Know How

1. INTRODUCTION

Many philosophers since Ryle agree that we attribute knowledge how to individuals and that attributions of know how are as explanatorily relevant to our understanding of mentality as our attributions of beliefs and desires. For example, David Carr states 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)

Work in cognitive science focuses upon human ability as often as representations of the world which may be true or false. When computer scientists like Herbert Simon model the human ability to play chess in programs such as CHREST or CHUMP, they model the human ability to play chess rather than the capacity to have beliefs about chess or to represent current states of play. (de Groot, A. And Gobet, F., 1996; Gobet and Simon 1996, 1998; Waters, Gobet, and Leyden 2002) Research into human abilities produces programs best understood as models of know how, not models of beliefs. These programs are sets of condition action rules, bits of practical reasoning, not to be confused with representations of the world that are true and false. Yet, one evaluates such programs epistemically. A 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 generating 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, Cummins 1986) A similar acknowledgment of the importance of *know how* is evinced in cognitive psychology. Glass and Holyoak tell readers that "human intelligence depends on our ability to construct complex procedures for performing tasks." (Glass and Holyoak 1986, p.14) Stillings equates the pervasive and important declarative/procedural knowledge distinction in cognitive science with the knowing that/knowing how distinction in philosophy. (Stillings, et al. 1987, pp.18-21) In neuroscience, Cohen and Squire likewise equate knowledge how with procedural knowledge and propositional knowledge with declarative knowledge. (Cohen and Squire 1980)

Despite their seeming acceptance of the existence of know how and its apparent importance to the scientific study of human cognition, philosophers devote little attention to knowledge how. Few accept Ryle's dispositional theory, fewer still offer theories themselves. In the absence of a serious and informed theory of know how, philosophers and others lack theoretical 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 and others regarding knowledge how. One might well describe the need for a theory of knowledge how as paramount in the philosophy of cognitive science.

This paper attempts to remedy this neglect. I forward a reliabilist theory of knowledge how inspired by, and consistent with, computational theory and its applications within cognitive science. I defend my theory in several ways. I refute both arguments against my **approach** to knowledge how and arguments against my **theory** of knowledge how. I also outline the ways in which my theory captures numerous important phenomena associated with knowledge how. Specifically, my theory captures the task-specific nature of know how, allows one to understand ascriptions of know how in the absence of manifest ability, allows one to explain the absence of know how given ability, distinguishes between an activity and an achievement, provides the theoretic underpinnings of a competence/performance distinction, and proves consistent with the physicalistic, computational explanations of cognitive science.

2. COMPUTATIONAL EXPLANATION AND KNOW HOW

To understand my definition of knowledge how, one must understand the structure and goals of computational explanations, the concept of a "well-specified cognitive task," and how such task specifications determine processes and relevance classes. As a result, I must sketch an outline of computational explanations in cognitive science before offering my theory of know how.

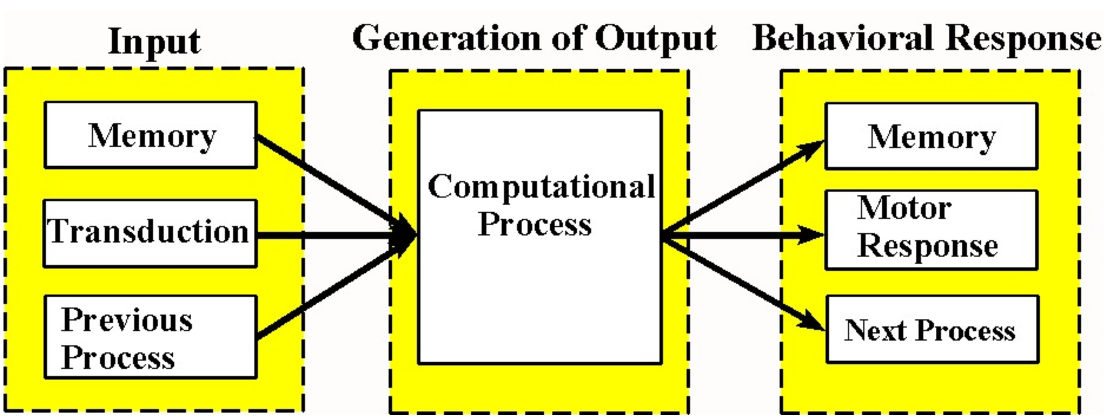

Computational explanations of ability involve three conceptual, though not necessarily temporally, distinct stages. Stage one involves the input of representations; from other processes, from the transduction of physical parameters (i.e., the direct conversion of physical parameters in the environment into representations of those parameters), from memory, or from some combination of the three. The system operates upon represented input, generating representations as output in the second stage. Cognitive scientists use a computational model for the second stage, though models vary from more traditional production systems and theorem provers to massively parallel and chaotic systems. Representations generated in the second stage guide the system's performance in the final stage involving some motor response, input to memory, and/or input to another process(es). For example, within one widely accepted explanatory framework for understanding how one successfully picks up objects, the computational cognitive psychologist refers to transduced input from the retina (from light reflected or projected onto the retina), the computational processes whereby one generates representations of the nature and location of the object from retinal input, and the motor response utilizing that object representation to perform the pick up.

Within cognitive science, one explains a system's performance in terms of a well-specified cognitive task: One shows 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] (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, a planning program has goal states and initial state descriptions as input and 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. (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) (3) Finally, a **WSTD** includes a specification of the relevant process 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.

Using the above notion of a well-specified task, one can address two of the main problems facing reliabilism.[2] Every reliabilist theory must solve the characterization and relevance class problems. Reliabilism ties knowledge to the reliability of the process or indicator generating the output. Because one can characterize a process in different ways, each having potentially different levels of reliability, reliabilists must forward a plausible hypothesis as to the appropriate manner of process characterization. For example, does my ability to see my hand result from the highly specific process linking the particular pattern of retinal stimulation to the particular activity pattern of my visual cortex and beyond? Or, is my ability to see my hand an instance of a process with wider applicability? Also, since the reliability of a characterized process depends upon the class of situations (alternatives) used to evaluate its reliability, reliabilists must forward a hypothesis as to how one determines the relevance class (the relevant alternatives). For example, the reliabilist must be able to determine if the large numbers of cases in which illumination is below the levels necessary for cone function [below which one cannot see colors] count against the reliability of color vision as alleged by Pollock (1984).

My account identifies the appropriate process characterization with the one instantiated by the system relative to the task description. A system instantiates a process when and only when it exhibits lawful state transitions between inputs and outputs across the task domain. As a result, the process characterization reflects the cognizer's lawful I/O relations across all and only those cases falling under the task specification. So in the absence of arguments that my seeing my hand involves different physiology or that conditions violate Marr's underlying physical assumptions regarding human vision, we can follow Marr and understand my seeing my hand as an instance of seeing, i.e., of a widely applicable process through which my brain computes "...shapes and their spatial organization in an object-centered coordinate frame,...." from information about reflected light as encoded in retinal activation patterns. (Marr 1981, p.37) One determines the set of relevant alternative situations by reference to the nomic correlations underlying the system's performance of the task. The task specification determines the system's environment of habitual performance, which is identified with the task domain. The task may change if the system's environment changes, altering the relevance class. One evaluates the effectiveness of the process (the strategy) relative to one's best hypothesis as to both the situations and relative frequency of situations the cognizer will encounter in performing this new task. So, one can eliminate low illumination cases as relevant alternatives for evaluating the reliability of color vision because they violate the **WSTD**. That is, one can eliminate cases of low illumination both because low illumination violates the preconditions for cone functioning and because people do not rely upon color vision when color vision becomes severely impaired by low ambient light levels.

My theory of knowledge how provides one with a unique process characterization and task domain given a **WSTD**. However, formulations of task are based largely upon empirical investigation and argumentation as to, for instance, the system's laws of operation, the covering laws operant in the hypothesized domain, and the relative distribution of potential situations within the hypothesized domain. While one will not find widely disparate task descriptions, and hence one will not find widely disparate hypotheses about relevance classes and process characterizations, there remains the possibility of disputes among informed individuals as to the exact formulation of a task. Likewise, task descriptions may vary in their level of specificity. A natural outcome of my theory is that the less well-specified a task, the more vague the ascription of know how.

Except for its task-specific and nomic nature, I forward a theory of *knowing how* much in the spirit of Gilbert Ryle's *Concept of Mind*:

> In judging that someone's performance is or is not intelligent, we have,...in a certain manner to look beyond the performance itself. For there is no particular overt or inner performance which could not have been accidentally or 'mechanically' executed by an idiot, a sleepwalker, a man in a panic, absence of mind or delirium or even, sometimes, by a parrot. ...in looking beyond the performance itself,.... We are considering his abilities and propensities of which his performance was an actualization. Our inquiry is not into causes (and *a fortiori* not into occult causes), but into capacities,.... (Ryle 1949, p.45)

One gets a preliminary idea of what I intend by considering what knowledge how, as I use it, provides: The concept of knowledge how provides a ground-level epistemic evaluation of a system's ability to exploit a computational process relative to some task. It distinguishes systems that can do epistemic work in some domain from systems that cannot do that work. Specifically, knowledge how integrates sensation, knowledge, and action so as to respect constraints on task performance in a fluid and variable environment. Ascriptions of know how are made in explaining the difference between a consistent, reliable ability to perform some cognitive task and an inability and/or an unreliable ability. Indeed, cognitive scientists appeal to knowledge in explaining behavior because that behavior systematically and reliably outperforms chance for the mechanism type. In other words, if one lacks an ability to reliably perform some cognitive task when nothing about one’s basic laws of functioning guarantees such performance, one has no motivation to ascribe knowledge how. For instance, humans and other animals exhibit striking differences in performance of a wide variety of behaviors both within individuals across time and circumstance and across individuals. Humans, for instance, differ widely in their ability to direct the inelastic collisions of billiard balls in order to drive them into pool table pockets. A system's consistent, reliable ability, where no mere reductive conglomeration of operant laws serves to explain the phenomenon or its variability, implies that one knows something. There must be something which allows one to function at a significantly better-than-chance level. I further claim that one cannot make such evaluations by reducing knowledge how to a knowledge of practical rules, as epistemic evaluation turns upon *how* the system uses its information, not merely upon what information it represents.

Of course, computer scientists continue to develop tools for talking about the efficiency and power of programs. I claim that the cognitive scientist's interests differ from the computer scientist's in two ways: First, unlike the cognitive scientist, computer scientists are interested in which system can perform the task in the fewest steps, with the least input, etc.. Second, whereas computer scientists make their evaluations based upon the execution of written programs on digital machines, cognitive science needs a more abstract evaluation that does not presuppose any particular architecture or style of computation. The computer scientist's concerns are quite legitimate, but they are the interests of people wanting to maximize the work-to-time ratio, etc. for a specific type of system, not the interests of someone trying to explain how a system (possibly an inefficient system) works at all. With the above in mind, I define *knowledge how* as follows:[3]

Ψ *knows how* to Φ iff
(1) Ψ is a system,
(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 **WSTD** of Φ. The notion of instantiating a process, as stated earlier, is to exhibit lawful state transitions between input and output across the task domain. To satisfy the satisfaction conditions associated with the idealized target function, the input/output pair generated by the system must be elements of the input/output sets mapped by the function. The definition also refers to three other features. It requires tractability of the computational solution: One can attribute *knowledge how* to a system only if one can characterize the system's process, its lawful state transitions, as a tractable strategy for performing the task. The system, in other words, must instantiate a process that can generate outputs from inputs given the constraints inherent to the task.

A negative example best clarifies tractability of computational solutions: Ted writes a program for his standard desktop PC that employs an exhaustive search method to play chess. Ted's program computes its next move by first generating scenarios for every possible game given the state of play, then choosing from that set of scenarios a move sequence that results in a win. Ted has a winning strategy. Unfortunately, Ted has failed to notice that his computer must generate approximately 10^{120} different moves (the combinatorial result of the average number of legal moves per turn, 31.6, and the average moves per game, 40). However the number of moves Ted's computer must generate is 10^{120}, more moves than seconds since the big bang. Ted's program might "eventually" play a game of chess, but not on Ted's computer. For Ted's ordinary desktop computer the task of generating and storing that many moves in memory, and then of searching and evaluating them in time to finish a single game of chess before the end of the universe is impossible. Ted's strategy represents an intractable process for playing chess. The strategy is not *theoretically* unsound. It plays a smashing game of chess in theory. Rather, the strategy is *practically* unsound. Ted's system cannot instantiate Ted's strategy so as to play chess given the limitations of memory and time associated with the game.

Tractability of computational solutions is not a machine-independent notion. Rather, it is relative to the system and task-specification under consideration. A program may prove a tractable strategy given the speed and memory of a super computer, yet prove intractable for an early desk-top computer. Likewise, tractability is not merely a notion about speed and memory vs combinatorial explosion: A nonalgorithmic program proves a computationally intractable solution--even on a super computer. Additionally, one ought not to confuse tractability with reliability. A system may instantiate a tractable strategy without proving reliable. In such a case, the system would be capable of generating outputs given the constraints of the task, but incapable of generating outputs, the majority of which satisfy the satisfaction conditions associated with the target function.

Even a reliable, tractable strategy must still have access, via inputs, to the relevant parameters in the task domain. Inputs must satisfy the satisfaction conditions for inputs in order to generate outputs that satisfy the satisfaction conditions for outputs. The system must be able to utilize the output of the process to guide its behavioral response (i.e., it must be capable of storing the output in memory, using it to guide motor response, and/or using it as input to the next process). In short, the system must not merely instantiate a rationale in its systematic interactions with its environment, it must also *exploit* that rationale. The system must employ a practicable strategy, and it must operate in ways that allow it to interact with the world so as to take advantage of that strategy by accessing the relevant parameters within the task domain, and utilizing the output(s) in its response.

A third feature of my account of knowledge how is the evaluation (via task-specific reliability) of the epistemic effectiveness of a process in general ***and*** the effectiveness given that its inputs satisfy the satisfaction conditions associated with the idealized target function. I claim that in order to possess knowledge how a system must instantiate a strategy that reliably generates outputs satisfying the satisfaction conditions for outputs. Though the above-mentioned general reliability requirement may seem sufficient, a moment's reflection reveals that general reliability cannot stand alone:

Suppose a system, S^{gr}, instantiates a process inferring the presence of *B*s from the presence of *A*s by explicitly representing the conditional '($A \rightarrow B$)'. Suppose that the minimum level of general reliability required for knowledge how is .73, and that .6 of the cases within the task domain where there is an *A* are cases where there is a *B*. Taken on its own, S^{gr}'s inference does not prove reliable enough for its process to count as knowledge how. However, suppose that .9 of the cases within the task domain have a *B* present, and that S^{gr} falsely tokens '*A*'s 15 percent of the time. In the case as described, the system has a general reliability of .735, as the number of *B*'s S^{gr} tokens because of false *A*'s, but in the presence of actual *B*s, is 13.5 per hundred (.9 x .15 x 100). In other words, we have two cases where the rule causes the system to infer a B is present. In the first case, it correctly tokens an A and infers a B. .6 of those cases are correct inferences. In the second case, it incorrectly tokens an A from which it infers a B. Given the base rate (.9) this second process will result in correctly tokening a B .135 percent of the time. Thus, the total percent of correct tokenings of B is .6 + .135, or .735.

If one takes general reliability as the sole epistemic measure, S^{gr}'s process counts as knowledge how. However, one cannot trace S^{gr}'s achievement of the minimum level of reliability (.73) to the effectiveness of its inference strategy. The additional .135 reliability is a side effect of S^{gr}'s faulty tokening of *A*'s in an environment rich in *B*s. In other words, the domain's large population of *B*s explains S^{gr}'s success, not its inference strategy. Requiring the conditional reliability of the process be greater than or equal to the general reliability controls for such side effects, and insures that the process works as an effective epistemic engine.

A less abstract example of the above point comes from early medical diagnosis. A nineteenth century doctor in New York was widely known for his diagnostic prowess with regard to hepatitis. Using his technique of tongue palpitation, he could predict with great reliability the eventual development of hepatitis long before the appearance of any symptoms. His success, however, can be explained not as the result of knowledge how, but as the result of his acting as a vector for the spread of the disease. Since hepatitis was very common at that time, the doctor would proceed through the ward palpating tongues until he came to a person who had the disease, from that point on he would spread the disease to each new patient.

Could one evaluate epistemic effectiveness solely in terms of conditional reliability? Suppose that a process has a conditional reliability of .9, but it also has six independent inputs, each of which has a .9 reliability. It has a .9 conditional reliability. Nevertheless, the process generates veridical output only 47.10069 percent of the time. A system can have a high conditional reliability, but still fall victim to equivocation. Hence, a general measure of reliability acts to insure that the system performs the task.

If one is low on funds at a bar, one can exploit the accumulation of possible error inherent in conditional probability as follows: Bet someone that they cannot guess the serial number of a bill, even if you give them nine guesses for each number. The person will quickly realize that the chances of guessing each number is .9 and will often take the bet. However, as there are eight numbers in the serial number the odds of their getting every number correct are .43, giving you an advantage which accumulates wealth over time.

One can, of course, simply set the minimum level of conditional reliability high enough to control for equivocation. But, one has to know the minimum level of general reliability associated with successful performance of the task in order to determine the level of the conditional reliability. Hence, one must again determine the minimum level of general reliability. Deleting general reliability from the definition of knowing how deceptively masks the role of general reliability in epistemic effectiveness.

3. ARGUMENTS AGAINST MY APPROACH

Many philosophers deny that one can explicate knowledge how in terms of a system's functioning.[4] One debunking strategy appears repeatedly in the literature. Practitioners of this strategy would eliminate knowledge how, arguing that all ostensive cases of knowledge how are misdescribed cases of explicitly represented knowledge, i.e. propositional knowledge. Specifically, debunkers attempt to understand knowledge how as a special subset of explicitly represented knowledge: They equate knowledge how with practical knowledge, contrasting practical knowledge with theoretical knowledge. Knowledge how, according to this strategy, involves knowing some particular set of condition-action instructions or theoretical knowledge of cause-and-effect. For instance, Carr characterizes knowledge how 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 a *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)

Within the framework of this debunking strategy, knowledge how becomes a theoretically interesting subset of represented knowledge. Each knowledge, knowledge how and propositional knowledge, has its own characteristic logic. Knowing, "to start the car, turn the key," is practical propositional knowledge. Knowledge that, "combustion engines generate force by igniting compressed gases within cylinders," counts as theoretical propositional knowledge.

These debunkers often note that many computer scientists apparently share their view to some extent. Computer scientists sometimes dismiss knowledge how because in writing programs, ultimately, knowledge how amounts to getting the right propositional knowledge into the system--getting the system to explicitly represent knowledge. Such and argument strategy looks particularly effective if all of the systems with which one deals operate by means of recursive applications of the same simple control mechanisms (i.e., the same simple operations on representations). Of course, not all computer scientist approach the modeling of cognition using the big data base strategy. Perhaps the most notable exception to the practice of modeling cognition through large data bases is the work of Rodney Brooks (1991, 1991a, 2002). Indeed, Brooks claims that

> The relationship to traditional academic robotics and traditional artificial intelligence is examined. In the new approaches a tight coupling of sensing to action produces architectures for intelligence that are networks of simple computational elements which are quite broad, but not very deep. Recent work within this approach has demonstrated the use of representations, expectations, plans, goals, and learning, but without resorting to the traditional uses, of central, abstractly manipulable or symbolic representations. (Brooks 1991, p. 1227)

> The fundamental decomposition of the intelligent system is not into independent information processing units which must interface with each other via representations. Instead, the intelligent system is decomposed into independent and parallel activity producers which all interface directly to the world through perception and action, rather than interface to each other particularly much. (Brooks, 1991a, p.139)

Moreover, reducing knowledge how to something akin to propositional knowledge remains a strategy only, a strategy that is largely a side effect of the sort of computational systems that computer scientists most often use as their models (i.e., standard digital computers). Not all computational systems operate using a small set of simple control mechanisms. Nor do all computational systems operate using the same control mechanisms. Production systems may operate using a simple matching algorithm and backtracking. Theorem provers might compute the same function using resolution and a unification algorithm. If one is to understand these systems, one must understand the sense in which these systems differ as well as the sense in which they compute the same function. Furthermore, it remains unclear that the best, or the only, method of learning for computational systems involves generating an appropriate set of representations.

Accordingly, replacing knowledge how with represented knowledge is a more complex, less powerful move than it first appears. In fact, when one finishes reducing the notion of knowledge how, something remains--even in the strongest cases. One cannot completely reduce knowledge how for two reasons: First, some systems implement an algorithm by instantiating a complex set of control mechanisms and *absolutely no represented instructions*. There are, in other words, systems that do not explicitly represent any knowledge, yet manifest the same knowledge how as systems that explicitly represent knowledge. Such is arguably the case for many connectionist systems, and is clearly the case for Turing machines not implemented on universal Turing machines. In these systems, physics acts to control the inferences by disciplining the state transitions. Second, even in cases where a system has relatively few control mechanisms and explicitly represents its knowledge as a set of complex instructions, the efficacy of those instructions depends directly upon the basic control mechanisms.

Altering basic control mechanisms while keeping the representations fixed can have a dramatic negative effect upon performance.

Moreover, the possibility (mentioned above) of systems that lack representations of procedural knowledge is not merely theoretical. Cognitive scientists have long noted that humans operate with few representations and complex control mechanisms in many cases. Vision researchers claim that the human visual system knows how to, for instance, infer edges from intensity values in retinal arrays, group edges into objects, and place those objects relative to one another in three dimensional space. Nevertheless, researchers do not suppose that these human visual abilities manifest themselves as the result of propositional knowledge of procedures for making such inferences explicitly represented in the brain. Rather, they claim that the visual system instantiates such procedural knowledge without representing it as instructions: The visual system's lawlike behavior as observed at each level of analysis ranging from gross behavioral level to the level of single neurons constitutes the basis for the ascription of knowledge how. (Hubel & Wiesel 1962 and 1968) Recall that connectionists like Rumelhart 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) I conclude, therefore, that though one cannot ignore representations of complex instructions in understanding a system's knowledge how, explicitly representing knowledge is insufficient as well as unnecessary for manifesting knowledge how. An adequate theory of knowledge how must be consistent with both possible implementations of knowledge how. My theory is consistent.

This sort of argument against knowledge how has recently been revived by Stanley and Williamson (2001) and Carl Ginet (1975) in the form of a criticism of Ryle's argument against intellectualism.[5] According to Stanley and Williamson (S&W), Ryle's argument has two premises:

(1) If one F's, one employs knowledge of how to F.
(2) If one employs knowledge that p, one contemplates the proposition that p. (Stanley and Williamson 2001, p.416)

S&W restrict (1) to intentional acts. As should be clear, I think S&W's restriction is not the appropriate one (either for Ryle or for knowledge how generally), but differences as regards (1) do not matter to the analysis of this argument.

S&W focus upon (2), suggesting that Ryle's regress argument relies crucially upon what I'll call the "employment is contemplative" premise. However, as shown above one need not and should not assume that Ryle depends upon the employment is contemplative. Indeed, Ryle is not really arguing against propositional knowledge by assuming (2). As I read Ryle, he is really arguing as follows: If propositional knowledge is the only type of knowledge, and knowledge guides all intelligent action, then one must explain the generation of intelligent action using only propositional knowledge. That is, "to act rationally was to have one's non-theoretical propensities controlled by one's apprehension of truths about the conduct of life." (Ryle 1949, p.26) The model for such causation of behavior via propositional knowledge is conscious reasoning. In the practical realm, actions are thought to be guided by conscious practical reasoning. Therefore, if practical reasoning cannot causally account for all intelligent actions, then propositional knowledge cannot be the only type of knowledge.

The real thrust of S&W's criticism is directed against (2), the employment is contemplative premise. But, Ryle only considers (2) because (2) exhausts the causal theories offered by the intellectualist tradition with regard to intelligent action. Thus, the characterization of the argument by S&W is a straw man, and (2) is a red herring. **Ryle's real argument against propositional knowledge is simply that it cannot, by itself, provide one with a causal explanation of intelligent behavior in all cases.**

But suppose S&W's arguments were not a straw man. Even if Ryle commits himself to (2), S&W's argument does not go through. S&W want to argue that (2) is false. Their argument really is just an adoption of Ginet's claim:

> I exercise (or manifest) my knowledge-that one can get the door open by turning the knob and pushing it (as wel1 as my knowledge-that there is a door there) by performing that operation quite automatically as I leave the room; and I may do this, of course, without formulating (in my mind or out loud) that proposition or any other relevant proposition (Ginet 1975, p.7).

According to S&W, Ginet's observation, "...brings out [sic] that employments of knowledge-that are often unaccompanied by distinct acts of contemplating propositions." (Stanley and Williamson 2001, p.415)

The problem with the above passage, and with construing it as an argument against (2), lies in the fact that Ginet's just quoted description of the events is blatantly question-begging. There is clearly a difference between acting in ways consistent with a belief one may have and "exercising (or manifesting) a belief." I take it that the latter, but not the former, supposes that the belief actually played a causal role in the production of the behavior. For example, I act in ways consistent with my beliefs about universal gravitation all the time, but I do no so act because of my beliefs about universal gravitation. If, as Ginet claims, he has no conscious awareness of the belief when performing the action, he has no reason to suppose that the belief played a causal role in generating the behavior. Thus, Ginet's example establishes only that he regularly acts in ways consistent with his beliefs about knobs and doors. It does not establish that his beliefs regularly cause his door-opening behaviors without his contemplating their propositional contents. Ginet merely supposes that his beliefs cause his behavior in such cases, precisely because he is biased towards such explanations of his behavior. Worse still, Ginet and S&W are, in fact, wrong to suppose that their door-operation beliefs are causally relevant in the production of the behavior. Neurological evidence regarding such behavior, though perhaps obscure to epistemologists and philosophers of language, clearly implicates areas of the brain other than those associated with propositional knowledge (inferior temporal lobe) in the causal generation of such behavior (e.x. the basal ganglia and the motor areas). (See for instance, Doya 2000, Mogenson, Jones, and Yim 1980, Pascual-Leone, Grafman, and Hallett, 1994, Schacter 1989, Smeets, Marin, and Gonzalez 2000.)

Worse still appeal to beliefs does not provide one with a causal story for belief. In fact, the sorts of stories that one tells about beliefs (i.e., functional/dispositional accounts) are those which involve the brain **knowing how** to operate upon the inert structures of representations to reliably produce appropriate behaviors. This point has long been standard fare among philosophers of mind, even staunch advocates of belief-desire psychology. Hence, Fodor's appeal to elementary operations that "...the nervous system is specifically wired to execute," in *The Appeal to Tacit Knowledge in Psychological Explanations.* (Fodor 1968/1981, p. 66), and Cummins' claim that cognizer's possess "*...inexplicit Information*, i.e., the information that exists in a system without benefit of a symbolic structure having the content in question. (Cummins 1986, p. 117).

So, neither computational theory, the use of computational theory in cognitive science, nor observations about the alleged non-contemplative role of beliefs in generating intelligent action support the complete reduction of knowledge how to propositional knowledge. What about other traditional philosophical arguments against knowledge how that one might adapt against my approach to knowledge how? Consider the sort of criticisms of Ryle typified by David Carr (1981): Carr critiques Ryle's dispositional analysis of knowledge how on the grounds that Ryle must, without resorting to trivializing *ceteris paribus* clauses, specify the conditions under which one manifests one's dispositional knowledge how. If one analyzes knowledge how to play chess in terms of dispositions to move chess pieces, one must specify the conditions under which a player manifests those dispositions. Such specifications are difficult since a person's reaction to a chess board and an invitation to play is subject to all manner of physical and psychological preconditions. If the person has been injected with curare, they will likely not react. If they believe themselves to be overmatched, they may demur.

However, the difficulties of dispositional analyses do not seem relevant to my own approach to knowledge how. My approach does not explicate knowledge how in terms of dispositions to manifest behavior. My approach to knowledge how explicates knowledge how it terms of the laws of functioning of the system within a well-specified domain and the effectiveness of that functioning given the system's typical interactions within the domain. My theory asks two questions: (1) Do the laws of a suspected chess system (within the domain) allow one to characterize certain of its state transitions as a process(es) computing legal, effective moves from representations of states of play? (2) Do those laws hold in the cases (if any) where the system has veridical access to states of play? Hence, my approach to knowledge how is not committed to either Ryleian dispositions or specifying the conditions for their manifestation.

The last criticism of my approach I'll consider again stems from Stanley and Williamson. In arguing against Ryle's account (as well as Bechtel and Abrahamsen 1991)

S&W make something like the following two arguments. The first argument is really a denial of an argument (not made by me) from the premise of the syntactic difference between sentences ascribing knowledge how and sentences ascribing propositional knowledge: Arguments for the distinctness of knowledge how allege a syntactic difference between sentences ascribing knowledge how and sentences ascribing propositional knowledge. For instance, consider the following two sentences:

(1) Ted knows how to tango.

(2) Ted knows that the tango is a Latin dance.

According to the syntactic argument, (2) requires a proposition, whereas (1) requires an infinitive. These theorists conclude on the basis of this alleged syntactic distinctness, that there is a real ontological difference between knowledge how and propositional knowledge. According to S&W statements ascribing knowledge how are not, in fact, "importantly" syntactically different from sentences ascribing propositional knowledge. Specifically, both classes of sentences consist of a subject, verb, and a tensed or untensed clause, bracketed in the sentences below:

(1) Ted knows [how to tango.]

(2) Ted knows [that the tango is a Latin dance.]

Theorists analyze the syntax of both classes of attribution sentences in terms of a phonological null pronoun (silent) in the subject position of the tensed or untensed clause as below:

(1) Ted knows [how PRO to tango.]

(2) Ted knows [that PRO the tango is a Latin dance.]

Hence, "...it is incorrect to take 'know how' as a constituent in sentences such as [(1)]." (Stanley and Williamson 2001, p.418)

I am in full agreement with S&W's conclusion regarding the syntactic dissimilarity argument, though I'm persuaded of the conclusion by reasons other than those offered by S&W. Arguments from syntactic similarity or dissimilarity are, on my view, quite misguided in that to establish a significant ontological difference purely on the basis of a syntactic difference they must hold something like the following: **Syntactic differences occur in all and only those situations where there are real ontological differences.** The "all" guarantees that absent a syntactic distinction, there is no ontological difference**.** The "only" in the claim guarantees that distinctions, in fact, indicate ontological distinctions. Nothing like this view strikes me a remotely plausible. After all, according to this view there would be no significant ontological differences among any views expressed using standard subject, verb, object construction. Likewise, linguistic constructions such as cases would *prima facie* constitute serious ontological markers. Thus, the use of der Zug and die Bank tell us that any theory of such objects must take into account the inherently feminine nature of the latter and the inherently masculine nature of the former. Similarly, if it was true that syntax always indicates important things about the world. Then we would not expect it to differ very widely in its structures across different languages or times. Yet, languages are neither structurally homogenous nor do they exhibit perfect or near-perfect structural stability across time. For instance, take cases in languages: Russian has 6 cases, German has 4 cases, and Hungarian has a startlingly elaborate system of 20 cases to express prepositional meaning!!! One might modify or qualify the above premise or the conclusions one draws from it so as to make it more plausible. However, such modifications would prove derivative upon our being able to recognize significant ontological differences on independent grounds. Rumfitt (2003) provides readers with a forceful and perceptive argument against S&W by making precisely these sorts of points with regard to the syntax of sentences ascribing knowledge how across different languages.

Of course, this is not to say that syntactical differences can supply one with no clues or insights into the nature of the world. Rather, it is to assert that such clues and insights one might glean from studying the syntax of languages serve only as an investigatory starting point, a single piece of evidence to be weighed in light of other evidence.

As noted, S&W have two arguments. The second argument turns upon a second bit of evidence, a theory regarding the semantics of the two respective classes of attribution sentences. The semantics argument uses the general semantic framework developed by Lauri Karttunen to provide a propositional interpretation of statements attributing knowledge how to individuals. I ignore both the details of the development of S&W's framework and their claims that it can easily be adapted to other semantic frameworks to yield the same results. S&W's penultimate conclusions, then, is 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 knowledge of how to X to an individual 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 point of S&W's position is that it is not, in fact, 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 both Koethe (2002), though I press the point harder here that Koethe did his brief but insightful critical comment.

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 doubt the latter point. 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, 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) Second, ascriptions of knowledge how often involve 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. Third, as noted above 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. (Again, 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, knowledge how occurs in their absence–even when the person's beliefs actually explicitly contradict S&W's hypothesized beliefs, they are implausible posits on their own merit in many cases of knowledge how, and they are not causally operant in 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 encoded explicitly as particular linguistic or quasi-linguistic expressions of propositions 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), and Charlotte Katzoff

(1984). These are likewise the claims of relevance to epistemologists and to such programs in the philosophy of mind as belief-desire psychology.

4. ARGUMENTS AGAINST MY THEORY

Neither the reduction of knowledge how to explicitly represented knowledge, nor Carr's objections to Ryle's dispositional account undermine to my approach. My **approach** seems a legitimate tack. Nevertheless, one might still object to my **theory** of knowledge how. One might object to my definition of knowledge how because, in accordance with the task-specific nature of computational explanations in cognitive science, the definition is task-specific. I think, however, that reflection upon the ordinary notion of knowledge how supports a task-specific approach. For example, suppose that Bill's cakes always turn out perfectly. Bill uses normal baking methods, which he acquired in a normal manner. One would surely say that Bill knows how to bake a cake. Moreover, one would surely not deny Bill his cake baking know how simply because he does not know how to alter his method were he to move from the coast to a high altitude. Neither would one deny his cake baking know how because he cannot bake a cake in a microwave oven. Yet, such cake baking possibilities are well within the realm of culinary possibility. Bill's cake baking know how is task-specific.

One might claim that I assume that **all** knowledge how occurs in the context of a cognitive task when it is unobvious that **all** skills and capacities one might call knowledge how are involved in performing cognitive tasks. This objection is an oxymoron. Knowledge is definitive of the cognitive. Knowledge how is knowledge of how to do something. Hence, any ascription of knowledge how is *de facto*, an ascription of a cognitive capacity.

A different criticism claims that, by tying knowledge how to success conditions, my definition of knowledge how ignores the distinction between an activity and an achievement.

Activities that lack a criterion for success or have an ill-defined criterion can still require knowledge how. For instance, one can know how to play chess without being able to win many games.

My response to this criticism is twofold. First, I deny that one appeals to knowledge how in understanding activities with ***absolutely no*** success criteria. For example, before one is thought to know how to play chess, one must reliably demonstrate a familiarity with the rules for moving pieces, the starting positions, and etc.. Before one is thought to know how to play chess well, one must also be capable of winning regularly against opponents of some skill themselves. There is a distinction between knowing how to play chess and knowing how to play chess well. However, in either case one finds success criterion, so my theory is capable of underwriting an ascription of knowledge how. Indeed, I claim that without a criterion for success, one cannot understand the system's performance as appropriate or inappropriate for the circumstances, making it unnecessary to appeal to the system's knowledge of how to perform to explain its ability to reliably generate appropriate responses.

Second, and perhaps more importantly, nothing in my theory prevents the theory's application to cases where success conditions are ill-defined. A natural outcome of my theory is that ill-defined success conditions result in a corresponding vagueness in the ascription of knowledge how. This consequence is born-out in cases in which success conditions are ill-defined: When one says that one knows how to rock-climb, one's statement is ambiguous between a general familiarity with equipment and an ability to climb $5^{.5}$ (easy routes requiring relatively little technique), and an ability to climb $5^{.13}$ (extremely difficult routes requiring sophisticated techniques and careful choice and ordering of movements). Hence, one often asks, "How well do you know how to climb?" to distinguish the potentially dangerous braggart from the modest expert.

In a similar fashion, one might object that my theory, by equating knowledge how with the ability to reliably perform, ignores the competence/performance distinction. As a result, I cannot acknowledge that there are many ways to impair one's ability to reliably perform without undermining one's claim to knowledge how (competence). I take it that objectors of this last stripe realize that my definition can handle such cases as one's continued knowledge of how to ski in the absence of snow. Because preconditions for the performance of the task, like snow on the ground, are built into the task specification, failures to perform involving violations of the task specification do not undermine one's claim to knowledge how. The cases to which the objection must refer, then, are the standard sort of systematic errors labeled "performance errors" in linguistics. One philosophical gloss on such errors is that an average native English speaker, Betty, knows how to parse English sentences, i.e., she knows the grammatical rules by which the sentence was generated and how to apply those rules to reveal the sentence's structure. However, Betty cannot parse English sentences that are, for instance, a page long. Her failure is not one of competence on the standard linguistic account, i.e., knowledge how. Rather, she fails to parse them because she, for instance, lacks sufficient working memory to remember the sentence long enough to apply the rules.

I claim that the above objection, understood in the second way, is based upon a confusion regarding the sense in which one can legitimately be said to know how to parse a sentence. I claim that Betty does not know how to parse page-long English sentences. I base my claim upon the fact that Betty's process (using the grammatical rules of English to parse sentences) is ***in principle inadequate for Betty to parse page-long sentences***. Her knowledge of the rules of grammar and their application represents an intractable solution to the parsing problem on the standard stories wherein Betty's performance deviates from her competence because of factors such as the inherent limitations of Betty's memory. However, there is a relevant counterfactual sense of knowledge how allowed by my theory: Betty knows how to parse page-long English sentences if she has sufficient memory and time. My theory allows one to understand the claim that Betty knows how to parse page-long English sentences as a claim about what she would know how to do given enough time and working memory, etc.. Nevertheless, it is always false that Betty has sufficient memory and time, so it is false that she actually knows how to parse such sentences. It is no more plausible to claim that she knows how parse page-long English sentences based upon her counterfactual abilities than it is to claim that she knows how to get around Boston on the basis of the counterfactual knowledge how she would have if Boston were Chicago. Just as getting around Boston would be a different problem if Boston were Chicago, parsing page-long English sentences given one's resources, one's typical environment of performance, etc. is a different problem than parsing those sentences if one has unlimited time, memory, absent normal environmental confounds, etc.. Such counterfactual ascriptions of knowledge how, though understandable on my account, 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, or their being, the actual circumstances in which the subject is placed. My theory explicates competence in terms of the possibly counterfactual, yet explanatorily relevant, ability to perform a task given enough time, memory, and in the absence of interfering factors (e.g. distractions). Hence, my theory is consistent with a competence/performance distinction despite equating knowledge how with ability.

In addition to underwriting a competence/performance distinction, the task-specific, counterfactual nature of knowledge how in my theory allows for the sorts of ascriptions of knowledge how in the absence of ability one finds in the philosophical literature. For example, Carr (1981, p.53) discusses one's desire to attribute, at least in a weak sense, know how to a temporarily incapacitated gymnast or an arthritic piano teacher despite their current inability to perform. One can understand such cases as task-specific ascriptions of knowledge how in which the presupposed underlying conditions for the manifestation of knowledge how are currently violated. The gymnast does not know how to perform his routine with a broken leg. Nevertheless, he knows how to perform his routine given his normal physical abilities. The piano teacher does not know how to play the piano at concert level with arthritic hands. She would know how to play at that level if, contrary to fact, she had her former dexterity. For this reason, she can teach those having full movement of their digits to play.

One might also object to my theory along a different line:[6] Suppose that Ted, the unsuccessful programmer, develops a serum that he believes would make a person raise their arm in the presence of radiation. He puts it in my coffee, places samples of uranium around the department, and observes my behavior. To his delight, my arm raises every time I come close to uranium. Surely, the objection goes, one would not want to say that you know how to detect uranium as a result to Ted's subterfuge. Yet, on your account, you instantiate a process that reliably tokens outputs in the presence of radioactivity.

Scrutiny reveals two closely related problems with the example as formulated. First, it remains unclear that one would want or need to give a cognitive explanation of my behavior any more than one might want to give a cognitive explanation of my bleeding when you cut me with a knife. My arm-raising, like my bleeding, is a physiological response to the imposition of external force, and as such is best explained physiologically. If the behavior is best explained as a physiological response, and not as the performance of a cognitive task, then the seeming lack of knowledge how and the fact that the relevant Φ is not a cognitive task make both sides of the biconditional in my definition false (making the definition true).

Second, on my account of knowledge how I do not, in fact, instantiate and *exploit* a process in the above story. Though Ted's elixir changes my laws of operation to alter my behavior in the presence of radioactivity, it is clear that I cannot *exploit* change without further alterations of my laws of operation: Since radioactivity results only in a transient response, I do not store information about radioactivity in memory. My laws of operation do not allow me to use a representation of radioactivity to guide my hand-raising activities. That is, my response to radioactivity does not constitute a rationale for using radioactivity to guide and control arm-raising in the way that, for instance, I use a visual representation of an object in my visual field to guide my motor control response when interacting with that object. Nor does the arm-raising constitute a representation making information about radioactivity explicit to me so as to serve as input to another process. So, while Ted can take advantage of my hand raising to detect radioactivity, I cannot.

The requirement of my theory that a system instantiate and exploit a process in order to know how also allows my theory to explain cases of absent know how given ability as they occur in the philosophical literature. Carr (1979, pp.406-8) discusses the case of a dancer who performs a dance called improvisation #15. Unbeknownst to the dancer, the dance is a movement for movement semaphore rendition a Grey's "Elegy." The dancer seems to have the ability to perform a semaphore rendition a Grey's "Elegy," but he does not know how to perform a semaphore rendition of Grey's "Elegy." That is, he would respond only with confusion if such a request were made of him.

My theory offers the following analysis of this situation: 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 understood. It is this lack of a functional connection that prevents knowledge how in the case of the dancer.

I claim that knowledge 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 the ability and the eliciting conditions for the ability. What prevents the dancer from performing the semaphore Gray's "Elegy" is that he lacks a conceptual understanding that links the ability to the request as described. In other words, the cue, and hence the conceptual context, he associates with the movements--is "Improvisation #15," not "a semaphore Gray's `Elegy'."

Finally, one might claim that my theory allows that thermostats know how to regulate ambient room temperature as they reliably process representational inputs of the ambient room temperature to generate furnace control signals. Thermostats are patently nonmental systems. Only mental systems can posses know how. Hence, my theory proves false.

This argument presupposes that my theory is about or has ramifications for the mental/ nonmental distinction. I deny this claim. My theory is not a theory of the mental/nonmental distinction, nor does it necessarily have ramifications for the mental/nonmental distinction. There are many possible positions regarding the mental/nonmental distinction that defuse this objection and are consistent with my theory. For example, one might claim that thermostats fail to posses knowledge how. Though they represent ambient temperature in the environment and use this representation to control the ambient temperature via activation of the furnace, thermostats do not know how because knowledge how, like cognition, requires mental representation. Distinguishing between mental and nonmental systems on the basis of mental representations is common in philosophy of cognitive science and perfectly applicable. (Dretske 1988, Fodor 1986) In contrast, one might assert that knowing how requires no mentality in our sense, but is merely one property strongly associated with mentality--just as reproductive ability, metabolism, etc. are strongly associated with life. Thermostats know how to regulate ambient temperature, but this poses no threat to a mental/nonmental division as knowing how is only one of a set of properties associated with mentality, none of which is either necessary or sufficient in themselves for mentality.

5. CONCLUSION

I conclude from the above discussion that my theory withstands attacks against my approach to knowledge how as well as attacks against the theory. Moreover, my theory has remarkable explanatory power both with respect to cognitive science and philosophy. The theory captures the task-specific nature of ascriptions of knowledge how within cognitive science as well as within ordinary discourse in a framework that is consistent with the physicalistic, computational explanations of cognitive science. The theory provides the basis of a competence/performance distinction. It discriminates between an ability and an achievement. The theory captures both cases of knowledge how in the absence of ability and cases of absent know how given ability in the philosophic literature.

BIBLIOGRAPHY

Armstrong, D. (1968). *A Materialist Theory of Mind*. London: Routledge and Kegan Paul Ltd.

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

Brooks, R. A., (1991). "New Approaches to Robotics", in *Science* 253: 1227–1232.

Brooks, R. A., (1991a) "Intelligence Without Representation", in *Artificial Intelligence Journal* 47: 139–159.

Brooks, R. A., (2002). *Flesh and Machines*. NY: Pantheon Books.

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

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.

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.

Cole, K.J., and Rotella D., (2002) "Old Age Impairs the Use of Arbitrary Visual Cues for Predictive Control of Fingertip Forces During Grasp, " in *Experimental Brain Research*. 143(1): pp.35-41.

Corkin, S. (1968). "Acquisition of motor skill after bilateral medial temporal-lobe excision," in *Neuropsychologia*, 6:255--265.

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

Cummins, R. (1986). "Inexplicit Information" in Miles Brand (ed.) *Representation Of Knowledge and Belief*. University of Arizona Press.

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

Doya, K. (2000). "Complementary Roles of Basal Ganglia and Cerebellum in Learning and Motor Control," *Current Opinion in Neurobiology*, 10 (6).

Dretske, F. (1988). *Explaining Behavior*. Cambridge, MA: MIT Press.

Fodor, Ja. (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.

Fodor, J. (1986). "Why Paramecia Don't Have Mental Representations," in *Midwest Studies in Philosophy*. 10: 3-23.

Ginet, C. (1975) *Knowledge, Perception, and Memory*. Boston: Reidel.

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

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

Glass, A and Holyoak, K (1989). *Cognition*. New York, NY: Random House.

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.

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

Hubel, D. H. & Wiesel, T. N. (1962). "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex" in *Journal of Physiology*. 160: 106-154.

Hubel, D. H. & Wiesel, T. N. (1968). "Receptive Fields and Functional Architecture of Monkey Striate Cortex" in *Journal of Physiology*. 195: 215-243.

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(6): 325-328.

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

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

Mogenson, G. Jones, D., and Yim, C. (1980). "From Motivation to Action--Functional Interface Between the Limbic System and the Motor System," in *Progress in Neurobiology.* 14 (2-3): 69-97.

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

Nöe, A. (forthcoming). Analysis.

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.

Pollock, J. (1984). "Reliability and Justified Belief," in *Canadian Journal of Philosophy*.

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

Rumfitt, I. (2003). "Comments and Criticism: Savoir Faire" in *The Journal of Philosophy*. 100(3): 158-166.

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.

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.

Smeets WJ, Marin O, Gonzalez A. (2000). "Evolution of the basal ganglia: new perspectives through a comparative approach." in *Journal of Anatomy.* 196 (Pt 4):501-17 .

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

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

Tulving, E. (1984). "Precis of Elements of Episodic Memory" in *Behavioural and Brain Sciences*, 7, 223 – 268.

Tulving, E., & Schacter, D.L. (1990). "Primary and Human Memory Systems," in *Science*, 247, 301 – 306.

Waters, A., Gobet, F., & Leyden, G. (2002). "Visuo-spatial abilities in chess players" in *British Journal of Psychology*. 30: 303-311.

ENDNOTES

[*]This paper was developed in large part during my tenure as a Mellon Post-Doctoral Fellow at the University of Rochester and as a participant in a 1993 NEH summer seminar at the University of Arizona. I am grateful to my colleagues at Rochester and Arizona, especially Robert Cummins and Ted Sider. An early version of this paper was developed as part of my dissertation, *Representation, Knowledge, and Structure in Computational Explanations in Cognitive Science*. I am grateful for comments made then by Ronald Giere.

1. (1) and (2) correspond to Marr's "Computational Theory" (1981, p.24-29). (2) and (3) correspond to his "Representation and Algorithm" and his "Hardware Implementation" levels.

2. I discuss problems associated with reliabilism in relation to Alvin Goldman's theory, and offer a more complete version of my solutions in "Truth-Ratios, Process, Task, and Knowledge," in *Synthese* (Feb. 1993). The comments here are intended merely to give those aware of these problems a sense of how I address them, and to give the reader an outline of how my theory works and its relationship to a well-specified task.

3. In my dissertation, the conditional probability mentioned in the second clause is expressed in terms of a naturalized definition of knowledge for inputs. While I prefer that definition, I do not have time to present and argue for another definition here.

4. Among the debunkers with regard to my notion of *knowledge how* one might include 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 Anthony Quinton (1970). Though I deal with a computer science inspired argument in the next few paragraphs, I also have Dr. Katzoff in mind. My objections there, are objections to her as well.

5. I note here that S&W seem, like many others, to suppose that Ryle argued that knowledge how and propositional knowledge were two real categories. Ryle, of course, was a reductive behaviorist and wanted to establish the legitimacy of knowledge how, not in contrast to propositional knowledge, but as the first step in reducing all knowledge, including propositional knowledge, to knowledge how. Nöe makes a similar observation regarding S&W's reading of Ginet's example.

6. This example as well as other helpful comments were given to me by my then colleague Ted Sider at the University of Rochester.

ABSTRACT

I forward a reliabilist theory of *knowledge how* inspired by, and consistent with, computational theory and its applications within cognitive science. I defend my position by refuting arguments against my approach to *knowledge how* as well as refuting arguments against my theory of *know how*. I also delineate the way in which my theory captures the task-specific nature of *knowledge how*, allows one to understand ascriptions of *know how* in the absence of manifest ability, allows one to explain the absence of *know how* given ability, distinguishes between an activity and an achievement, provides the theoretic underpinnings of a competence/performance distinction, and proves consistent with the physicalistic, computational explanations of cognitive science.