

THE ROLES OF PHILOSOPHY IN COGNITIVE SCIENCE¹

Tim van Gelder
Department of Philosophy
University of Melbourne
Parkville VIC 3052
Australia
tgelder@ariel.unimelb.edu.au
<http://timpc.arts.unimelb.edu.au>

What does philosophy contribute to cognitive science? This question is addressed indirectly, by describing some of the many roles philosophers play. These include Pioneer, Building Inspector, Zen Monk, Cartographer, Archivist, Cheerleader and Gadfly. As a preparatory exercise, philosophers are characterized in terms of their primary methods: argument, conceptual analysis, and historical perspective. The various roles philosophers in fact play are seen to follow naturally from this way of characterizing the philosopher.

Cognitive science is always described as an interdisciplinary enterprise, and philosophy is almost always listed as one of the participating disciplines. Yet in the daily practice of cognitive science philosophy is usually less visible than other major contributors such as computer science, cognitive psychology and linguistics. This is reflected in what might be described as a “balance of citations deficit”—philosophers cite others much more than others cite them.² The suspicious might suspect that philosophy mainly gets listed for the same reason “postmodernists” always claim to draw on queer theory, feminist epistemology, poststructuralist theory, and the rest—viz., it sounds good. It might therefore be of some service to philosophy, and cognitive science more generally, if we pause a moment to take stock of philosophy’s actual contributions.

This turns out to be harder than it sounds. Philosophers have been practicing cognitive science, in one form or another, for millennia. Cognitive science overlaps with philosophy in many of its most central areas: the theory of knowledge, philosophy of mind, philosophy of language, logic. Most major topics within the other disciplines of cognitive science give rise to philosophical questions, and philosophers have found plenty of their own problems to worry about. And finally, philosophers disagree among themselves about most issues, so there is very little by way of current consensus to report.

¹ Appearing in *Philosophical Psychology*, 1998.

² I have no hard data to justify this assertion, but would be prepared to place good money on it.

Rather than try to summarize philosophy's contributions, this paper will try to describe the philosopher's methods, and some of the roles philosophers play within cognitive science. Along the way there should be enough examples for the general flavor, at least, of philosophy's particular contributions to become apparent.³

1. Who is the Philosopher?

Arguably, the best way to identify philosophers—or cognitive scientists who are thinking philosophically—is not by the subject matter they are discussing, nor by where they are housed, but by their *method*. Psychologists standardly run and analyze experimental studies, neuroscientists play around with brains, linguists gather and systematize linguistic data, and computer scientists write programs. Similarly, what philosophers primarily do is *argue*, though they also rely heavily on *conceptual clarification*, and to a lesser extent on *historical perspective*. Of course, other cognitive scientists can and often do use these methods as well; the point is that philosophers are unique in *specializing* in them.

1.1 Argument

Ever since Western philosophy's emergence, with Plato, out of professional advocacy and debating, argument has always been philosophy's most important tool. Thus, when trying to establish a point, the philosopher typically leans back and selects some premises, hopefully less controversial than the point at issue, from which that point supposedly follows. In assessing a position, the philosopher will either analyze the logic of the arguments that support it or search around for possible arguments against it. One thing the philosopher typically *doesn't* do, by contrast, is conduct any demanding empirical investigations into the way the world is; to do that would be to become some kind of scientist for a while. (Of course, the philosopher can take advantage of such work by others by using their findings as premises in arguments.)

For example, suppose we want to know whether the human mind is a machine of some kind - surely one of the most basic questions in cognitive science. Since the advent of digital computers, the most sophisticated machines ever built, this question has come into a new focus as the problem of whether the mind is nothing more a computer running an appropriate program. Now, the psychologist would presumably proceed to run subjects in an endless variety of tasks to see if

³ There are some previous chapter-length introductions to the philosophy of cognitive science that attempt a more issue-oriented approach. Perhaps the best available is (Garfield, 1987); see also (Harman, 1989). For accessible book-length introductions to those areas of philosophy of mind most relevant to cognitive science, see (Bechtel, 1988; Churchland, 1988).

human performance data can resolve the issue. The computer scientist would try to write the program, and run plausible candidates on a more convenient kind of computer. The neuroscientist would take a good hard look at the inner machinery. And so on. The philosopher, however, takes an easier course. All you really need is a good argument. John Searle provided the following classic, one of the most famous and infamous in the philosophy of cognitive science:⁴

Premise 1: Computer programs are formal (syntactic).

Premise 2: Human minds have mental contents (semantics).

Premise 3: Syntax by itself is neither constitutive of nor sufficient for semantics.

Conclusion: Programs are neither constitutive of nor sufficient for minds.

Quod erat demonstrandum. (Philosophers often like to throw in snippets of Latin.) All that work by psychologists, computer scientists, neuroscientists, etc., might shed a lot of light on human performance, computer programs, the brain, etc., but what it can't do, if this argument is right, is show that minds are nothing more than computers running appropriate programs.

Can progress in cognitive science really be this easy? Well, not exactly. The argument is short but the thought that went into formulating it was not trivial. Moreover, there have been dozens of attempted refutations; Searle and his followers have devoted considerable energy to defending the argument. So the philosopher's task is not really easier than anyone else's, its just different.

This argument, like all arguments, is a set of premises and a conclusion. The premises are meant to state uncontroversial truths - facts that most reasonable people already accept. (If any one is not, then of course it should in turn be supported by premises which *are* uncontroversial.) So what Searle is really trying to do is show that the idea that minds are just computer programs is inconsistent with other ideas that we all generally accept, and hence should be rejected. In general, the process of formulating and analyzing arguments can be seen as the way we work out the patterns of consistency and inconsistency among our current beliefs, and between current beliefs and new claims (e.g., ones like "minds are just computers running programs"). From this point of view the philosopher, whose primary technique is argument, is in the business of straightening out the cupboard of knowledge. If you ask a philosopher to settle an issue, he'll do it not by consulting the world directly but by seeing how it fits in with other things we know.

⁴ This version of the argument is quoted directly from (Searle, 1990). For other presentations see (Searle, 1980; Searle, 1989).

Since argument has always been the philosopher's primary technique, they have paid a lot of attention to the nature of argument itself. This has given rise to a whole sub-discipline, *logic*, which has turned out to be quite important for the development of cognitive science. However, reflecting on some of the most basic distinctions made by logicians reveals why philosophers don't use *just* argument: there are other techniques with which pure argument must be supplemented.

At the most general level arguments divide into two kinds, the deductive and the inductive. This distinction is a matter of the kind of support that the premises give the conclusion. In deductive arguments the premises are supposed to guarantee their conclusions, in the sense that *if* the premises are true, then the conclusions *must* be true. This guarantee is called *validity*. In inductive arguments, on the other hand, the premises are merely supposed to make their conclusions more likely; if the premises are true, then the conclusion is more probable.

Deductive arguments are in an obvious way preferable to inductive ones. Very often, however, it's difficult to tell whether a given deductive argument really is valid. Logicians have produced detailed rules and methods for this, rules and methods which apply if the argument is expressed in certain special *formal* languages, such as propositional or predicate calculus. For example, by following the rules it is simple to determine that the following argument is deductively valid in predicate calculus:

P1:	$(\forall x)(Cx \supset Sx)$	(If anything is C then it is S)
P2:	Ca	(A particular thing, a, is C)
C:	Sa	(a is S)

A key point here is that you don't need to know what the predicate-symbols C and S actually mean in order to know that this argument is valid. In arguments expressed in the right formal languages, deductive validity turns solely on the *structure* of the argument, where structure is a matter of having the logical symbols (\forall , \supset , etc.) interspersed in the right kind of pattern with nonlogical ones (C, S, etc.).

1.2 Conceptual Clarification

Now, philosophers would deal with nothing but deductively valid formal arguments if they could. The trouble is that most arguments of any real interest, in cognitive science as in most other areas, cannot be stuffed into this narrow format. Consider Searle's argument again. It is easy enough to provide a reasonable formal approximation, in predicate calculus, to the first two premises; for example,

Premise 1: $(\forall x)(Cx \supset Sx)$ (if anything is a computer program then it is syntactic)

But what about Premise 3? How do we formalize the claim that syntax by itself is neither constitutive of nor sufficient for semantics? Perhaps surprisingly, there is no way to express the idea that property A is not constitutive of nor sufficient for property B in standard first order predicate calculus.⁵ Consequently, we cannot rely on that set of rules to deliver a verdict on the validity on Searle's argument.

Nevertheless, most philosophers feel confident that they can decide whether Searle's argument is valid or not, and they make this decision without actually formalizing the argument, whether in predicate calculus or any other formalism. Since these decisions are not made purely on the basis of form or structure alone, there must be something more going on. That "something more" is the meanings of the terms involved. Our realization that the argument is valid depends, at least partly, on *our understanding what Searle is talking about*. And the same goes for virtually all interesting deductive arguments in cognitive science: the guarantee that the premises impart to the conclusion depends on what the premises and the conclusion *mean*.

A similar moral applies in the case of inductive arguments. There is, in fact, no formal language within which inductive arguments can be judged strong simply in virtue of form. Assessing the strength of an inductive argument always requires knowing what the argument is about, and hence what its terms mean. Consider for example the following inductive argument schema:

- P: In three trials on Task T, new-fangled AI program CoGNizer delivered the correct answer.
- C: CoGNizer is capable of Task T.

The amount of support the premises will actually offer the conclusion in any real case depends crucially on what T and CoGNizer are. Thus suppose T is *parsing English sentences*, and CoGNizer has, internally stored, the parse trees for exactly three sentences - just the ones that it was going to be tested on. Then the premise

⁵ The problem is the notion of "constitutive of". One could try formalizing Premise 3 as $(\neg x) \rightarrow (\neg Sx \rightarrow \neg SEx)$ (it is not the case that if any thing is syntactic then it is semantic) but this only captures the idea that syntax is not sufficient for semantics. Searle, in his formulation of the premise, is careful to say that syntax is "neither constitutive of nor sufficient for" semantics. Presumably this is to deny the existence of some intimate connection between syntax and semantics which is not the same as logical sufficiency but which would nevertheless satisfy enthusiasts of Strong AI. However, there is no obvious way to render "constitutive of" in standard first order predicate calculus. There may be senses of "constitutive of" that can be rendered in some other, less familiar formal languages, but even if this is true, one thing is certain: philosophers and others evaluating Searle's argument don't as a matter of fact translate it into any such formalism, and yet are prepared to deliver verdicts on its validity.

offers virtually no support for the conclusion. If on the other hand, the three trials were on randomly chosen sentences of English, then the premise offers quite impressive support, despite such a small number of trials. So, you can't tell simply by looking just at the *form* of this argument whether it is a good one or not.

Of course, many of the factors that are relevant to the goodness of an inductive argument can be made explicit and entered as premises, and there are plenty of statistical techniques for estimating the reliability of the conclusion given the information in the premises. The problem is that despite the best efforts of logicians for centuries, such techniques are in general not capable of turning the task of assessing the goodness of a real-life inductive argument into a pure matter of inspecting form alone. There is an ineliminable component of non-formalizable *understanding* and *judgment* which enters into assessing inductive arguments.

There is yet another major reason why assessing the quality of an argument requires understanding the meaning of the terms involved. Thus far I have been talking mostly about the strength of the connection between premises and conclusion, but in good arguments the premises must also be *true*. Skirting some philosophical difficulties, we can say that a sentence is true if it properly describes the way things actually are, and this depends in part on the *meaning* of the words in the sentence. Thus, as everyone knows, whether Searle's third premise

Syntax by itself is neither constitutive of nor sufficient for semantics. is true depends on what "syntax," "semantics," and even "constitutive of" actually mean. Determining what these terms mean is a notoriously tricky business, and the difficulties here largely account for the great amount of disagreement over whether Searle's conclusion must be accepted.

So in at least three major ways, dealing with argument effectively crucially depends on another process, that of getting clear on the meanings of the terms involved. Now, there is a very close link between the meanings of terms and *concepts*: basically, clarifying the meaning of a term just is clarifying the concept which corresponds to that term. Thus an alternative way of describing this activity is as *conceptual clarification* or *conceptual analysis*. And so, since argument is the philosopher's primary technique, it follows that philosophers are going to be spending a lot of time doing conceptual clarification as well.

How do philosophers go about clarifying concepts? There are at least three standard methods. Probably the most common is to consult intuitions. That is, philosophers first assume that they already *have* the concept in question. By virtue of being ordinary people, able to cope in ordinary ways with other people, language, etc., philosophers already have the concepts of *mind*, or *belief*, or *pain*, etc.; or by

virtue of being immersed in the world of cognitive science, they already have certain technical concepts such as *syntax*, or *semantics*. The problem is simply to get *clear* on the concept, i.e., to make its structure totally explicit. This is where intuitions come in. You entertain a series of propositions involving the concept, and decide whether you'd be inclined to accept them. Some are found acceptable, and some unacceptable, simply by asking yourself whether, "intuitively," you agree with them. Either way, some light is thought to be shed on the concept itself. The idea is that once you know what all the intuitively true propositions in the vicinity of the concept are, and what all the intuitively false ones are, then you have made the nature of the concept (at least, the one you had in mind) explicit.

For example, suppose you want to clarify the concept of *meaning* or *semantics*. You might, as Searle invites you to do, entertain fanciful propositions of the following kind:

If I (not a Chinese speaker) were in a room manipulating Chinese symbols according to a rule-book, those symbols would mean nothing to me.

This proposition is clearly in the vicinity of the concept of *semantics*. Moreover, most people seem to have a pretty strong hunch about whether it is true or not, and they have that hunch by virtue of an "intuitive" prior grasp of the concept. Your verdict on a proposition like this thus sheds light on the concept as you hold it; think of it as a kind of probe into its structure. The idea is, that by using many such probes, you can bring the concept itself into clear focus; and having done that, you may be in a position to formulate some succinct general claims about it, such as

Mental phenomena have intrinsic meaning.

Such claims are the intended reward of successful conceptual clarification via intuitions, for with them in hand you are in a position to formulate strong and interesting claims and arguments; indeed, to undertake such ambitious enterprises as refuting the idea that minds are simply computers.

However, conceptual clarification via intuitions must be handled carefully. Used appropriately, it can be a valuable resource. There are times when "thought-experiments"—consulting intuitions about hypothetical situations—are the only available way of exploring some aspect of a concept. However, philosophers often make the mistake of treating intuition-gathering as sufficient on its own for complete clarification of a concept. They have a tendency to assume that one's "intuitive" verdicts on the acceptability of propositions involving a concept are both accurate and exhaustive as guides to the nature of that concept. Yet they have never made much of an effort to justify this assumption, and there are reasons to place it under suspicion. The best view of concepts is that they are a kind of abstraction from the

totality of a community's linguistic, and even non-linguistic, practices. For example, my concept of *up* is an abstraction from all the ways in which I use the word "up" and related words in sentences and inferences, and even from how I orient objects such as my body and a chair. Now, any particular person's gut-level acceptability judgments on a selected set of decontextualized propositions involving the concept are going to form *at best* a narrow window onto that concept. Further, in order to bring out what they take to be the most revealing intuitions, philosophers commonly contrive propositions describing unusual and even quite bizarre situations: planets that are duplicates of Earth in all respects except that water isn't really H₂O, neuroscientists brought up from birth in black and white environments, people conducting conversations by manipulating Chinese symbols according to giant rule books, and so forth. But what guarantee is there that intuitive judgments about such situations accurately limn the actual concept as it occurs in more mundane circumstances?

The second method of conceptual clarification is to pay close and systematic attention to the very practices from which the concept is an abstraction in the first place: the observations, sentences, inferences, and non-linguistic practices of those who deploy it. Of course, this is a rather demanding empirical project, requiring skills and patience of the kind possessed by field linguists, anthropologists and the like, and philosophers are no more likely to undertake it systematically than they are to go about gathering scientific evidence about the nature of reality in the first place. What they *are* often prepared to undertake is a modest linguistic version: studying the way particular words and phrases are standardly used in certain contexts. This process is known as linguistic analysis.

Consider, for example, the relatively novel cognitive science concept of *distributed representation*. What is it? The armchair intuitions of the average philosopher are basically worthless on this topic, not least because they are usually insufficiently familiar with the concept. The way to proceed in clarifying the concept is to pay close attention to how practicing cognitive scientists *actually* use the terms "distributed" and "distributed representation" in a myriad of *real* everyday situations. Performing this kind of linguistic study reveals a concept of representation that is multi-faceted, richly textured, and of considerable potential theoretical significance (van Gelder, 1991). It is highly unlikely that all this would have been apparent if one had been relying only one's "intuitive" verdicts on the acceptability of bizarre propositions.

Both forms of conceptual clarification discussed so far—consulting intuitions, and linguistic analysis—have in common the idea that there already is, somewhere in our minds or in our practices, an adequate concept just waiting to be clarified; the

only problem is to get a clear and explicit *grasp* of it. But what if this is not true? What if people happen to be just plain confused about what semantics, say, actually is? These two modes of conceptual clarification will, at best, just bring this out; at worst, they will perpetuate confusion by presenting partial perspectives masquerading as faithful clarifications. What they would not be able to do is provide the kind of straightforward conceptual order that we need if we want to, say, determine whether Searle's argument is valid. In such circumstances, conceptual clarification requires *reconstruction*. Under such circumstances getting clear on the concept means *making* the concept clear, forcibly restructuring it in whatever ways necessary. It is to engage in a *prescriptive* rather than merely descriptive activity. It is to prescribe that one should think about the topic a certain way, or use the corresponding words only in a certain regimented fashion.

For example, one might argue that, whatever else people might have been thinking, the only phenomenon of deep theoretical interest in the vicinity of semantics is that of *systematic interpretability*: if you can systematically, consistently and non-trivially interpret the symbols in some reasonably intricate structure as having a certain content, then they really do have that content. If the concept of semantics is clarified in this reconstructive way, then Searle's third premise is probably false: having (the right) syntax may well be sufficient for having semantics. since having (the right) syntax is what underwrites systematic interpretability.

1.3 Historical Perspective

Since the philosopher's primary method is argument, most of the philosopher's training involves developing argument skills. This is achieved by drilling the student on a selection of the best or most significant arguments available, a majority of which are drawn from the history of philosophy. Whereas psychology, computer science, linguistics, and neuroscience offer their students, at most, cursory sketches of the history of their concerns, philosophical training involves continual revisiting of past ideas and movements. The student comes to be quite familiar with the ideas and arguments of such past greats as Aristotle, Aquinas, Hume, Kant, Frege and Wittgenstein. Consequently, philosophers tend to be more thoroughly versed in the history of thinking on any given general topic than any of their counterparts in cognitive science. This learning translates into a distinctively philosophical way of approaching an issue: bringing to bear the accumulated insights and arguments of philosophical history.⁶

⁶ Note that, relative to many other disciplines in the humanities, Anglo-American analytic philosophers are usually, in certain ways, quite ahistorical in their

This *historical perspective* facilitates progress in at least two ways. First, it brings to the fore important previous arguments and counter-arguments bearing on the issue at hand. This can help save cognitive scientists from retracing old ground, something that still happens with alarming regularity. Second, it heightens appreciation of the dimensions and subtleties of the problem. This may not only help prevent pitfalls, it can suggest interesting new lines of thought.

For example, if the issue is whether minds are just machines of some sort, it is usually the philosopher who is expected to be familiar with what Great Minds of the Past have had to say on the matter. The doctrine of mechanism (the universe is just a complex machine, as are minds within it) has certainly had a long and honorable tradition in philosophy. Back in 1748, for example, Julien Offray de la Mettrie had been urging that

thought is so little incompatible with organized matter, that it seems to be one of its properties on a par with electricity, the faculty of motion, impenetrability, extension, etc. ((de la Mettrie, 1748/1927) p.64)

A century earlier, in fact, Thomas Hobbes had described thinking in a way that is now often taken to be the first encapsulation of the stance that Artificial Intelligence and mainstream cognitive science takes on the nature of all human cognition:

REASON...is nothing but *reckoning*, that is adding and subtracting, of the consequences of general names..." (Hobbes, 1651/1962) .

Relatively recently, the idea that thought processes are the orderly manipulation of internal symbols gained a large measure of credibility from Alan Turing's demonstration that some symbol manipulating machines are so powerful as to be what he called "universal" (i.e., able to be programmed so as to mimic any other discrete state machine) and his subsequent arguments that such a machine could behave intelligently (Turing, 1950). Turing's arguments lent support to mechanism because, if symbol manipulating machines can in fact be intelligent, our own intelligence might simply be the operation of such a machine, sometimes known as a "syntactic engine".

Now, we can see Searle as dampening all this enthusiasm by suggesting that if all you have is mechanistic symbol manipulation, a central aspect of mentality—namely, meaning—is being left out of the picture. But he was by no means the first to raise such a concern; Hobbes' contemporary Descartes had already

approach to problems. The claim that historical perspective is a distinctively philosophical approach to an issue in cognitive science is only true relative to the almost completely ahistorical character of the various other sub-disciplines.

seen this kind of problem. In a now famous passage from his *Discourse on Method*, he argued that a machine constructed to mimic a person

could never use words, or put together other signs, as we do in order to declare our thoughts to others. For we can certainly conceive of a machine so constructed that it utters words, and even utters words which correspond to bodily actions causing a change in its organs (e.g. if you touch it in one spot it asks what you want of it, if you touch it in another it cries out that you are hurting it, and so on). But it is not conceivable that such a machine should produce different arrangements of words so as to give an appropriately meaningful answer to whatever is said in its presence, as even the dullest of men can do... (Descartes, 1637/1988), p.44

No machine, in other words, could really converse in natural language. Why not? Because conversing means saying something appropriate to what was said before, and this appropriateness is a matter of the *meaning* of that previous utterance, and whatever meaning is, it is not something straightforwardly physical that has causal effects that can be built into the machine.⁷

John Haugeland has described this problem as the Paradox of Mechanical Reason (Haugeland, 1985), and has argued that twentieth-century philosophy and computer science have managed to come up with a certain kind of solution, one that requires that the machine be a *computer*—i.e., a machine which manipulates symbolic structures purely on the basis of their syntax, but in doing so manages to "respect" semantic properties. Searle, however, is not satisfied; he is arguing that *even if you can* build a computer that converses in natural language, there would *still* be a problem with meaning; namely, the internal symbolic "thoughts" of the machine would fail to have the kind of *intrinsic* meaning that is possessed by human thoughts. So Searle's argument can be seen as a contemporary way of reformulating Descartes' point that mechanical systems can't handle meaning, a reformulation in the light of our improved understanding of the nature of computational systems and the relation of meaning to symbolic representations. The Cartesian point now is that syntax alone doesn't get you what Searle thinks of as full-blooded semantics.

2. The Roles Philosophers Play.

Thus far philosophers have been characterized as distinguished by their typical methods, and that these methods are argument, conceptual clarification and historical perspective. But if this is right, what are philosophers actually qualified to

⁷ For illuminating philosophical discussions of this argument, see (Chomsky, 1990;

do? What roles do philosophers in fact play in cognitive science? This section briefly discusses some of these roles. This list is not meant to be definitive; philosophers contribute to cognitive science in many and diverse ways, and any reader familiar with cognitive science can probably think of interesting additions or alternatives. It is constructed with two major purposes in mind; first, to bring out some of this diversity, and second, to generate a sense of the importance of philosophical activity to the wider enterprise.

The Pioneer. It has been said, of philosophy generally, that its job is to struggle with those problems that nobody else quite knows how to address. This insight applies to cognitive science in a fairly direct way. Thus, suppose you had the general intellectual skills of the philosopher, but none of the particular skills of the psychologist, the linguist, or the computer scientist. What issues would you be best able to address? Clearly, it would not be specifically psychological, linguistic or computational issues. It would be issues that, for whatever reason, are not amenable to treatment by means of those other specialized disciplines. Now, an important subset of these issues are ones which can, with enough hard work, be *turned into* questions that can be addressed by specialized scientific techniques. An important role for philosophers, then, is to be the *pioneers*: to tackle problems that nobody else knows how to handle yet, in the hope of transforming them into scientifically tractable questions.

Historically, the role of philosophers as pioneers in cognitive science is indisputable. Virtually all the major topics of cognitive science—the nature of intelligence, of knowledge, perception, action, imagination, concepts, mental representation, the relationship of thought processes to the brain, etc. etc.—were first opened up and addressed by philosophers. For example, the existence and nature of "the language of thought" —a key theoretical postulate of mainstream cognitive science (Fodor, 1975)—was extensively debated in medieval times (Normore, 1990; Spade,). If there are now non-philosophical methods of tackling these issues, these methods generally developed out of the prior philosophical treatments and testify to the *success* of philosophers (whoever they may have been) in dealing with them.

Further, the most basic ideas around which contemporary cognitive science is constructed are all philosophical achievements. For example, the materialist metaphysical stance that is unthinkingly accepted by the vast majority of cognitive scientists—the idea that human thought processes are the operation of a certain kind of physical system—was first proposed, debated and refined by philosophers. The

foundational hypothesis of classical cognitive science—that thought is, more precisely, a form of symbolic computation—was, as mentioned above, a philosophical invention. And even the notion of computation itself was first carved out by philosophers. Cognitive science owes its existence, in part, to the pioneering efforts of philosophers such as Leibniz, Frege, Russell & Whitehead, and Turing, who are collectively responsible for the fact that we now understand quite well how the orderly manipulation of appropriately structured physical objects can, in a certain way, bridge the apparent gap between the physical world of causes and effects on one hand and semantic and even mental properties such as meaning, truth, and intelligence on the other (Pylyshyn, 1984).

Perhaps philosophers were pioneers in the *origin* of cognitive science, but don't play any such role *now*. Is there reason to believe that philosophers play an ongoing pioneering role? Yes; there are plenty of examples of issues within recent cognitive science that were first tackled by philosophers, and where those initial labors were subsequently turned to practical advantage by other cognitive scientists. Philosopher Clark Glymour has pointed out that many developments in machine learning in artificial intelligence have been a matter of adapting and implementing theories of scientific knowledge developed by philosophers (e.g., the Dendral and Meta-Dendral expert systems, which applied Vienna Circle philosopher of science Carl Hempel's theory of explanation and his instance-based approach to hypothesis confirmation). Another quite different example is the more recent upsurge of interest in, and progress on, issues of modularity in cognitive architecture (i.e., the extent to which the human cognitive system is made up of quite distinct modules which independently handle distinct cognitive tasks). This current attention to modularity is due in significant measure to Jerry Fodor's pioneering monograph *The Modularity of Mind* (Fodor, 1983). In that book he re-opened an important topic that had lain relatively dormant for many years and, in a fiery crucible, heated equal parts of argument, conceptual clarification, and historical perspective to produce a new strong modularity thesis that has become the primary point of orientation for all subsequent work in the area (see, e.g., (Shallice, 1988), (Farah, 1994)).

Further, there are plenty of important issues in cognitive science that are *still* such that nobody else has all that much to say about them yet, and so it is up to philosophers to try making such progress as they can. These include the *mind-body problem* (how, in general terms, is the relation between the mental and the physical best described?), the nature of *intentionality* (what does the meaningfulness, or "aboutness", or representational character, of mental states consist in?) and the problem of *consciousness* (what is it?). Philosophers have been busily working on

these issues while other cognitive scientists tackle more concretely specifiable problems. Of course, whether these philosophers' current work is pioneering in fact is a judgment that it will only be possible to make in retrospect, but there are some reasons for optimism; witness, for example, the favorable reception in wider cognitive science circles of Daniel Dennett's recent book on consciousness (Dennett, 1991).

The Building Inspector. Descartes once likened human knowledge to a house constructed on foundations of questionable integrity (Descartes, 1641/1986). Adapting this famous image to cognitive science, we can think of each specific scientific mode of enquiry as proceeding within the terms of a set of basic theoretical and methodological assumptions which constitute its foundations, and which ought sometimes be placed under scrutiny. The first stage is simply articulating the basic assumptions themselves; the second is inspecting them for various kinds of defects (vagueness or incoherence, inconsistency with each other, outright falsity). The third and most difficult stage is to reconstruct the foundations, if necessary. Now, since these foundations are what makes everyday kinds of scientific inquiries possible, inspecting foundations cannot be part of those everyday inquiries; it is a very different kind of intellectual activity. The tools of the philosopher—argument, conceptual analysis, and historical perspective—turn out to be very useful here (though they must be supplemented with the kind of subtle *insight* that only comes from deep familiarity with the area itself). Consequently, anyone engaged in inspecting foundations must be, if not a philosopher already, at least temporarily setting aside normal scientific methods and pursuits in favor of philosophical reflection; and one of the most common roles that philosophers actually do play is that of inspecting foundations.⁸ Extending Descartes' image, we can think of philosophers as the *building inspectors* of cognitive science.

For example, cognitive psychologists typically assume (a) that the performance of individual human subjects on cognitive tasks such as memorizing is the direct result of the way relevant information is mentally represented and manipulated; (b) that subjects are all basically alike in their mental organization *vis a vis* such tasks; and (c) that, consequently, one can infer the general structure of that

⁸ Symptomatic of this is the fact that so many articles or books in the philosophy of cognitive science use the term "foundation" in their title: e.g., Jay Garfield's introduction to the philosophy of cognitive science "Philosophy: Foundations of Cognitive Science" (Garfield, 1987), and Zenon Pylyshyn's classic book *Computation and Cognition: Toward a Foundation for Cognitive Science* (Pylyshyn, 1984).

mental organization on the basis of data on performance gained by averaging over trials on many individual subjects. Assumptions like these form part of the foundations of contemporary cognitive psychology; specific hypotheses and experiments within the discipline are constructed upon them, and psychologists standardly pay them little attention as they go about their business. However, it is not *obvious* that such assumptions are correct or unproblematic, and they deserve to be carefully scrutinized. To some extent the proof will be in the pudding; the value of the assumptions will be borne out by the empirical success or failure of actual work in experimental cognitive psychology. But for a number of reasons we cannot rely solely on this empirical measure. On one hand, those who accept a certain set of foundational assumptions can always explain away an apparent lack of empirical success by claiming that more research time and effort is still needed. On the other hand, a mode of scientific enquiry can *appear* quite fruitful in its own terms, while hindsight reveals that this fruitfulness was largely an illusion sustained by unquestioned acceptance of the very assumptions at issue. And even if some mode of scientific enquiry does grind to a halt, indicating severe problems with its foundational assumptions, this tells us only *that* something was wrong with those assumptions, and not *what* was wrong. Philosophical reflection is required in order to understand and eliminate the problem.

Thus, by the nineteen fifties it was increasingly apparent to many psychologists that the basic assumptions of behaviorism rendered experimental psychology a largely sterile and futile business, in most areas at least. Yet it took careful philosophical investigation to articulate those assumptions, demonstrate their inadequacy, and propose new principles that might guide more fruitful work. This is exactly what Noam Chomsky did in his famous critique of Skinner's book *Verbal Behavior* (Chomsky, 1957; Skinner, 1957). Though Chomsky is best known in cognitive science circles as a linguist, that review was not primarily a piece of linguistics; it was, rather, one of the classics of the philosophy of cognitive science. Likewise, the foundational assumptions of contemporary cognitive psychology need to be articulated, clarified, and evaluated; and although this might best be done by someone who is normally a cognitive psychologist, that person will be temporarily be stepping outside everyday activity in cognitive psychology to engage in some philosophizing.

The Zen Monk. The fact that philosophers can bring some *general* intellectual skills to issues that nobody else can handle is no guarantee that the philosophers will make any progress. What often happens instead is that they get caught up in arcane and

irresolvable debates that, as time goes on, become increasingly removed from everyday work in cognitive science. The philosopher becomes like the Zen monk—a figure supported by the community to ponder those imponderable issues that everyone thinks should be thought about by someone, but for which nobody else has time or patience. In theory, the philosopher-monk eventually reaches a state of enlightenment, but unfortunately that enlightenment is necessarily incommunicable to those who have not undertaken the requisite prolonged course of meditation and asceticism.

Consider perception. Suppose you and I are lost in a game park somewhere in Costa Rica and simultaneously catch sight of an approaching and obviously hungry Tyrannosaurus Rex. Somehow or other we both come to think *Oh no! I'm about to be eaten by a T-Rex!* and we both run for cover. Now, there are thousands of questions that cognitive scientists might ask about these events. How did our visual systems discriminate the dinosaur in front of us? How were we able to categorize the object as a T-Rex? How did we infer that we were about to be eaten, or that we should flee? How was the desire to flee translated into appropriate limb motions? And so on. All of these questions are naturally suited to various of the specialized techniques of cognitive science. The operation of our visual systems in discriminating the dinosaur from the trees, for example, is a question for the computational neuroscientists. But philosophers have come up with their own quite different questions. They notice that when you and I see the T- Rex, even though in some sense the same thing is going through our minds, we actually have quite different thoughts: I think that *I* am about to be eaten, and you think that *you* are about to be eaten. They notice, moreover, that we would be having different thoughts even if our cognitive systems were in all other relevant respects identical; indeed, we would be having different thoughts even if we were molecule-for-molecule duplicates of each other! Even if all the other cognitive scientists had completely finished answering all their questions about how our cognitive systems work, there would still be this question left over: how is it—and how does it matter—that you and I are really thinking different things when, in one obvious sense, we're both thinking the same thing?

The first step towards philosophical enlightenment here is to see that the problem is one of thought *content*: our thoughts differ not in their causal roles in our cognitive systems but in their *meaning*. The next step is to distinguish two kinds of meaning - what philosophers call narrow content and wide content. Narrow content is that aspect of the meaning of our thoughts which the two of us share and accounts for the idea that we are thinking the same thing. Wide content is that aspect of

meaning which goes beyond narrow content and accounts for the idea that we are thinking different things. Wide content is usually thought to be determined by narrow content *plus* something from "outside" the system: thus the wide content of my thought that *I am about to be eaten by a T-Rex* is its narrow content plus the fact that it is *me* that is thinking this thought rather than you.⁹ And it is often thought that cognitive scientists need to take narrow content into account, since that's the notion of meaning that is relevant to understanding how my inner mental representations manage to get me to do what I do, but they don't need to worry about wide content, since that depends on factors "outside" the cognitive system and hence cannot be relevant to the causal operation of the system (Fodor, 1980).

So far so good. Philosophers then rightly go on to explore these issues in greater depth. What exactly are narrow content and wide content? How are they related? What kind of content should matter to the cognitive scientist—either or both or neither? and so on. However, the discussion quickly becomes rather obscure, and every conceivable position seems to have been taken up and defended by one philosopher or another. It is increasingly difficult to tell what real progress philosophers have made in this area. In the meantime, the rest of cognitive science goes about its business casually ignoring the distinction between wide and narrow content. Somehow it just doesn't seem to make a difference, though everyone is agreed that someone should be worrying about these things. That's one reason cognitive science takes philosophers on board.

The Cartographer. The philosopher Wilfred Sellars once characterized philosophy as aiming "to understand how things, in the broadest possible sense of the term, hang together, in the broadest possible sense of the term." Applied to the philosophy of cognitive science, this means that one role of philosophers is understanding and describing how all the various elements of cognitive science fit together (or conflict, as the case may be). In doing this they produce large-scale conceptual maps of the discipline; we can thus think of philosophers as the *cartographers* of cognitive science.

Again, it follows straightforwardly from the characterization in Section 1 that cartography will be up to the philosopher. Of course, someone who is normally, say,

⁹ It does, admittedly, sound a little strange to say that the fact that it is *me* rather than *you* thinking a certain thought is a fact that is "external" to me. Clearly, however, it is not something *internal*; you and I might be, internally, absolutely identical, right down to the molecular structure—i.e., no internal differences at all—and we would still be having different thoughts. Generally speaking, philosophers have found it most appropriate to treat this kind of "indexical" fact as just another feature of the external context.

a psychologist, can also be a cartographer, but in doing so she will not be using any of the distinctive skills of the psychologist, such as experimental design, statistical analysis, etc.; these skills are obviously irrelevant to the task. She will have to fall back on general skills such as argument and conceptual clarification, i.e., on the methods of the philosopher.

In its own modest way, this chapter is itself an exercise in cartography. But for better examples, one need only look at many of the classics of the philosophy of cognitive science literature. In the last ten to fifteen years numerous books have been written on Artificial Intelligence. Most of these have been technical introductions to AI programming, written by and for computer scientists. The philosopher John Haugeland's book *Artificial Intelligence: The Very Idea* (Haugeland, 1985), by contrast, aims at an understanding of AI in terms of its wider historical and conceptual context. He sets out by sketching the history of certain key ideas and problems, from Copernicus and Galileo through Descartes, Hobbes and Hume. He then spends two whole chapters just articulating what a computer *is*, in terms of concepts such as formal systems, digitalness, medium independence, algorithms, automation, meaning and interpretation. An overview of different kinds of computer architecture leads into a conceptual survey of various AI research programs: machine translation, heuristic search, microworlds, frames, and an assortment of important potential difficulties: the frame problem, pragmatic sense, imagery, feelings, ego involvement. The point here is the kind of understanding of AI the reader is given. It is not the sort of inside, hands-on knowledge of the technical practitioner; rather, it is an understanding of what AI is in terms of how it *fits* into a larger framework of ideas and philosophical ambitions.

Drawing up a good map requires knowing the various regions that it covers pretty well, and when the map is supposed to cover some significant chunk of a domain as vast and heterogeneous as cognitive science, this is a pretty tall order. By practical necessity, philosopher-cartographers almost always do not know the particular topics that they are trying to fit into the map as well as the specialists on that topic, and as a result the specialists often regard the cartographer as superficial or misinformed. (In the worst case, *all* relevant specialists will take this attitude; the philosopher becomes something of a *dilettante*, purporting to know about many things but having no really deep familiarity with any of them.) However, the appropriate response to this inevitable problem is not to stop drawing up maps altogether, but to continue striving to draw up better maps, which means cooperating with all the specialists in order to develop the best way of describing their particular region and how it fits with other parts of cognitive science.

The Archivist. Closely related to the role of cartographer is that of *archivist*. It is the philosopher who, more than anyone else in cognitive science, is expected to be the repository of accumulated wisdom: to have understood, seen the significance of, and remembered ideas and research programs of the past both good and bad. That philosophers in particular would play this role is just what you'd expect if, as suggested in Section 1, a distinctively philosophical way of approaching an issue in cognitive science is that of historical perspective. After all, much of that accumulated wisdom is philosophical work itself, and that which isn't the philosopher should be aware of in his role as cartographer anyway.

An excellent example of philosophers playing this archival role (among others, of course) to great effect is the devastating critique of AI attempts at natural language understanding by philosopher Hubert Dreyfus . Human processing of natural language is a complex business. At one end is "low-level" perceptual processing of auditory or visual stimuli; at the other is the fine motor control involved in producing an utterance or writing a sentence; and in between are a range of cognitive tasks such as parsing, knowledge retrieval and inference. Nobody can yet hope to produce a single unified account of how this entire complex operation takes place; cognitive scientists must partition the problem into manageable aspects. Artificial intelligence researchers have attempted to distill out and solve one particular problem of natural language *understanding*: given a few sentences of English (i.e., strings of words, not acoustical signals), the computer must figure out what those sentences are saying. The criterion of success is being able to answer simple questions about what was said. Thus, in a typical example, the computer is given the sentences

Mary saw a dog in the window. She wanted it.

and it might be asked to answer the question

What did Mary want?

This may sound trivial, but AI researchers quickly found out that answering such questions depends on possessing and bringing to bear just the right commonsense background knowledge. We all know that people rarely want windows when they see them, but may well be gripped by a desire to have a cute dog as a pet. Since understanding natural language means being able to handle not just these particular sentences but any that might come along, the general capacity requires having a vast amount of knowledge of this kind at one's fingertips. Any AI natural language understanding system that performs as well as the average human child must be given access to this same store of knowledge. At least some AI researchers have

consequently assumed that the success of AI hinges on being able to explicitly program all of ordinary human commonsense knowledge into a computer, and have set about trying to do this (Lenat & Feigenbaum, 1991).

Dreyfus's *What Computers Can't Do* is a profound book-length critique of this and other AI ambitions, and justice could not possibly be done to its arguments here. The point of immediate interest is just the way in which Dreyfus' critique is crucially based on his familiarity with past intellectual developments—in other words, on his playing the role of archivist for cognitive science. This is one of the key respects in which Dreyfus' philosophical treatment of the AI problem of natural language understanding differs from that of the computer scientists, most of whom seem to be unaware that there were any important ideas before about 1960.

Dreyfus' critique proceeds in two stages. In the first he situates the ideas underlying AI as the culmination of a certain tradition in Western philosophical thought:

GOFAI¹⁰ is based on the Cartesian idea that all understanding consists in forming and using appropriate symbolic representations. For Descartes, these representations were complex descriptions built up out of primitive ideas or elements. Kant added the important idea that all concepts are rules for relating such elements, and Frege showed that rules could be formalized so that they could be manipulated without intuition or interpretation. Given the nature of computers as possible formal symbol processors, AI turned this rationalist vision into a research program and took up the search for the primitives and formal rules that captured everyday knowledge... (Dreyfus, 1992) p.xi.

Dreyfus also knew that a number of twentieth-century philosophers had already presented powerful criticisms of aspects of this tradition. In the second stage, then, he used the insights of these philosophers to argue that AI could not hope to succeed. Dreyfus adopted from Heidegger, Merleau-Ponty, and the later Wittgenstein a series of arguments and phenomenological descriptions¹¹ which indicated that it is neither possible nor necessary to produce a complete formal specification of the commonsense knowledge which underlies our ability to understand language and behave intelligently in general. That background

¹⁰ "Good Old Fashioned Artificial Intelligence" (Haugeland, 1985).

¹¹ Here "phenomenological description" does not mean (as it sometimes does elsewhere) describing one's own introspectively accessible ideas, but rather the careful study and description of the way everyday things ordinarily are, including especially things that are usually somehow hidden from view, such as the way

"knowledge" is primarily a matter of non-represented skills, experiences and imaginative capacities that depend crucially on one's human embodiment and ongoing engagement in physical and social situations. While any particular aspect of this background knowledge might be made explicit in the form of a symbolically represented fact or rule, it is impossible in principle to make it all explicit and usable by a computer since the usability of any given piece of explicit knowledge (e.g., realizing that it is relevant to a given problem) depends on the intactness of the non-formalized background. Consequently, AI can never succeed in its goal of programming a computer with commonsense knowledge. The human mind in fact works in an entirely different way.

Dreyfus's arguments are subtle and controversial, but they gain some measure of support from the fact that over the last few decades AI seems to have been encountering exactly the kind of difficulties that Dreyfus predicted on the basis of these philosophical considerations. Now, there would have been nothing to prevent AI researchers from reading up on their Heidegger and Merleau-Ponty before setting out to write their programs, except that that is not the kind of thing computer scientists typically do. It happens to be the philosopher's role in cognitive science; the important thing is just that other cognitive scientists be prepared to benefit.

The Cheerleader. As the official repositories of accumulated wisdom, philosophers have acquired a certain measure of both authority and responsibility in determining what counts as a Good Idea. Consequently, when other cognitive scientists believe that their particular line of research has turned out to be Important, they like to have philosophers publicly bestow their official Seal of Approval, and even better, to explain to the world just how Significant that line of research really is. Of course, any philosopher who does see the Significance of that research must be quite perceptive. A pleasantly symbiotic relationship thus builds up between philosophers and cognitive scientists, with each pronouncing on the goodness of the other's work.

This can be put by saying that if the researchers are the players out there on the field of science, philosophers are sometimes the *cheerleaders*. This is not the place to discuss specific examples; suffice to say that while most major movements in cognitive science have had their share of philosophical cheerleaders, some have done particularly well; these include artificial intelligence, connectionism and computational neuroscience. On the other hand, it may be that some schools of thought, such as ecological psychology, have ended up less prominent in cognitive

your tennis racquet is for you as you concentrate on playing the game.

science than they might otherwise have been due to unfortunate deficiencies in this key area.

The Gadfly. If Oscar Wilde had been in the philosophy of cognitive science, he might have said that the only thing worse than being refuted is being ignored. Thus philosophers often advance positions that are so strongly and provocatively stated that other cognitive scientists feel compelled to respond. This is the philosopher as gadfly. An excellent example is the notorious paper on connectionism by Fodor & Pylyshyn (Fodor & Pylyshyn, 1988). The authors there stated bluntly, and argued deftly, that a whole new research program was either hopeless or an uninteresting variant on their own classical world-view. Cognition, they claimed, is *systematic*, and connectionist models either fail to explain this, or they must be implementing some computational cognitive architecture of the kind Fodor & Pylyshyn had been recommending for years. Many connectionist psychologists and computer scientists found these charges so galling that they couldn't resist taking up the challenge of refuting the argument. In the process, Fodor & Pylyshyn's paper became one of the most discussed in the philosophy of cognitive science, and a whole new research niche was created.

In fact, the Fodor & Pylyshyn paper can be seen as exemplifying almost all the roles discussed so far. It was pioneering, in bringing to the forefront of attention a previously neglected phenomenon, that of systematicity. It was inspecting the foundational assumptions of whole research paradigms in cognitive science. It was cartographic, in that it began with a high-level sketch of the lay of the land in issues of cognitive architecture. It was archival, in that it likened connectionism to the seventeenth-century empiricist David Hume and twentieth century behaviorists such as Hebb, Osgood and Hull, and claimed to be merely re-presenting thirty-year old arguments against such views. It was highly enthusiastic in its support for one school of thought over another. To its credit, however, the authors were not playing the role of Zen Monk: the arguments it contained were clear, powerful, and highly relevant to other branches of cognitive science.

References

- Bechtel, W. (1988) *Philosophy of Mind: An Overview for Cognitive Science*. Hillsdale N.J.: L. Erlbaum Associates.
- Chomsky, N. (1957) Review of B.F. Skinner's *Verbal Behavior*. In J. A. Fodor & J. J. Katz ed., *The Structure of Language*. Englewood Cliffs N.J.: Prentice-Hall,
- Chomsky, N. (1990) Linguistics and Descartes. In J.-C. Smith ed., *Historical Foundations of Cognitive Science*. Dordrecht: Kluwer, 71-79.

- Churchland, P. M. (1988) *Matter and Consciousness: A Contemporary Introduction to the Philosophy of Mind* (Revised ed.). Cambridge MA: MIT Press.
- de la Mettrie, J. O. (1748/1927) *Man a Machine*. Chicago: Open Court.
- Dennett, D. C. (1991) *Consciousness Explained*. Little-Brown.
- Descartes, R. (1637/1988) *Discourse on the Method of Rightly Conducting One's Reason and Seeking the Truth in the Sciences*. In *Descartes: Selected Philosophical Writings*. Cambridge: Cambridge University Press, 20-56.
- Descartes, R. (1641/1986) *Meditations on First Philosophy* (Cottingham, John, Trans.). Cambridge: Cambridge University Press.
- Dreyfus, H. L. (1992) *What Computers **Still** Can't Do: A Critique of Artificial Reason*. Cambridge MA: The MIT Press.
- Farah, M. (1994) Neuropsychological inference with an interactive brain. *Behavioral and Brain Sciences*, **17**, 43-61.
- Fodor, J. A. (1975) *The Language of Thought*. Cambridge MA: Harvard University Press.
- Fodor, J. A. (1980) Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences*, **3**, 63-73.
- Fodor, J. A. (1983) *The Modularity of Mind*. Cambridge MA: Bradford/MIT Press.
- Fodor, J. A., & Pylyshyn, Z. (1988) Connectionism and cognitive architecture: a critical analysis. *Cognition*, **28**, 3-71.
- Garfield, J. (1987) *Philosophy: Foundations of Cognitive Science*. In N. A. Stillings et al ed., *Cognitive Science: An Introduction*. Cambridge MA: MIT Press, 305-361.
- Gunderson, K. (1964) Descartes, La Mettrie, language and machines. *Philosophy*, **39**, 193-222.
- Harman, G. (1989) Some philosophical issues in cognitive science: Qualia, intentionality, and the mind-body problem. In M. A. Posner ed., *Foundations of Cognitive Science*. Cambridge MA: Bradford/MIT Press, 831-848.
- Haugeland, J. (1985) *Artificial Intelligence: The Very Idea*. Cambridge MA: MIT Press.
- Hobbes, T. (1651/1962) *Leviathan*. New York: Collier Books.
- Lenat, D., & Feigenbaum, E. (1991) On the thresholds of knowledge. *Artificial Intelligence*, **47**.
- Normore, C. G. (1990) Ockham on mental language. In J.-C. Smith ed., *Historical Foundations of Cognitive Science*. Dordrecht: Kluwer Academic Publishers, 53-70.
- Pylyshyn, Z. W. (1984) *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge MA: Bradford/MIT Press.
- Searle, J. R. (1990) Is the brain's mind a computer program? *Scientific American*, 20-25.

- Shallice, T. (1988) *From Neuropsychology to Mental Structure*. Cambridge: Cambridge University Press.
- Skinner, B. F. (1957) *Verbal Behavior*. New York: Appleton-Century-Croft.
- Spade, P. V. (1980) Synonymy and equivocation in Ockham's mental language. *Journal of the History of Philosophy*, **18**, 9-22.
- Turing, A. (1950) Computing machinery and intelligence. *Mind*, **59**, 433-460.
- van Gelder, T. J. (1991) What is the 'D' in 'PDP'? An overview of the concept of distribution. In S. Stich, D. Rumelhart, & W. Ramsey ed., *Philosophy and Connectionist Theory*. Hillsdale NJ: Lawrence Erlbaum Associates, 33-59.