8 Meaning Making and the Mind of the Externalist

Robert A. Wilson

1 Intentionality and the Mind

During the 1980s, many philosophers of mind, and even the occasional cognitive scientist, were very exercised about something called "the problem of intentionality." The problem was something like this. There are certain things in the world that appear to possess, through their operation and functioning, a special kind of property: intentionality. This is the property of being about something, of having content about that thing, of carrying information about that thing. The problem of intentionality was threefold: to explain what intentionality was; to delineate which things had intentionality (and so which things didn't); and to provide an account of just why they had not only intentionality, but the particular intentionality they had—their *content*. The third of these chores was the core one, the task of specifying in virtue of what certain things in the world were about the particular things they were about.

The problem of intentionality was especially pressing within the naturalistic view of the mind that motivated much of the discussion of the problem. The idea was to view naturalism as a kind of constraint on what could count as an acceptable endeavor to complete the core chore: that one's account of what made for intentionality could not itself rely on unexplained intentional or semantic notions. An answer to the problem of intentionality must be given solely in terms of "naturalistically acceptable" notions, such as causation, counterfactual dependence, material composition, biological function, or phylogenetic history.

To understand a little more about the problem of intentionality, we need to turn to its second part, the part that divides the world into things with intentionality and things without it. Two of the things that paradigmatically have intentionality are the language that people use to communicate,

and the thoughts that people have (often, but not only, while participating in those communicative acts). Of these, it is plausible to hold that the intentionality of language derives from that of thinking. That is because people mean something by making an utterance because that utterance conveys a thought (belief, desire, expectation, fear, memory, etc.) about that very thing. When I say to you that the kangaroo paws in King's Park are flowering, that's in part because I believe that very thing, having ridden by them earlier in the day. But more to the point, the utterance means what it does in part because of the meaning that the mental representations that led me to utter it have. The meaning of an utterance is "derived" from the meaning of an act of thought. Thought, however, has *original* intentionality in that the meaning that its constituent mental representations have does not in turn derive from the intentionality of anything else.

If that is true of the intentionality of spoken language and the mind, then perhaps it is true more generally. There are other things that have some kind of intentionality—road signs, gestures, written language, computer programs—and if the intentionality of speech derives from that of the mental representations that generate it, then it doesn't seem too much of a stretch to think that these too have much the same kind of derived intentionality. In short, underlying all of intentionality, perhaps, is the intentionality of the mind. This, together with the naturalistic constraint, created a certain kind of research program, one that aimed to provide an account of the intentionality of mind in non-intentional terms. Hence the rise of informational semantics (Dretske 1981, 1988), causal theories of representation more generally (Fodor 1984, 1987), and biosemantics (Millikan 1984, 1993)—all primarily accounts of the intentionality of mental representations and minds, and all putatively kosher from a naturalistic point of view.

There are many nuances to the problem of intentionality that this sketch ignores, giving rise to a variety of misgivings that could be expressed about proceeding without further elaboration. My main purpose in beginning with this thumbnail, however, is simply to say that there is one thing that this literature has been right to seize on, and one thing that it presupposed that, in hindsight, is more dubious. First, the Good News.

Whether or not one agrees with the reductive drift underlying the treatment of cognition as asymmetrically primitive when it comes to intentional phenomena, and whether or not one views the naturalistic constraint as, in effect, making the problem of intentionality unsolvable, the mind really is special when it comes to matters meaningful. Mental representation of some kind or other is involved in some way or other in all other

cases of intentionality. Language, conventions, social practices, images, causal chains of dependence, information transmission, biofunctions, historical anchoring, resemblance, and systematicity—all of which have been proposed as contributing or critical factors in responses to the problem of intentionality—have to pass the acid test of accounting for mental representation in order to do so. This is sufficient to support the point I am making here. But I also think that the mind is presupposed in the operation or application of each of these to the problem of intentionality.

One might wonder in just what sense mental representation is presupposed in positing, say, that the heart has the (bio)function of pumping blood, or in proposing that there is an informational relation between smoke and fire. Surely, one might think, both of these turn solely on whether the corresponding naturalistic grounding relations exist. In the case of the biofunction of pumping blood, these are historical facts concerning what ancestral hearts did, and in that of smoke and fire, the kind of causal dependence between the two. Yet at best, in the complexities of the actual world, natural selection is not fine-grained enough to distinguish the pumping of blood from myriad properties with which it is entwined (*sensu* Wilson 2003b). And similarly, the robust clustering of properties means that "carries information about" will be, at best, a one–many relation. What is needed, in both cases, is ultimately something mindful—either an individual mind itself, able to select and discriminate among the property clusters, or a convention, habit, or agreed-on practice that, in turn, presupposes individual minds that are able to do so.

In short, the mind is crucially implicated in the making of meaning, both directly (in nearly all cases) and indirectly (in the remainder). That is the Good News about how the problem of intentionality has been construed. By homing in on mental meaning, it has hit Intentionality Central.

The Bad News is that the task of solving the problem of intentionality has been complicated by an innocuous-enough sounding assumption, namely, that mental states are "in the head." In fact, they often are not. To see why, we need to remind ourselves of how the debate between individualists and externalists about cognition has developed over the past thirty years.

2   Externalism about Cognition and the Extended Mind

The predominant view of cognition throughout the short history of cognitive science and the longer history of the philosophy of mind has been *individualistic* in that cognitive processes have been understood by

abstracting away from the physical and social environments in which cognition takes place. That is, it is not simply that cognition is a property of individuals (see Zerubavel 1997; cf. Barnier et al. 2008), but that it is construed and investigated as if that individual were the only thing that existed in the world. This "methodological solipsism" (Fodor 1980) about cognition has been defended as a consequence of accepting a computational theory of cognitive processing (Devitt 1990; Egan 1995), and by appealing to broader claims about the nature of scientific taxonomy (Fodor 1987) and the metaphysics of mental causation (Crane 1991). Acceptance of the constraint of individualism has gone hand in hand with a view of the mind as possessing a rich, internal structure, whether that structure be innately hardwired or largely acquired through training or regimentation of some kind; modular or generalist in its functional decomposition; or rules-and-representational or associationist in its operational dynamics.

Those who reject individualism are *externalists* about the mind, and they have both responded to the arguments for individualism (Wilson 1992, 1995, part I) and, more recently, constructed positive visions of the study of cognition based on the rejection of individualism (Clark 2003; Rowlands 1999; Wilson 2003a, 2004). Those visions have taken the idea of externalism in more radical directions in recent years (e.g., Menary 2007; Wilson and Clark 2008; Wilson and Craver 2007; Chemero and Silberstein 2008; Clark 2008). The "radicalness" of these recent directions can perhaps be seen most clearly by reflecting briefly on how externalism began.

Doubts about individualism were first raised in the now classic arguments of Putnam (1975) and Burge (1979), both turning on the question of whether individualistic views of the mind could adequately account for meaning or mental content. These original challenges to individualism were cast in terms of whether psychological states, particularly intentional states, should be individuated or taxonomized in accord with the constraint of individualism. For this reason I have elsewhere (Wilson 2000, 2004) called the resulting forms of externalism *taxonomic externalism*; it has also been called "traditional externalism" and "philosophical externalism." The general idea was to continue to view content-laden mental states as some kind of internal state of the individual, but to argue that, nonetheless, because of their content, they did not supervene on, or were not metaphysically determined by, what fell within the physical boundary of that individual. Hence the debates over whether folk psychology was individualistic (Burge 1979; Fodor 1982; Loar 1988), whether the notion of

content used in Marr's celebrated theory of vision was internalist or externalist (Burge 1986a; Egan 1992; Segal 1989; Shapiro 1997), and the issue of the relationship between individualism and the normativity of the mental (Burge 1986b; Ebbs 1998; Pettit 1993).

More radical forms of externalism about the mind—what I (Wilson 2000, 2004) call *locational externalism*, Rowlands (1999) calls *environmentalism*, and, following Clark and Chalmers (1998 and this volume), many call the *extended mind thesis*—hold that the mind and the cognitive processes that constitute it extend beyond the boundary of the skin of the individual agent. The extended mind thesis identifies cognitive systems themselves as reaching beyond individuals into their physical and social environments. The thesis challenges individualism directly by implying that an individualistic psychology could only, at best, tell part of the story about cognitive processing: the inside story. An early gesture at such a view was my *wide computationalism* (Wilson 1994), the view that the computational systems that make up the mind can extend into, and include as a part of themselves, aspects of an organism's environment.

Locational externalism, environmentalism, and the extended mind thesis are radical forms of externalism in at least two ways. First, they do not rest on claims and intuitions about whether the content of a pair of states of two individuals in different environments (or one individual in two such environments over time) is the same or different, about how particular intentional states are taxonomized, or about the role of the physical or social environments in individuating such states. Instead, they appeal to the nature of psychological processing, to the arbitrariness of the head (or the skin) for bounding cognitive systems, and to what happens in real-life, online cognitive activity in the world. Thus, if the extended mind thesis is true, it is true in virtue of something implementationally deep about cognition, rather than some debatable view of mental content. Second, locational externalism is not simply a view of how we "talk about" or view cognition and the mind—about the epistemology of the mind, one might say—but about what cognition and the mind *are*—about the ontology of the mind.

If the mind is not, literally, "in the head" (and for resistance to radical externalism, see Adams and Aizawa 2008, this volume; cf. also Aizawa 2007), then maybe we need to reconceptualize the many issues that turn on the assumption that it is—such as the problem of intentionality. But it will take us some time to return to this point, for we first need to probe the idea of extended cognition, and arguments for accepting that large chunks of cognition are extended, more thoroughly.

3   On Arguments for Extended Cognition

As externalism has articulated and developed its own positive vision for
the study of the mind, it has also drawn, often implicitly, on arguments
that are not simply of the form "Arguments for my opponent's position are
no good." The argument I shall focus on concerns intentionality, although
it does not rely on intuitions about Twin Earth kinds of thought experi-
ments, as did the early arguments of Putnam and Burge. It belongs to a
family of arguments that I shall call *active cognition* arguments.

Active cognition arguments are so called because they all appeal to the
active exercise of cognitive capacities in the real world (see Clark 1997;
Haugeland 1998; Hurley 1998, 2001; Rowlands 1999, chaps. 5–8; Wilson
2004, chaps. 8–9). Active cognition arguments have appealed especially to a
range of work in perceptual psychology, ranging from Gibsonian approaches
(Gibson 1966; see also Rowlands 1999) through to Dana Ballard's animate
vision paradigm within computational psychology (Ballard 1991; Ballard et
al. 1997; see also Wilson 2004, chap. 7) and the more recent view of visual
experience as involving the animated exploration of one's environment and
in so doing displaying a sensitivity to sensorimotor contingencies in the
environment of the organism (O'Regan and Noë 2001). As one might expect,
given the diversity in this work, active cognition arguments for the extended
mind have also appealed to other areas of psychology: to problem solving,
in cases where this involves exploiting the visual or tactile features of one's
environment to complete the problem-solving task (Clark 1997, 2003; Wil-
son 2004, chap. 8); to actual and possible cases involving sensory inversion
(Hurley 1998, 2001); and to the mediational approach to cognitive develop-
ment pioneered by Vygotsky and Luria, championed in the contemporary
literature by Michael Cole (1996) and James Wertsch (1998, 2002).

These arguments all focus on determinate forms of a particular cogni-
tive ability (e.g., memory, attention, problem solving) as they are exercised
by individual agents. They view the integration of individuals with both
their biological and artificial environments as critical to their status as
cognitive agents with these particular capacities. With this focus on actual
agents and the abilities they act on, active cognition arguments try to pre-
empt the objection that "the extended mind" is merely a conceptual pos-
sibility or a *façon de parler*. The chief aim of active cognition arguments has
been to show directly that much of cognition as we know it is extended;
the real question for their proponents is just which aspects of cognition
are extended, and in what ways.

Active cognition arguments contrast, in these respects, with another kind of argument for the extended mind, due chiefly to Andy Clark, which I shall call *cyborg fantasy* arguments. These arguments, like the classic arguments of Putnam and Burge for taxonomic externalism, turn primarily on intuitions about imagined cases ("intuition pumps," as they are called). The best-known of these—from Clark and Chalmers 1998 (this volume, chap. 2), and discussed in several papers in the current volume, the case of Otto and Inga—focuses on a person's (Otto's) reliance on a notebook as a memory-storage device. Otto is a kind of cyborg, even if, in the basic case at least, his cyber-enhancement relies on an old form of technology—the book. Some versions of cyborg fantasy arguments (e.g., many of those in Clark 2003) imagine amplifications of dependencies on existing technologies—on cell phones, electronic implants, and telerobotics, for example—to argue that future cognition will likely be more radically extended than it currently is, with implications for the headier notions of the self, identity, and bodily integrity.

Cyborg fantasy arguments for the extended mind proceed by introducing an imaginative example in which an individual's cognitive performance is mediated by external forms of technology, typically arguing, through a comparison to cases in which the same kind of activity is performed without such mediation, to the conclusion that the boundary between what's inside the head and what is in the environment is irrelevant to whether a given agent has some particular cognitive capacity. The focus in cyborg fantasy arguments is on imagined cases (albeit ones that sometimes are or can seem close to actual cases), and there is an emphasis on techno-facilitation, rather than the individual's integration with her natural environment. The chief aim of cyborg fantasy arguments has been to establish the extended mind as a kind of conceptual default; they do so by shifting the burden of proof to internalists, challenging them to identify why the skin should be a relevant boundary for cognition at all.

Common to both active cognition and cyborg fantasy arguments for extended cognition is the idea that by examining just what is involved in the exercise of some particular cognitive capacity, one finds that it actually does or could well involve causal loops that extend beyond the body of the individual agent. In particular, these causal loops (do or may) pass through objects and other entities in the agent's environment, and it is only the whole, functioning, beyond-the-head causal system that constitutes the matter in motion that realizes the exercise of the capacity. In this system, some of the cognitive resources—what Hurley (1998) and Rowlands

(2003) would call cognitive *vehicles*—lie inside, and others outside the head. Just as not all of the resources used to build an organism's *developmental resources* are genetic, and not all of these nongenetic developmental resources are located inside the organism's existing boundary, so too with cognitive resources: not all of these are neural, and some of those that are not are located beyond the head of the individual.

Both active cognition and cyborg fantasy arguments have recently been attacked by critics of the extended mind thesis, with these critiques drawing broader conclusions about the tenability of radical forms of externalism on the basis of putative failings of particular arguments: Adams and Aizawa (2001) on Clark and Chalmers (1998), Grush (2003) on Haugeland (1998), and, most recently, Rupert (2004) largely on Rowlands (1999). For the most part, these critiques have to reconstruct, sometimes quite imaginatively, the arguments that they critique, leaving one with the feeling that externalists must surely have something more up their sleeves than what their critics draw from the hat. I think there are several explicit, active cognition arguments for externalism that have not thus far been considered in the literature; I shall focus on one of them.

4   The Argument from Meaning Making

Here is a bald statement of the particular argument that I want to explore in more detail, the *argument from meaning making*, together with a brief gloss on each premise:

(a) *Minds are intentional machines or semantic engines.* This is to identify an important and distinctive feature of minds, namely, that they are mechanisms or devices that operate on items that are intentional or semantic in nature, language-like entities such as propositions and propositional attitudes, but also nonlinguistic actions and objects that are imbued with meaning and significance.

(b) *Intentional machines or semantic engines detect and create meaning.* They detect meaning when it preexists in some structure in the world, whether it be a sentence that someone utters or an action they undertake; they likewise create meaning through both linguistic and nonlinguistic means.

(c) *Meaning detection and creation involve the sequestering and integration of internal and external cognitive resources.* Internal cognitive resources include individual mental symbols and rules for their combination, neural networks, categories, and schemas. External cognitive resources include individual spoken and written symbols and rules for their combination, social

networks, collective representations, and the bodily activities of oneself and other agents.

(d) *Internal cognitive resources are part of the structure of the intentional machine that detects and creates meaning.* This is a traditional assumption within individualistic cognitive science that contributes to explaining how postulated internal structures generate cognitive behavior.

(e) *External cognitive resources often play the same or similar functional roles in the detection and creation of meaning as do internal cognitive resources, or complement, compensate for, or enhance those roles.* External cognitive resources can replace internal cognitive resources (e.g., external memory) or can create capacities in agents that they would not otherwise have (e.g., Kanzi, the bonobo who has exhibited advanced linguistic capabilities). In either case, they are no less central to cognition than are internal cognitive resources.

Thus:

(f) *External cognitive resources, like internal cognitive resources, are part of the structure of the intentional machine that detects and creates meaning.* This follows from a parity assumption, namely, that things that have the same relevant properties vis-à-vis some cognitive process should have the same status in cognition as one another. Whether one is in the head and one is in the environment is irrelevant.

Therefore:

(g) *The extended mind thesis is true.*

There are analogues to premises (c)–(e) in other arguments for externalism, such as those that appeal to the nature of memory, or consciousness, or to some other aspect or property of cognition. Thus, a more thorough examination of at least these parts of the argument from meaning making will perhaps shed light on those arguments as well.

5 Meaning, External Resources, and Fundamentality

There are at least three points at which the argument from meaning making might be challenged by those working within the cognitive sciences and the philosophy of mind. First, one might claim that the very first pair of premises begs a crucial question by assuming a broad view of the kinds of entities to which intentional machines are sensitive. Traditional cognitive scientists might well balk at the idea that intentional machines process "actions" or entities in the world, or might think that the sense in which they do so is incompatible with the extended mind thesis. Second,

the very idea of an external cognitive resource, invoked in the third premise, might be taken to express an oxymoron, in that "cognitive resources," whatever else they might be, must be internal. Third, even those happy to make both of these concessions might well think that the final premise, (e), is indefensible, since there will always remain a crucial asymmetry between internal and external cognitive resources. Roughly speaking, the latter only gain purchase on cognitive activity via the former, and so internal resources remain fundamental to cognition in a way that vitiates the inference to externalism.

Although these objections are closely related, considering each in turn will allow us to discuss a variety of issues to which the debate over the extended mind is important, and to home in on some of those that separate externalists from their opponents. It will also provide some clues about how we might recast the problem of intentionality.

Does the First Premise Beg the Question?

The idea that minds are semantic engines commands relatively widespread assent within the cognitive science community, and there are particular ways of understanding how this idea is to be integrated with the claim that cognition is computational in nature. In traditional, AI-inspired cognitive science, the idea was roughly this: minds are semantic engines in virtue of being syntactic engines, together with the correlation between the syntactic strings that the rules governing such syntactic engines operate on and the meaning that we ascribe to those strings. More precisely, these syntactic strings *encode* meaningful mental representations and are themselves in turn *realized in* physical configurations in the brain. In connectionist cognitive science, matters are more complicated, in part because some connectionists are *eliminativists* about the idea of mental meaning, and so altogether reject the idea that minds are semantic engines, while others view themselves as providing an account of the realization of mental representations that simply dispenses with the correlation thesis. For those connectionists willing to hop on the mental representation bus, representations can be nonpropositional, subsymbolic, or distributed, for example, with the networks underlying representation departing from the traditional rules and representations of traditional cognitive science.

Although both traditional and connectionist forms of representationalism acknowledge that there is causal interaction across the boundary between world and head, both also hold that the intentionality of the mental lies firmly on the head side of this boundary, whether it derives from some kind of syntactic encoding or emerges more holistically from a

distributed system of interconnected neural nodes. The world can provide background conditions necessary for representations to have the specific content that they do, but it is not physically constitutive of representational content itself. Meaning is not out there in the world waiting for us to detect it, but created by the mind in its interactions with the world; meaning is *intrinsic* to in-the-head mental representation. If objects in the world are part of the representational content of our mental representations, that is because (let's assume) they are coded for by intrinsic properties of rich in-the-head representations, or because (again, let's assume) they are distributively represented by simple in-the-head representations.

My own view is that these general conceptions of meaning and representation are flawed (see Wilson 2000, 2004, chap. 7), but the question here is whether the argument from meaning making presumes this in a way that begs the question in favor of the extended mind. To see that it does not, we should construe premises (a) and (b) in a way that is neutral between various accounts of how mental representation operates. All that has to be conceded in accepting (a) and (b) is that (i) agents have some kind(s) of mental representation that play causal roles in their perception and behavior; (ii) whatever form(s) these mental representations take, and whether or not they are located exclusively inside the skin, they often causally derive from and in turn causally generate actions, objects, and events that are not so bounded. Both (i) and (ii) concern the phenomena to be explained, and as such should be granted by individualists and externalists alike.

This clarificatory response might be thought to place a heavier burden on premise (c), which now appears to make a substantial claim about how to conceptualize the detection and generation of meaning. If there is, however, something suspect about the very idea of an external cognitive resource, then the argument from meaning making is in trouble.

Is "External Cognitive Resource" an Oxymoron?

In introducing the term "cognitive resource" in section 3 I invoked an analogy between cognitive resources and developmental resources, an analogy that calls for more extensive explication. Within the philosophy of biology over the past two decades the predominant gene-centered view of inheritance and development has been challenged by a number of alternative paradigms. Developmental systems theory (hereafter, DST), deriving from the work of Susan Oyama (1985, 2000), is one of these paradigms. In the past decade DST has shifted from providing a critique of the genocentrism of developmental biology to articulating its own, positive vision

of research on inheritance and organismic development. The key idea in the DST paradigm is that the most fundamental units of agency needed to understand these biological phenomena are developmental *systems*, where genes are one among many elements in those systems. These elements are called developmental *resources*, but since such resources always operate within the context of some developmental system or other, they should be viewed as neither ontologically nor epistemically more fundamental than those systems. Although not all forms of DST have embraced the idea that developmental systems, and so cognitive resources, can extend beyond the boundary of the skin of the organism, some proponents of DST have advocated precisely this view, arguing that this is a consequence of the logic behind DST, together with the empirical details that are provided by a consideration of the inheritance and development of particular traits.

As I have suggested elsewhere (Wilson 2005, chap. 7), there is a fairly direct parallel between the relationship between DST and traditional views of inheritance and development, on the one hand, and locational externalism and traditional views of cognition, on the other. Just as (wide) DSTers in biology argue that developmental systems extend beyond the organism's boundaries, including developmental resources from both sides of those boundaries, externalists in psychology argue that cognitive systems extend beyond the head of the individual, including cognitive resources both inside and outside the skin. Indeed, the analogy between the two cases reaches as far as the connotations possessed by the most ubiquitous adjectival terms in each—"genetic" and "mental"—both of which convey the sense of an entity, process, or activity that goes on entirely inside an organism or agent. Hence, the idea of an extended mind sounds just as odd, perhaps, as does that of an extended gene (or even an extended genetic system).

This connation of "innerness" stems, I suggest, from a common source: from an overarching dominant framework for thinking about both organismic development and cognitive agency that is individualistic. To challenge that framework in either or both areas of inquiry, some less heavily loaded terminology is needed. "System" and "resource" fill this niche in both cases, together with a shift from "mental" and "genetic" to "cognitive" and "developmental" (or "heritable," when moving from ontogeny to intergenerational transmission).

Suppose, then, that we can accept the minimal readings of premises (a) and (b) that remain neutral about the mode of representation through which cognition operates, and that there is some rationale for shifting

from semantically loaded terms to those that allow the idea of an external cognitive resource to make no less sense than that of an external developmental resource. The claim in premise (c) is, roughly, that intentionality operates in the mind via the integrated functioning of internal and external cognitive resources. This itself remains compatible with individualism about cognition, provided that there is some way to single out only internal cognitive resources as themselves truly cognitive.

Are Internal Resources More Fundamental to Cognition Than Are External Resources?

One way to express the extended mind thesis is to say that, for at least a variety of cognitive activities, the physical configuration of the brain is not metaphysically sufficient for their performance *qua* cognitive activities. Something more is needed, and that something more involves the physical configuration of the world beyond the head. When locational externalism is expressed in this way, however, pointing to an asymmetry between what's in the head—Brain—and what's outside of it—World— seems easy. For Brain is always involved in cognition, whereas World is only sometimes involved. More to the point, World is only putatively relevant to cognition when it forms part of a causal chain that passes through Brain. But the converse is not true of Brain. In short, since World must be causally connected to Brain in order for World to result in cognition, but not vice versa, cognitive resources are properly all internal, and so restricted to Brain. Thus, the mind is not extended, and what I have been calling "external cognitive resources" are more accurately labeled in their traditional way: they are *inputs to*, not parts of, cognition.

That brain activity is always involved in cognition is not something that I wish to dispute, since I'm happy to stick my neck out and say that it is actually *true*. The significance of this for the debate over the extended mind, however, is what is up for grabs. Consider again the analogy to developmental systems theory in biology. One should grant that gene expression is always causally involved (however minimally or indirectly) in the development of any particular phenotypic trait an organism possesses, whereas other developmental resources sometimes are and sometimes are not causally active. But whether that provides reason to endorse accounts of ontogenetic development exclusively in terms of genes and their immediate expression remains a further issue. (In this case, this is not least of all because other developmental resources are always required for genes not only to be expressed as organism-level phenotypic traits, but for genes to do anything at all.)

To maintain our analogy with the case of cognition, however, we need to consider not simply genes but the full range of internal developmental resources—methyl groups, RNA complexes, mitochondria, cytoskeletal structures (such as actin fibers and microtubules), and so on. And there certainly are developmental systems—the chromatin marking system would be one—whose operation does consist solely of changes within the organism. Such processes are internalist, and though they do not operate in a contextual vacuum, the environment is properly viewed an providing either inputs to these processes or as specifying background conditions necessary for their functioning. The same is true, I think, of some cognitive processes, such as some acts of visual imagining, some kinds of problem solving, and some decision making. Yet this does not establish the requisite asymmetry between internal and external resources, in either the biological or the psychological case, because it is not true for *all* developmental and cognitive processes. Consider two examples that I have discussed previously (Wilson 2004, chap. 8) that differ in an important way.

First, consider the bonobo Kanzi's planning, thinking, and decision making that utilizes a 256-symbol keyboard, which Kanzi has, over the years, learned to use to communicate his beliefs and desires. Bonobos in the wild surely have desires (e.g., for bananas), but just as surely they don't have the kind of sophisticated, cooperative desires that Kanzi expresses, such as the desire to be taken by a particular person to a sequence of locations, or to do one activity first and then another. What the symbol board has done is to reconfigure Kanzi's capacity for belief and desire, much as our using pen and paper reconfigures our mathematical ability by augmenting the in-the-head capacity we have for multiplication. Both are cases in which an external symbol system becomes integrated with preexisting cognitive capacities in ways that significantly modify the nature of those capacities. We can, of course, distinguish between the *parts of* those capacities that are internal and those that are not, but this is already to concede that the overall cognitive process itself is extended.

Second, suppose that you are playing a visual problem-solving game, such as Rush Hour, or completing a jigsaw puzzle. (Rush Hour involves a square board loaded with cars and trucks in various positions, and the object is to move these in a sequence so as to allow a designated car to leave the board through the only exit.) Here the problem solving in part consists of internal mental operations (let us suppose), but also in part it consists of the active physical manipulation of pieces on a board or pieces within the puzzle. One might suppose that these manipulations—rotations of puzzle pieces, or trials of car move sequences—were simply pre-

sented to you in a computer simulation, or through an automation of trial moves from which you select the best outcome. I suspect that this would change the problem-solving task in a significant way; but note that even here problem solving requires active engagement with a part of the world beyond the head, namely, the various visual displays from which one must select. In both cases, the problems are solved by utilizing, exploiting, or manipulating a set of resources, some of which are outside of the head. These are not simply inputs to those that are inside the head, because the very process of problem solving involves them as much as it involves resources inside the head. There is nothing bounded by the skull that counts as solving these kind of problems.

In both cases it is not simply environmental structures that somehow magically make for cognition, but the causal integration of these with onboard capacities that organisms already have. In general this causal integration can be ontogenetic or phylogenetic, individual or collective, cultural or biological, and the cognitive abilities such integration generates are often genuinely novel. Some of these compensate for organismic cognitive limitations (say, in short-term memory capacity) through external augmentation; some may actually further diminish such internal capacities (as, for example, the shift from oral to literary traditions has likely done with respect to auditory recall). The point here is that there is a lot of variation in the kinds of transcranial cognitive processes that the extended mind thesis embraces, and a corresponding breadth to the range of cases of which it is true.

This brings me to the chief difference between these two examples that I want to draw explicit attention to. This is that whereas the first involves the causal integration of explicit symbols located in an organism's environment into that organism's cognitive regime, the second appeals to the cognitive incorporation of nonsymbolic aspects of that environment. Much of the discussion of the extended mind has focused exclusively on cases of just the former kind. In combination with the emphasis on cyborg fantasy arguments for the extended mind thesis, this has created (however inadvertently) the impression that the scope of the thesis is narrower than it actually is. It is primarily a view neither about technological enhancement of existing cognitive abilities, nor about the conceptual possibility of the mind being extended. Rather, it is a view of what cognition is for at least as long as it has been a trait of human beings, before we had writing systems (6,000 or so years ago), and before we developed pictorial representation (25,000 or so years ago).

I would hazard a guess that the cognitively most significant forms that our nonsymbolic environment takes are, broadly speaking, *social* in

nature. (I also think that these are, in many respects, more fundamental than the technological innovations that have also reshaped our cognitive architecture, but that is not my plaint here). These social features include the interpersonal relations found in human social groups—from dyads through to face-to-face communities; the group-level structures, such as kin groups and social hierarchies, that human agents are sensitive to; and the range of social practices, customs, and rituals—such as song, burial, and marriage—that can be found stretching back various distances through the history of the species. Put roughly, the idea is that it is because we are socially oriented creatures in a fairly rich and deep sense that we have biologically evolved and culturally developed extended minds that incorporate these very aspects of our social world.

This isn't to say that technological development and appropriation is unimportant in thinking about extended cognition, or that it is never a primary shaper of the extended mind. Rather, it is to claim that we need also to take the nonsymbolic environment seriously in articulating the extended mind thesis, and that the most important place to look in doing so is the social realm. The social and the technological are both significant aspects of extended cognition—evolutionarily, historically, biologically, culturally—and I suspect that it is in tandem that they have sculpted human cognition over tens of thousands of years to its present level.

6   Rethinking the Problem of Intentionality

I began by identifying three parts to the problem of intentionality—to say what intentionality is, to delineate things that have intentionality from things that don't, and to provide an account of in virtue of what it is that representations have the particular content that they do. The problem, especially this final part of it, has been subject to a naturalistic constraint, and the intentionality of the mental has been taken to be the core phenomenon to be explained. At the end of section 1 of the chapter I suggested that the problem of intentionality was complicated by internalism; at the end of section 2 I said that externalism would provide us with the means for reconceptualizing that problem. How so? In at least three ways.

First, the extended mind thesis amplifies the number of kinds of mental representations that there are by including at least some external representations as cognitive in nature. One might think that this makes the problem of intentionality worse, in that instead of being faced with a rela-

tively well-circumscribed set of entities to puzzle about—internal mental representations—we now have a hodgepodge comprised of almost anything we make perceptual or cognitive contact with. But I think just the opposite is true. Recognizing the variety of mental representations embraced by the extended mind thesis undermines the very terms in which the problem of intentionality has been formulated. It is not so much that the question "In virtue of what does internal state R represent condition (state, object, property) C?" has been replaced by one that substitutes "internal or external" for "internal," as that this is just the wrong kind of question to ask about representation. It is to ask for something like the *essence* of mental representation, and one of the consequences of the extended mind thesis, I am suggesting, is that this is to ask for something that does not exist.

But what *is* the right question to ask here? Does the promised reconceptualization of the problem of intentionality go beyond this diagnosis of error? This brings me to a second way in which externalism helps: by shifting our focus from representational essences to representation in practice. Although I have distinguished between "active cognition" and "cyborg fantasy" arguments for the extended mind, both arguments take the dynamic aspects to cognition seriously and appeal to how we engage in cognitive practices that make use of representations. The shift is one from a focus on "things," such as representations, to a concern with "activities," such as the act of representing. Such activities are often bodily, and are often world-involving in nature. A version of the problem of intentionality formulated so as to apply to them—"In virtue of what is activity A the representation of C?"—seems hardly pressing at all. Rather, what cries out for discussion is the question of just what forms these activities take, and just how they bring about the effects they do.

And so a third way in which the extended mind thesis allows us to reconceptualize the problem of intentionality is that it suggests a methodological reorientation. We no longer seek the essence of representations, for there is none to be found; we focus instead on activities or acts of representing. But the methodology appropriate here, in the sense of likely being the most fruitful in outcome, is not traditional conceptual analysis but an interdisciplinary, pluralistic motley. The task is to understand a variety of representational practices and wherein they are representational. The means we employ in doing so will be various: historical analysis of their emergence, sociological analysis of the conditions under which they operate, experimental psychological analysis of representational gaps

and gluts, anthropological analysis of practices of symbolization, evolutionary analysis of social environments and our sensitivity to them.

This is very much to echo John Sutton's (2004, this volume) call for a more adventuresome interdisciplinary development of the extended mind thesis (see also Chemero and Silberstein 2008; Barnier et al. 2008; MacIver 2009; MacIver and Wilson in progress), one that doesn't simply augment cognitive science with a few snippets from studies of human–computer interaction or speculative tales about the evolution of cognition. When cognition is extended, intentionality is extended and the traditional problem of intentionality transformed. Precisely what that transformation results in will turn largely on how thoroughly the philosophical naturalism driving the work on the extended mind becomes integrated with the kind of interdisciplinary adventures mentioned above. As that integration proceeds, we will gain a deeper sense of what intentionality is and a clearer idea of what things have it and what things don't; and we will be as moved by the final, "in virtue of" part of the problem of intentionality as we are by the question of in virtue of what gravity operates, or in virtue of what matter exists.

Acknowledgments

References

Adams, F., and Aizawa, K. (2001). The bounds of cognition. *Philosophical Psychology*, *14*, 43–64.

Adams, F., and Aizawa, K. (2008). *The Bounds of Cognition*. New York: Blackwell.

Adams, F., and Aizawa, K. (2009). Why the mind is still in the head. In P. Robbins and M. Aydede (eds.), *The Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press.

Aizawa, K. (2007). The embodiment of perception. *Journal of Philosophy*, *104*, 5–25.

Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, *48*, 57–86.

Ballard, D. H., Hayhoe, M. M., Pook, P. K., and Rao, R. P. N. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences*, *20*, 723–742.

Barnier, A., Sutton, J., Harris, C., and Wilson, R. A. (2008). A conceptual and empirical framework for the social distribution of cognition: The case of memory. *Cognitive Systems Research*, *9*(1–2), 33–51.

Burge, T. (1979). Individualism and the mental. In P. French, T. Uehling Jr., and H. Wettstein (eds.), *Midwest Studies in Philosophy*, vol.4: *Metaphysics*. Minneapolis: University of Minnesota Press.

Burge, T. (1986a). Individualism and psychology. *Philosophical Review*, *95*, 3–45.

Burge, T. (1986b). Intellectual norms and foundations of mind. *Journal of Philosophy*, *83*, 697–720.

Chemero, A., and Silberstein, M. (2008). After the philosophy of mind: Replacing scholasticism with science. *Philosophy of Science*, *75*, 1–27.

Clark, A. (1997). *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press.

Clark, A. (2003). *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. New York: Oxford University Press.

Clark, A. (2008). *Supersizing the Mind*. New York: Oxford University Press.

Clark, A., and Chalmers, D. (1998). The extended mind. *Analysis*, *58*, 7–19. Reprinted as chapter 2 of this volume.

Cole, M. (1996). *Cultural Psychology: A Once and Future Discipline*. Cambridge, MA: Harvard University Press.

Crane, T. (1991). All the difference in the world. *Philosophical Quarterly*, *41*, 1–25.

Devitt, M. (1990). A narrow representational theory of mind. In W. G. Lycan (ed.), *Mind and Cognition: A Reader*. Cambridge, MA: Blackwell.

Dretske, F. I. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

Dretske, F. I. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.

Ebbs, G. (1998). *Rule-Following and Realism*. Cambridge, MA: Harvard University Press.

Egan, F. (1992). Individualism, computation, and perceptual content. *Mind*, *101*, 443–459.

Egan, F. (1995). Computation and content. *Philosophical Review*, *104*, 181–203.

Fodor, J. A. (1980). Methodological solipsism considered as a research strategy in cognitive psychology. *Behavioral and Brain Sciences 3*: 63–73. Reprinted in Fodor 1981.

Fodor, J. A. (1981). *Representations*. Sussex: Harvester Press.

Fodor, J. A. (1982). Cognitive science and the Twin-Earth problem. *Notre Dame Journal of Formal Logic*, *23*, 98–118.

Fodor, J. A. (1984). Semantics, Wisconsin style. *Synthese*, *59*, 231–250.

Fodor, J. A. (1987). *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.

Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Boston: Houghton-Mifflin.

Grush, R. (2003). In defense of some Cartesian assumptions concerning the brain and its operation. *Biology and Philosophy*, *18*, 53–93.

Haugeland, J. (1998). Mind embodied and embedded. In his *Having Thought: Essays in the Metaphysics of Mind*. Cambridge, MA: Harvard University Press.

Hurley, S. (1998). *Consciousness in Action*. Cambridge, MA: Harvard University Press.

Hurley, S. (2001). Perception and action: Alternative views. *Synthese*, *129*, 3–40.

Loar, B. (1988). Social content and psychological content. In R. Grimm and D. Merrill (eds.), *Contents of Thought*. Tucson: University of Arizona Press.

MacIver, M. (2009). Neuroethology: From morphological computation to planning. In P. Robbins and M. Aydede (eds.), *The Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press.

MacIver, M., and Wilson, R. A. (in progress). Taking an extended view of embodied cognition: Evolution, control, and computation.

Menary, R. (2007). *Cognitive Integration: Mind and Cognition Unbounded*. London: Palgrave Macmillan.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.

Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.

O'Regan, J. K., and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*, 939–1031.

Oyama, S. (1985). *The Ontogeny of Information* (2nd ed.). Durham, NC: Duke University Press, 2000.

Oyama, S. (2000). *Evolution's Eye: A Systems View of the Biology–Culture Divide*. Durham, NC: Duke University Press.

Pettit, P. (1993). *The Common Mind: An Essay on Psychology, Society, and Politics*. New York: Oxford University Press.

Putnam, H. (1975). The meaning of "meaning." In K. Gunderson (ed.), *Language, Mind, and Knowledge*. Minneapolis: University of Minnesota Press. Reprinted in H. Putnam, *Mind, Language, and Reality: Philosophical Papers* (vol. 2). New York: Cambridge University Press, 1975.

Rowlands, M. (1999). *The Body in Mind: Understanding Cognitive Processes*. New York: Cambridge University Press.

Rowlands, M. (2003). *Externalism: Putting Mind and World Back Together Again*. Chesham: Acumen.

Rupert, R. (2004). Challenges to the hypothesis of extended cognition. *Journal of Philosophy*, *101*, 389–429.

Segal, G. (1989). Seeing what is not there. *Philosophical Review*, *98*, 189–214.

Shapiro, L. (1997). A clearer vision. *Philosophy of Science*, *64*, 131–153.

Sutton, J. (2004). Representation, reduction, and interdisciplinarity in the sciences of memory. In H. Clapin, P. Staines, and P. Slezak (eds.), *Representation in Mind: New Approaches to Mental Representation*. Amsterdam: Elsevier.

Wertsch, J. (1998). *Mind as Action*. New York: Oxford University Press.

Wertsch, J. (2002). *Voices of Collective Remembering*. New York: Cambridge University Press.

Wilson, R. A. (1992). Individualism, causal powers, and explanation. *Philosophical Studies*, *68*, 103–139.

Wilson, R. A. (1994). Wide computationalism. *Mind*, *103*, 351–372.

Wilson, R. A. (1995). *Cartesian Psychology and Physical Minds: Individualism and the Sciences of the Mind*. New York: Cambridge University Press.

Wilson, R. A. (2000). The mind beyond itself. In D. Sperber (ed.), *Metarepresentations: A Multidisciplinary Perspective* (pp. 31–52). New York: Oxford University Press.

Wilson, R. A. (2003a). Individualism. In S. Stich and T. A. Warfield (eds.), *The Blackwell Companion to Philosophy of Mind* (pp. 256–287). New York: Blackwell.

Wilson, R. A. (2003b). Pluralism, entwinement, and the levels of selection. *Philosophy of Science*, *70*, 531–552.

Wilson, R. A. (2004). *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. New York: Cambridge University Press.

Wilson, R. A. (2005). *Genes and the Agents of Life: The Individual in the Fragile Sciences: Biology.* New York: Cambridge University Press.

Wilson, R. A., and Clark, A. (2008). How to situate cognition: Letting nature take its course. In P. Robbins and M. Aydede (eds.), *The Cambridge Handbook of Situated Cognition*. New York: Cambridge University Press.

Wilson, R. A., and Craver, C. (2007). Realization. In P. Thagard (ed.), *Handbook of the Philosophy of Science*, vol. 12: *Philosophy of Psychology and Cognitive Science* (pp. 81–104). Amsterdam: Elsevier.

Zerubavel, E. (1997). *Social Mindscapes: An Invitation to Cognitive Sociology.* Cambridge, MA: Harvard University Press.