascriptions conspire to create a false intuition regarding the dancer's actual know how. In similar fashion, the intuition is undermined when one considers questions designed to increase the perceived distance between the actual and counterfactual cases:

1.) Suppose that the conventions of semaphore developed differently in the actual world, so that the dancer's movements are not a semaphore rendition of Grey's "Elegy" in **actual semaphore**. Would the dancer still know how to perform a semaphore recital of Grey's "Elegy" according to the now counterfactual semaphore conventions?

2.) Would the dancer know how to perform a semaphore recital of Grey's "Elegy" had

semaphore **never been discovered and utilized anywhere on Earth**?

3.) Does the dancer know how to do everything one can imagine his actions being utilized to accomplish? For instance, suppose that aliens in another galaxy created a motion sensitive chamber that could be used to translate the dancer's movements into steering commands. The dancer's motions could be translated into a series of commands that pilot a ship (created by different aliens) through an asteroid field located in still another galaxy. Does the dancer know how to pilot the alien ship through that asteroid field?

In the above series of questions, each question moves the counterfactual connection between the dancer's potential abilities, knowledge, and circumstances farther away from those had by the dancer in the actual world. Everyone to whom I have put these questions reports that the ascriptions of knowledge how seem decreasing plausible with the increasing distance between the actual and counterfactual situations. For instance, one wouldn't want to say that I know how to read this: lwkl qnwkhuhi ruhldp. Even though there is a mapping from message letters to alphabet letters that decodes it. One can thus adapt the above diagram to capture the relationship between strength of intuition and perceived distance from normal ability, knowledge, and circumstance as follows:

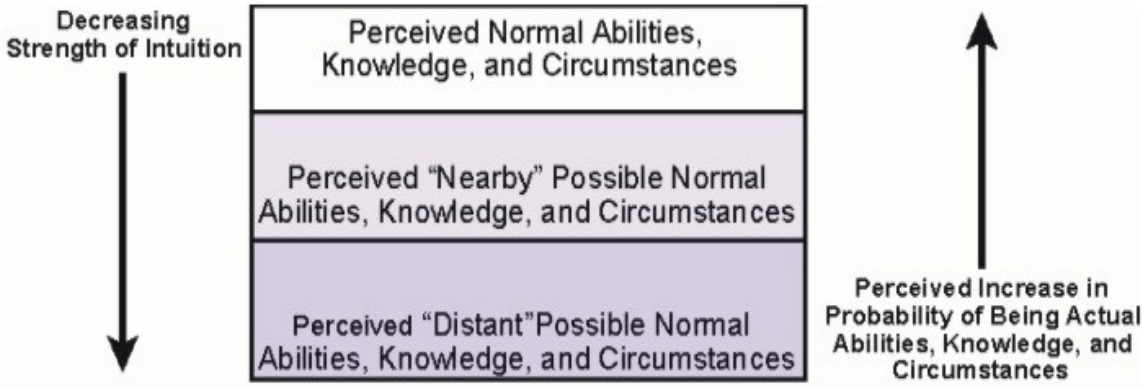

The suggestions made here with regard to intuitions and attributions of knowledge how, though not systematically empirically validated are consistent with the literature on the availability and simulation heuristics in that imagination as well as rule-based hypotheses influence (correctly and incorrectly) the strength of knowledge how attributions. (Tversky & Kahneman, 1973, 1974, and 1982) Likewise, the suggestions are consistent with the "dual-representation" approach to mind reading widely acknowledged in the recent literature. (Stich & Nichols 1997; Goldman, 2002; Saxe, 2005)

To summarize with respect to the dancer case, I claim that knowing how requires both ability and the capacity to exploit an ability given the task-specific context. That is, knowledge how requires a functional connection between one's abilities, knowledge, and the circumstances under which it is appropriate to exercise that ability. What prevents the dancer from performing the semaphore Gray's "Elegy" is that he lacks a conceptual understanding linking the ability to the request as described. In other words, the conceptual context, he associates with the movements--is the one defined by Improvisation #15, not a semaphore Gray's "Elegy." By claiming that he needs to know that he knows how to perform the semaphore as otherwise described, one merely asserts that the dancer needs a functional connection between the movements as he understands them and the conceptual context of the movements as a semaphore rendition. I suspect that those who have contrary intuitions regarding the dancer case are motivated by the perception of closeness of the counterfactual world in which understanding semaphore provides a mechanism integrating the dancer's ability, knowledge, and context. Indeed, the intuition seems to be dampened when one asks people questions which make the counterfactual abilities more distant from the dancer's actual

abilities, knowledge, and circumstances.

An important class of counterfactual ascriptions of knowledge how are those based upon the possession of linguistic information about how to perform tasks. Language represents a unique and important means by which humans acquire and transmit knowledge how. Our linguistic abilities greatly facilitate the acquisition of other abilities and hence knowledge how. In many cases linguistic information can immediately or almost immediately result in knowledge how. For example, all one usually needs to know how to get to the university is good directions. Because linguistic information can often translate with relative ease into knowledge how, people often overemphasize the significance of the possession of linguistic knowledge of how to perform a task and suppose that possession of linguistic information is sufficient for knowledge how. For instance, Stanley and Williamson suggest that "... a ski instructor may know how to perform a certain complex stunt, without being able to perform it herself." (Stanley and Williamson 2001, p.416) Interestingly, this claim seems unlikely to be true on S&W's account. Recall that in order for the ski instructor to know how to perform the stunt for S&W, the instructor must stand in the knowledge-that relation to the Russellian proposition that *w* is a contextually relevant way for the instructor to perform the stunt, and the instructor entertains this proposition under a practical mode of presentation, where thinking of a way under a practical mode of presentation entails the possession of certain complex dispositions. However, if the instructor cannot perform the stunt, it seems dubious to suppose that she possesses both the belief and the complex dispositions which she entertains under a practical mode of presentation. Hence, the instructor does not know how to perform the stunt after all on S&W's account. If, on the other hand, one grants S&W that she does

have the belief and the complex dispositions entertained under the practical mode of presentation while lacking the ability to perform the stunt, then it seems as though the instructor's belief that *w* is a contextually relevant way for her to perform the stunt is actually false. Hence, the instructor does not know how to perform the stunt after all on S&W's account.

Another way in which linguistic information creates the illusion of actual knowledge how results from the fact that even in cases where the person can never have the required abilities, they can often transmit the linguistic information to others so as to create knowledge how in these people. Thus, the expression, "those who can't do, teach." That is, an important class of counterfactual knowledge how ascriptions are ascriptions based upon an individual's ability to teach another person how to perform a task. Thus, individuals may continue to ascribe knowledge how to Reeves or Alborn despite even permanent incapacities in so far as they can use their residual abilities and knowledge to help others to gain a knowledge of how to ski or horse jump. As with other types of counterfactual ascriptions, the degree to which such ascriptions are relevant corresponds to the potential that such linguistic information has to create a reliable ability in the individual or someone to whom the individual can transmit the information. Thus, one can again create decreasing levels of intuitive confidence in knowledge how ascriptions by explicitly manipulating the counterfactual distance between linguistic information and actual ability:

1.) A logic professor can recount how to prove Gödel's incompleteness theorem.

2.) A logical novice can recount how to prove Gödel's incompleteness theorem.

3.) A theater student having no familiarity with or knowledge of formal logic memorizes how

to prove Gödel's incompleteness theorem as lines for a play.

4.) A theater student having no familiarity with or knowledge of formal logic memorizes phonetically how to prove Gödel's incompleteness theorem in Spanish as lines for a play.

Were it the case, mere possession of linguistic information were sufficient for knowledge how, one would not expect to see the decrease in intuitive confidence that the person knows how to prove Gödel's incompleteness theorem as one progresses through the cases (i.e., as linguistic information is separated from actual ability).

So, overestimating the importance of linguistic information to knowledge how often leads to intuitions that people know how to perform a task in ordinary usage. Furthermore, people often overestimate the amount of linguistic information possessed by individuals who do know how to perform tasks. Thus, people often suppose that individuals have a general knowledge of how to perform a task based upon past–even current ability. For example, one might suppose that individuals who are fluent readers of their native language and by any standard quite literate would have a great deal of propositional knowledge related to reading the English language in print. However, none of my literate colleagues who were cornered in various hallways throughout the university knew the number of phonemes in the English language. Nor could they list a significant number of the estimated 40–45. Likewise, though it has long been known that, for example, readers exhibit associative priming effects facilitating both lexical decision and naming tasks, no one supposes that individuals can state propositionally or exhaustively those associations. (Perea and Gotor, 1997, Perea and Rosa, 2002) Indeed, though literate individuals posses a great deal of know how when it comes to reading English, linguists have long known that much of it is not necessarily readily

translatable to propositional knowledge or to other modalities. (Reber and Lewis, 1977) This is likely one reason why a disturbing large an increasing number of literate people who become blind do not learn to read in braille. (Wells-Jensen, S., Wells-Jensen, J., and Belknap, G., 2005). It is, in fact, somewhat difficult to learn to read in braille, even when you were a good or excellent reader before losing your vision.

In a similar fashion, people occasionally claim that the permanently blind woman mentioned earlier knows how to read, she merely cannot manifest that know how. Indeed, Stanley and Williamson claim that this sort of example clearly indicates that knowledge how does not require ability. S&W's claim here is ironic for two reasons. First, looks false according to their own theory. Specifically, the blind woman will either continue to "...entertain the Russellian proposition that *w* is a way for her to read under a practical mode of presentation," or she will stop. If she stops believing the proposition, then she no longer knows how to read on S&W's account. If she continues to believe the proposition, then, since she no longer has the relevant abilities, the belief is false, and the blind woman no longer "...stands in the knowledge-that relation to the Russellian proposition..." as required by S&W's view. (Stanley and Williamson 2001, p. 430) The second ironic aspect of S&W's assertion is that, if one asks people why they suppose that the blind woman can read, they invariably link the ascription to a counterfactual ability. That is, they say, "she could read, if she had the ability to see." Thus, rather that separating ability and knowledge how, the example seems to link the two. Like the other intuitions discussed in this section, this objection is based upon a confusion between a counterfactual and an actual sense in which one can legitimately be said to know how to do something. There is a sense, allowed by my position, in which the blind

woman can be said to know how to read; she would know how to read if she could see. This is a counterfactual claim about what the blind woman would know how to do if she had the ability to see. However, the objector falsely supposes that this counterfactual sense of knowing how warrants ascriptions of actual know how. The blind woman cannot read, and does not know how to read given her actual physical abilities. While such an ascription is relevant when applied to someone who has temporarily lost their sight, the ascription of knowledge how is both false and uninteresting when applied to a permanently blind person. In fact, the permanently blind woman no more knows how to read based upon her counterfactual abilities than I know how to get around Peking on the basis of the counterfactual know how I would have if Peking had exactly the same lay out as Vancouver. Just as getting around Peking would be a different problem if Peking were laid out exactly like Vancouver, reading when one is blind is a different problem than reading when one has vision. In my view, then, know how coordinates the interaction of various abilities and knowledge to facilitate consistent and reliable performance of a cognitive task. Hence, knowledge how presupposes that the abilities are actual in the context in which the person knows how, and that that context falls within the normal range of abilities, knowledge, and circumstances for the individual. Counterfactual ascriptions of know how, though understandable, are relevant, intuitive, useful, appropriate, etc. to our understanding or ascriptions of actual knowledge how roughly in proportion to the likelihood of their having been, their being, potentially being the actual abilities, knowledge, and circumstances of the subject. Thus, it seems silly for a blind person to refuse to learn braille because, they say "I know how to read."

The case of the blind woman does bring out two important factors in people's use of

counterfactual ascriptions of knowledge how. (1) People often suppose individuals possess more propositional knowledge than those individuals actually have. (2) Possession of propositional knowledge is often thought to warrant knowledge how in the absence of ability. However, many of these cases are actually false when understood as literally true and not as interesting, relevant, etc.. counterfactual ascriptions.

4.) CONCLUSION

To sum up, Carr's necessary conditions fail to capture his distinction between knowledge how and ability. They wrongly tie necessary conditions for knowledge how to conscious intention and an explicit knowledge of rules. Though know how can be instantiated by conscious reasoning using consciously accessible beliefs about practice, it can also be instantiated by purely physical processes which require no explicit, conscious representation of intent or of rules. A theory of knowledge how must remain neutral with regard to these possible instantiations. Finally, Carr has trouble capturing his own intuitions about various cases because he fails to see that ascriptions of knowledge how are necessarily task-specific, presupposing certain underlying conditions for the manifestation of the ability. Hence, knowledge how ascriptions have a counterfactual supporting structure, which explains certain ascriptions of knowledge how in the absence of manifested ability. The weaknesses in Carr's position are inherent in any position that redescribes knowing how in terms of practical reason, i.e., by requiring a causal connection between conscious intention to ϕ and ϕing, and ties know how to the explicit representation of general, task-independent conscious reasoning about practical matters. All such positions are challenged by the cases of tacit knowledge, unreliable introspective access to higher order cognitive processes,

somnambulists, profoundly amnesic patients, animal know how, and unconscious cognitive processes.

BIBLIOGRAPHY

Anderson, J., Kline, P, & Beasley, C. (1979). "A General Learning Theory and its Application to Schema Abstraction," in G. H. Bower (Ed.) *The Psychology of Learning and Motivation.* New York: Academic Press.

The American Psychiatric Association (1987). *Diagnostic and Statistical Manual of Mental Disorders: Third Edition-Revised*. Washington: American Psychiatric Association.

Bargh, J. and Chartrand, T. (1999). "The Unbearable Automaticity of Being," in *American Psychologist.* 54: 462-479.

Baseler, H., Morland, A., and Wandell, B.(1999). "Topographic Organization of Human Visual Areas in the Absence of Input from Primary Cortex," in *Journal of Neuroscience*. 19: 2619-2627.

Bechtel, W. and Abrahamsen, A. (1991). *Connectionism and the Mind: An Introduction to Parallel Processing in Networks*. Cambridge: Blackwell Press.

Bliss, J. and Ogborn, J. (1994). "Force and Motion from the Beginning," in *Learning & Instruction*, 4(1), pp. 7-25.

Brown, D. (1970). "Knowing How and Knowing That, What," in Oscar P. Wood and George Pitcher (eds.), *Ryle*. Garden City: Anchor Books.

Behrmann, M., Marotta, J., Gauthier, I., Tarr, M., and McKeeff, T., "Behavioral Change and Its Neural Correlates in Visual Agnosia After Expertise Training," in *Journal of Cognitive Neuroscience.* 17: 554-568.

Carr, D. (1979). "The Logic of Knowing How and Ability," in *Mind*. 88: 394-409.

Carr, D. (1981). "Knowledge in Practice," in *American Philosophical Quarterly*. 18: 53-61.

Chervinskaya, K. and Bekhterev, V. (2000). "Some Methodological Aspects of Tacit Knowledge Elicitation," in *The Journal of Experimental & Theoretical Artificial Intelligence*, 12(1), pp. 43-55.

Cohen, N. J., & Squire, L. R. (1980). "Preserved learning and retention of pattern-analyzing skill in amnesia: dissociation of knowing how and knowing that," in *Science*, 210:207-209.

Cummins, R. (1983). *The Nature of Psychological Explanation*. Cambridge: MIT Press.

de Groot, A. & Gobet, F. (1996). *Perception and Memory in Chess. Heuristics of the Professional Eye*. Assen: Van Gorcum.

Ellis, N. (2005) "At the Interface: Dynamic Interactions of Explicit and Implicit Language Knowledge," in *Studies in Second Language Acquisition*, 27(2), pp. 305-352.

Fodor, J. (1986). "Understanding Without Knowing How to Behave," in Zenon Pylyshyn and William Demopoulos ed. *Meaning and Cognitive Structure: Issues in the Computational Theory of Mind*. New Jersey: Ablex Publishing.

Ford, D., Bergen, N. and Sterman, J. (1998). "Expert Knowledge Elicitation to Improve Formal and Mental Models," in *System Dynamics Review*, Vol 14(4), pp. 309-340.

Geach, P. (1957). *Mental Acts*. New York: Humanities Press.

Geach, P. (1966). "Dr Kenny on Practical Inference," in *Analysis*. 26: 76-9.

Gilbert, D. (1989). "Thinking Lightly About Others: Automatic Components of the Social Inference Process," in J. S. Uleman & J. A. Bargh (Eds.), *Unintended Thought* (pp. 189–211). New York: Guilford Press.

Goldman, A. (2002). "Simulation Theory and Mental Concepts" in *Advances in Consciousness Research*. 45:

Gobet, F, and Simon, HA (1996). "The Roles of Recognition Processes and Look-Ahead Search in Time-Constrained Expert Problem Solving: Evidence From Grand-Master-Level Chess," in *Psychological Science.* 6: 381-381.

Gobet, F., and Simon, H. A. (1998). "Expert Chess Memory: Revisiting the Chunking Hypothesis", in *Memory*, 6:225-255.

Griffen, D. (1992). *Animal Minds*. Chicago: University of Chicago Press.

Hansell, M. (1972). "Case Building Behavior of the Caddis-fly Larvae," in *Lepidostoma Hirtum Journal of Zoology*. 167: 179-92.

Hansell, M. (1984). *Animal Architecture and Building Behavior*. London: Longman.

Haugeland, J. (1981). "Semantic Engines: An Introduction to Mind Design" in J. Haugeland (ed.) *Mind Design*. Cambridge: MIT/Bradford Books.

Hintikka, J. (1974). "Practical vs Theoretical Reason--An Ambiguous Legacy," in Stephan Korner (ed.) *Practical Reason*. New Haven: Yale University Press.

Kahneman, D., & Tversky, A. (1982). "The Simulation Heuristic," in D. Kahneman, P. Slovic, & A. Tversky (Eds.) *Judgment Under Uncertainty: Heuristics and Biases* . Cambridge: Cambridge University Press.

Katzoff, C. (1984). "Knowing How," in *The Southern Journal of Philosophy*. 22: 61-70.

Kenny, A. (1966). "Practical Inference," in *Analysis*. 26: 65-75.

Koethe, J. (2002). "Stanley and Williamson on Knowing How," in *The Journal of Philosophy*. 99: 325-328.

Pearson, D. and Laird, J. (Forthcoming). "Incremental Learning of Procedural Planning Knowledge in Challenging Environment," in *Computational Intelligence.*

Luce, G and Segal, J (1966). *Sleep.* New York: Coward McCann.

Mackie, J. (1974). "A Reply to Jaakko Hintikka," in Stephan Korner (ed.) *Practical Reason*. New Haven: Yale University Press.

Mandler, G. (1975). "Consciousness, Respectable, Useful, and Probably Necessary" in R. Solso (ed.) *Information Processing and Cognition*: *The Loyola Syposium*. Hillsdale, NJ: Earlbaum.

Marr, D. (1981). *Vision*. New York: Freeman and Company. .

Nisbett, R. and Wilson, T. (1977). "Telling More Than We Can Know: Verbal Reports on Mental Processes," in *Psychological Review*. 84:231-259.

Pascual-Leone, A., Grafman, J., and Hallett, M. (1994). "Modulation of Cortical Motor Output Maps During Development of Implicit and Explicit Learning," in *Science*. 263:1287-1289.

Perea, M. and Gotor, A. (1997). "Associative and Semantic Priming Effects Occur at Very Short Stimulus-Onset Asynchronies in Lexical Decision and Naming," in *Cognition*. 62: 223-240.

Perea, M. and Rosa, E. (2002). " The Effects of Associative and Semantic Priming in the Lexical Decision Task," in *Psychological Research*. 66: 180-194.

Quinton, A. (1970). "Ryle on Perception," in Oscar P. Wood and George Pitcher (eds.), *Ryle*. Garden City: Anchor Books.

Reber, A. and Lewis, C. (1977). "Implicit Learning: An Analysis of the Form and Strucuture of a Body of Tacit Knowledge" in *Cognition*. 5: 333-361.

Rosenbloom, P., Laird, J., McDermott, J., Newell, A., & Orciuch, E., (1985). "R1-Soar: An Experiment in Knowledge-Intensive Programming in a Problem-Solving Architecture", in *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 7: 561-569.

Rumelhart, D. (1989). "The Architecture of the Mind: A Connectionist Approach," in M. Posner (ed.) *Foundations of Cognitive Science*. Cambridge: MIT Press.

Ryle, G. (1949). *The Concept of Mind*. New York: Barnes and Noble.

Saxe, R. (2005). "Against Simulation: The Argument From Error," in *TRENDS in Cognitive Sciences.* 9: 174-179.

Schacter, D. (1989). "On the Relation Between Memory and Consciousness: Dissociable Interactions and Conscious Experience," in H. Roediger and I. Craik (eds.) *Varieties of Memory and Consciousness: Essays in Honor of Endel Tulving*. Hillsdale: Lawrence Erlbaum Associates.

Sun, R., Slusarz, P., and Terry, C. (2005). "The Interaction of the Explicit and the Implicit in Skill Learning: A Dual-Process Approach," in *Psychological Review* 112 (1) pp.159–192.

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

Sternberg, R. and Horvath, J. (1998). "Cognitive Conceptions of Expertise and Their Relations to Giftedness," in Friedman, R. and Rogers, K. (Eds) *Talent in Context: Historical and Social Perspectives on Giftedness*. pp. 177-191.

Stich S. & Nichols, S. (1997). "Cognitive Penetrability, Rationality, and Restricted Simulation," in *Mind & Language.* 12: 297–326.

Stillings, N., Feinstein, M., Garfield, J., Rissland, E., Rosenbaum, D., Weisler, S., and Baker-Ward, L. (1987). *Cognitive Science: An Introduction*. Cambridge: MIT Press.

Thalen, E. (1995). "Time-Scale Dynamics and the Development of an Embodied Cognition," in Robert Port and Timothy Van Gelder (eds.) *Mind as Motion*. Cambridge: MIT Press.

Tversky, A., & Kahneman, D. (1973). "Availability: A Heuristic for Judging Frequency and Probability," in *Cognitive Psychology.* 5: 207-232.

Tversky, A., & Kahneman, D. (1974). “Judgment Under Uncertainty: Heuristics and Biases,” in *Science.* 185: 1124-1131.

Unger, P. (1967). "Experience and Factual Knowledge," in *The Journal of Philosophy*. 64: 152-73.

Waters, A., Gobet, F., & Leyden, G. (2002). “Visuo-Spatial Abilities in Chess Players” in *British Journal of Psychology*. 30: 303-311.

Wells-Jensen, S., Wells-Jensen, J., and Belknap, G. (2005). “Changing the Public's Attitude Toward Braille: A Grassroots Approach” in *Journal of Visual Impairment & Blindness*. 99: 133-140.

ENDNOTES

1.) The people plausibly interpreted to hold a version of this position include D.M. Armstrong (1968), D.G. Brown (1970), David Carr (1979, 1981), P.T. Geach (1957, 1966), Jaakko Hintikka (1974), A.J. Kenny (1966), Charlotte Katzoff (1984), Anthony Quinton (1970), Stanley and Williamson (2001).

2.) Though intuitions regarding many thought experiments in the know how literature are not completely univocal. I have found intuitions regarding the cases I discuss are widespread. Moreover, at the end of the paper I discuss the sources of contrary intuitions, noting how my positive theory naturally predicts and explains these intuitions without attributing any veracity to them.

3.) Not everyone agrees that a permanently blind person no longer know how to read. In fact, Stanley and Williamson blithely assert that blind people do know how to read. (Stanley and Williamson 2001, p. 416) I address such counter intuitions in sections 7 and 8. In the case of Carr's (CD3), he must qualify his conditions to avoid contradicting his own reasoning. Carr wants to allow that those who cannot currently perform, like an elderly arthritic piano teacher, can still have *knowledge how*. He asserts that "Perhaps it might be considered necessary in order to know how to ϕ that one should *at some time* have possessed the physical power to ϕ, but present possession of the ability does not appear to be necessary." (1981, p.53) I deal with worries about present ability in the next section.

4.) Readers of Haugeland (1981) should recognize this sort of thought experiment.

5.) For examples of this literature see Ford, D., Bergen, N. and Sterman, J. (1998), Chervinskaya, K. and Bekhterev, V. (2000), and Sternberg, R. and Horvath, J. (1998). Stanley and Williams (2001) offer an alternative analysis on such cases in terms of beliefs about "ways". Koethe (2002) offers one, I think, successful criticism of this view.

6.) Carr's suggestion (1981, p.54) that "...we may adequately say all we need to about canine capacities by means of expressions like 'is able to' and 'can' understood in the physical power sense," strikes me as particularly inadequate. Sheep dogs are not merely physically able to herd sheep, they know how to herd sheep. Rats do not merely possess the physical power to negotiate mazes after learning, they know how to negotiate those mazes.

7.) People can indirectly control object boundary detection by, for instance, closing their eyes. However, they cannot directly decide either to detect object boundaries or not to detect them.

8.) Indeed, one can imagine the following conversation under such circumstances: "Do you know how to dance a semiphore rendition of 'Gray's Elegy'?" "Well, in some sense I do. I've been told that one of the many dances I know how to dance is supposed to be a semiphore rendition of 'Gray's Elegy'. But, I don't really know how because I don't know which one."

For instance, one wouldn't want to say that I know how to read this:
lwkl qnwkhuhi ruhldp