

Continuing Commentary

Commentary on Annette Karmiloff-Smith (1994). *Précis of Beyond modularity: A developmental perspective on cognitive science*. BBS 17:693–745.

Abstract of the original article: *Beyond modularity* attempts a synthesis of Fodor's anticonstructivist nativism and Piaget's antinativist constructivism. Contra Fodor, I argue that: (1) the study of cognitive development is essential to cognitive science, (2) the module/central processing dichotomy is too rigid, and (3) the mind does not begin with prespecified modules; rather, development involves a gradual process of "modularization." Contra Piaget, I argue that: (1) development rarely involves stagelike domain-general change and (2) domain-specific predispositions give development a small but significant kickstart by focusing the infant's attention on proprietary inputs. Development does not stop at efficient learning. A fundamental aspect of human development ("representational redescription") is the hypothesized process by which information that is *in* a cognitive system becomes progressively explicit knowledge *to* that system. Development thus involves two complementary processes of progressive modularization and progressive "explicitation." Empirical findings on the child as linguist, physicist, mathematician, psychologist, and notator are discussed in support of the theoretical framework. Each chapter concentrates first on the initial state of the infant mind/brain and on subsequent domain-specific learning in infancy and early childhood. It then goes on to explore data on older children's problem solving and theory building, with particular focus on evolving cognitive flexibility. Emphasis is placed throughout on the status of representations underlying different capacities and on the multiple levels at which knowledge is stored and accessible. Finally, consideration is given to the need for more formal developmental models, and a comparison is made between representational redescription and connectionist simulations of development. In conclusion, I consider what is special about human cognition by speculating on the status of representations underlying the structure of behavior in other species.

How far beyond modularity?

Luca Bonatti

Laboratoire des Sciences Cognitives et Psycholinguistique, 75006 Paris, France. wca@cogito.lscp.ehess.fr

Abstract: I question (1) whether Karmiloff-Smith's (1994a,r) criticisms of modularity hit the target and (2) how much better the representational redescription model is. In both cases, "the mystery of the cognitive clock" is problematic for her account.

Fodor's (1983; see also multiple book review of Fodor's "The modularity of mind" in *BBS* 8 (1) 1985) modularity thesis (M) is a central point of reference for cognitive scientists, but it is not dogma. There are at least two critical attitudes toward it. One is fairly simple: it consists in showing that M is false for some core domains. This can be done by providing evidence that, in an allegedly modular domain, information flows freely at any moment during processing. For example, if contextual information is always accessible on line during lexical access or syntactic processing, then M is false for these domains. Many people have raised doubts about the truth of M in specific cases; the debate is still open and fruitful.

Another, more ambitious, way is to offer an alternative or a more comprehensive framework that explains data and intuitions in favor of M and also covers phenomena that M leaves unexplained. Karmiloff-Smith (1992) in *Beyond modularity*, has taken the second route. It is not unimportant to determine whether her attempt is correct, because, as some have argued, if it were viable it might have substantial consequences in philosophy, in psychology (see Dartnall 1994; Losonsky 1994), and perhaps even in pedagogy (Estes 1994). Her attempt certainly deserves admiration; however, I find it unsuccessful.

To state my point briefly, I don't think *Beyond modularity* goes beyond modularity. Other commentators on Karmiloff-Smith's *précis* (1994t) have hinted at such a conclusion in one form or

another (de Gelder 1994; Foster-Cohen 1994; Ohlsson 1994), but I don't think a fully explicit argument has been advanced and, necessarily, Karmiloff-Smith has only partially addressed those partial arguments (see Karmiloff-Smith 1994r, pp. 734–35). Briefly stating the argument can help in developing a brief and full answer, to everybody's satisfaction.

I think that Karmiloff-Smith's proposal is not an advance over M for at least two reasons. First, I do not see the precise nature of her criticisms of M; second, I am unclear about what gains the new framework would or could ensure. Both reservations are tied to what I will call the "mystery of the developmental clock": many cognitive abilities follow a strict order of development, both functional and temporal, regardless of the radical differences in inputs from child to child. For example, all normal neonates distinguish linguistic from acoustic nonlinguistic signals (Bertoncini et al. 1989) and can discriminate among different classes of languages (Mehler et al. 1988); however, they lose the ability to discriminate some contrasts irrelevant to their natural language at around 10 to 12 months (Werker & Tees 1984). Similarly, all normal children pass the "false belief task" at around 4 years of age, but not a few months earlier. Ten-month-olds seem unable to exploit properties of objects they nevertheless encode in order to draw sortal distinctions, but 12-month-olds succeed (Xu & Carey 1996). There are many cases like these. The mystery is not a real mystery, but only a problem, if one takes a strong nativist stand; the developmental clock is a biological clock, with the environment acting as a trigger. But then one must acknowledge that there is no real cognitive development "in the sense that developmental cognitive psychologists have in mind" (Fodor 1983). Karmiloff-Smith wants to resist this conclusion, but then she must eliminate the mystery and I do not think she can.

Does modularization go beyond modularity? Here is the first reason for my doubts. According to Karmiloff-Smith, there are two pieces of data one must account for: initial brain plasticity and final rigidity. For Karmiloff-Smith, M is consistent with final rigidity but at odds with initial brain plasticity. To encompass both aspects,

she proposes to abandon strict M and to account for final rigidity through progressive modularization. Of course, if modularization were only the maturational unfolding of a genetic program plus or minus a bit, there would be no need to go beyond M: if nativism makes any sense, development in this sense is fully nativist. But for Karmiloff-Smith, modularization is something quite different, namely, the result of a true process of interaction with the environment (1992, p. 10; see also 1994a, p. 733). Notice that M and modularization do generate different explanations. Consider the example: Karmiloff-Smith thinks that at the outset language processing is not modular, but becomes modularized because, *inter alia*, it is processed very fast: "Attention biases and some innate predispositions could lead the child to focus on linguistically relevant input and, with time, to build up linguistic representations that are domain-specific. Since we process language very rapidly, the system might with time close itself off from other influences – i.e., become relatively modularized" (1992, p. 36). On the other hand, M would give an opposite explanation: we process language very rapidly because there is a language module. Hence it is important to see which paradigm is correct.

I see two problems in Karmiloff-Smith's claim that M should be substituted with modularization. First, even if M is compatible with brain localization of function at birth – for which there is indeed evidence (Bertoncini et al. 1989; Best 1988) – it does not require localization. M is a thesis about the functional organization of the mind, and not directly about the brain; hence, simple brain plasticity is no basis for rejecting it (as Karmiloff-Smith also recognizes, *contra* Quartz & Sejnowski 1994). M would indeed be in trouble if the initial state and the time course of development showed no functional modular architecture, but Karmiloff-Smith offers no evidence for that conclusion. Instead, there is much evidence (albeit inconclusive) that very specific signal processors are present at birth and are tuned for very specific features of the input (e.g., for syllables – Bijelac-Babic et al. 1993 – but not for morae – Bertoncini et al. 1995). Such specificity and its developmental schedule seem to be much more than the generic innate predispositions Karmiloff-Smith grants. Second, if her modularization is more than the genetic unfolding of a program in which the environment plays a triggering role, we would expect typical individual and group differences, determined by sharp differences in the experiential histories of the organisms, that we do not see. Of course, one can still pay lip service to development and insist that the unfolding of modules is not really M but a process of modularization, but then in such perspective the developmental clock becomes a mystery. Consider American Sign Language (ASL), for example. Because deaf children reach full linguistic mastery even if deprived of normal acoustic stimuli, Karmiloff-Smith concludes that ASL learning by the deaf child supports modularization against M. However, what is striking in the ASL case is not that deaf children learn a natural language: once again, this is not an argument against the existence of modules, but at most an indication of the level of abstractness of their inputs. What is striking is rather that the development of their linguistic abilities closely matches that of normal children regardless of the sharp difference in inputs (e.g., Petitto & Marentette 1991). This is a prediction for M, but it is only a lucky coincidence for modularization. The developmental clock makes modularization either highly implausible or just another name for M.

Can representational redescription be an alternative to nativism? The second reason for my doubt concerns Karmiloff-Smith's specific positive proposal, the representational redescription (RR) model. Deep down, Karmiloff-Smith proposes a modified Piagetianism integrated by "some innately specified, domain-specific predispositions that guide epigenesis" (1992, p. 172). There are other important differences between Piaget and Karmiloff-Smith, but in this context what counts is that both assume that a domain-general mechanism is responsible for cognitive development – RR in this case (1992, pp. 25 and 167). And the trouble is, RR is too Piagetian, both for the vagueness of its specification and for its centrality, and I feel that it is going to have the same problems as

the Piagetian developmental theory. Let me state first what I think the real issue is, and then why I think RR fails in elucidating it.

Karmiloff-Smith raised a very interesting problem. During development, information implicit in the mind becomes knowledge to the mind, either explicitly or (sometimes) consciously: What is the role of these changes? RR is supposed to cast light on this issue. Now, a first reservation concerns what the real problem is exactly, and whether RR is really useful even to describe it. Notice that RR is proposed as a model of changes in the format of the internal representations, but it is not clear that the issue at hand really concerns representational *formats*. Suppose that a parser incorporates a grammar with explicit rules, which it can consult in the process of syntactic analysis. In that case, information is explicitly encoded in the parser, but may simply not be accessible to other cognitive departments. With time, we may imagine that some (but, crucially, not all) of the rules of the grammar are marked as accessible, hence open to the child's metareflection. Notice that in such a scenario no redescription or change in the format of representation is necessary, and yet the same phenomena Karmiloff-Smith mentions can be explained. In short, the accessibility of representations and their formats are orthogonal issues, and I don't see arguments for assuming that development engenders change in the latter. If I am right, then the important problem Karmiloff-Smith raises is not that the format of internal representations changes – it may well remain the same across the board – but that some piece of information can become explicit (at different levels) and some other piece is not and cannot. And then, under this reading, the interesting questions are detailed ones concerning what information becomes explicit, why, and what difference it makes (if any) for cognition. For example, 6-year-olds reach metalinguistic awareness of the category "word," and 10-year-olds of the possessive determiner system, but no child or untrained adult ever reaches conscious awareness of C-Command. Why, and what differences would it make if it were otherwise? Now, strictly speaking, RR cannot answer these types of questions because they are about representations changing their formats, and not about representations becoming more or less accessible.

One can easily modify RR and transform it into a taxonomy for different degrees of accessibility of representations (even if, in that case, it loses much of its "revolutionary" aspect). But then a second reservation comes to mind: in the absence of a specification of the mechanisms that should lead to the phases of redescription/increased accessibility postulated in the RR model, one is still unsure about its real explanatory role. In her response, Karmiloff-Smith remarks that nativists are no better off, since they often describe only scantily the mental operations that would embody the alleged innate structures in the child's mind. She surely makes a point here, but not one that is of much consolation for a theory such as hers, which is supposed to specifically vindicate *cognitive development*. Such a notion does not play any real role in a nativist framework, but it *is* the core of a constructivist's theory. Piaget already postulated something close to Karmiloff-Smith's redescription, the *abstraction réfléchissante*. The trouble was that he never explained what that was. The best we can say is that one must wait to see whether RR will be empirically richer than Piaget's *abstraction réfléchissante*.

I can grant Karmiloff-Smith that the above criticism is unfair and that it may be premature to ask from RR a detailed description of the mechanism of cognitive development. Then the crucial question for evaluating RR becomes: Can a more detailed version of it account for the phenomena of cognitive development any better than the current RR model? Even here, I doubt that it can be successful. My last reservation has to do with the structure of an RR-like model. We know that explaining cognitive development through unconstrained theory construction from unconstrained data is hopeless, even for developmental facts much less problematic than the cognitive clock. Now, RR is a domain-general mechanism, *just like theory construction*. How can it be any better? Karmiloff-Smith tries to make it different by conceding a

point to the nativist: RR does not start from a *tabula rasa*, but bootstraps from some initial innate attention biases. But notice that making room for selective attention does not render a model different from general theory construction. Even when we build a theory of, say, microphysics, we have selective attention, because we disregard data about ancient history. To construct a theory involves precisely the use of a domain-general mechanism applied to a domain-specific input. Add to it that the domain-specific initial biases granted by Karmiloff-Smith are really very poor: she concedes a “fairly limited amount of innately specified domain-specific predisposition” (1992, p. 4), or, as she writes in her response, only “some *minimal* predispositions” (1994, p. 733). So the tools she allows herself are a set of very poor domain-specific predispositions plus a domain-general mechanism, namely, RR. It looks like the available tools are even poorer than those of theory construction. But if that is all, how to square development with the uniqueness of its outcomes and its tight time course? Nobody expects people to come up with the same theories at the same time precisely because the looseness with which the elaboration of a theory is tied to its data eliminates the possibility of making such predictions. Likewise, we would expect “re-descriptions” (as well as modularizations) to occur at any moment, almost unpredictably during cognitive development; we would expect broad variations in, say, phonological, semantic, or grammatical abilities that we do not find. And so the cognitive clock would still be a mystery.

Karmiloff-Smith is sensitive to this problem, so she adds another factor in order to explain the absence of variability. She writes: “The fact that development proceeds in similar ways across normal children does not necessarily mean that development must be innately specified in detail, because it is also true that all children evolve in species-typical environments. Thus, it is the interaction between similar innate constraints and similar environmental constraints that gives rise to common developmental paths” (1992, p. 172). For the sake of argument, let me grant this point as well. Now the question becomes: Are poor innate biases plus a domain general mechanism plus similar environmental constraints sufficient to cope with the cognitive clock? I doubt it. The notion of similar environmental conditions is too generic to cut any ice. What would it mean for language, for example? There are children who grow up in a monolingual environment and end up with the right grammar, the right phonology, the right semantics, and the same degree of metalinguistic awareness, roughly at the same ages. Other (indeed most) children grow up in multilingual environments right from the start, and yet never end up with mixed linguistic systems, or with wrong phonologies, or with higher or lower degrees of explicit grammatical knowledge. The linguistic environment of these two groups is surely similar in certain respects – for example, almost all children are exposed to motherese, and so on – but the trouble is that most of these factors are known to be irrelevant: the notion of “similar environment” is far too coarse-grained to account for children’s regularities in the details.

In sum, I do not yet see how the RR model can offer an alternative to the nativist paradigm, in which general learning, theory construction, or epigenesis play a severely constrained role. If Karmiloff-Smith wants to pursue her project, she must find a way to remove the air of mystery from the developmental clock.

ACKNOWLEDGMENTS

Thanks to Jerry Fodor and Jacques Mehler for useful comments on a first version of this commentary. When not otherwise indicated, references are to Karmiloff-Smith’s *Beyond modularity*.

Putting knowledge to work

Derek Browne

Department of Philosophy, University of Canterbury, Christchurch, New Zealand. d.browne@phil.canterbury.ac.nz

Abstract: Representational redescription (Karmiloff-Smith 1994a; 1994) translates implicit, procedural knowledge into explicit, declarative knowledge. Explicit knowledge is an enabling condition of cognitive flexibility. The articulation and inferential integration of knowledge are important in explaining flexibility. There is an interesting connection to the availability of knowledge for verbal report, but no clear explanatory work is done by the idea of knowledge that is available to consciousness.

Knowledge at the “entry level,” Level I, is implicit, procedural, domain-specific, proprietary to particular skills, chunked into unarticulated wholes, and not transportable to other cognitive operations. Animals equipped only with this kind of knowledge are incapable of the cognitive, and so the behavioural flexibility that mature human beings display. To explain this versatility, we posit knowledge that is explicit, declarative, available across all domains, and articulated into elements that are independently transportable. Level-E3 knowledge is maximally explicit. It is available for verbal report. At the E2 level, it is (just) consciously accessible. It is minimally explicit when it is (just) available for multiple uses: that is, when it ceases to have a proprietary application. The process that translates implicit knowledge into explicit knowledge is “representational redescription” (Karmiloff-Smith 1994a; 1994). There are a lot of ideas here, some doing more explanatory work than others. I will argue that the articulation and inferential integration of knowledge are crucial to the development of cognitive flexibility. There is a plausible explanatory connection between cognitive flexibility and knowledge available for verbal report, which I identify as declarative knowledge in a strict sense. But the idea that some of our knowledge is consciously accessible does not do the required explanatory work.

1. Implicit/explicit. In a hierarchical cognitive system, knowledge that is explicit at one level is often hidden from operators at higher levels. In particular, knowledge that is fully explicit at the level of more elementary computational processes, and which is explicitly consulted by the microagents (“homunculi”) who work at that level, is inaccessible at the personal level. The only reason for describing this knowledge as “implicit” is that it is personally inaccessible. The only reason for describing some other knowledge as “explicit” is that it is personally accessible. There is a different, Rylean sense in which knowledge might be described as implicit. In this sense, a skilled performance is executed *as if* the operator is following a written recipe, except that there is not literally a recipe to be read. But mere as-if knowledge cannot explain any actual activity. If Level I knowledge is to explain activity, it had better not be intrinsically nonexplicit. It had better just be hidden from higher-level operators.

2. Procedural/declarative. A variety of distinctions might be marked with these terms. I doubt the relevance of most of them. The useful distinction is this: knowledge is “procedural” if it is only available for use in the execution of a specific procedure, that is, if it is proprietary to that procedure. Knowledge is “declarative” if it is available for verbal report, that is, for declaration. Again, the important facts are facts about the access that cognitive operators have to bits of knowledge.

In the beginning, knowledge is distributed among microagents, who have proprietary use of it. Subsequent cognitive development produces both articulation and integration. Articulation produces finer-grained partitions in the mind’s representation of things. Representational elements that are originally embedded in the control processes for independent skills, but which overlap semantically, are extracted from those processes and translated into more abstract formats that capture common (or overlapping) meanings. Articulation is necessary for the transportation of knowledge into new domains, for the development of the recombinatorial power that is at the core of higher thought and reasoning. Integration

occurs when independent (that is, articulated) representations are jointly available for a common task: that is, when they can appear together in the premises of a single inference. Under extreme idealization, the inferential integration of knowledge yields Fodor's central cognition (Fodor 1983). Inferential integration is what representational redescription achieves, above all else. It is a process whereby knowledge that was distributed among arrays of microagents becomes centrally available. It is knowledge that is (in principle) poised to enter inferential liaisons with any other knowledge contained in the central store.

3. Representational redescription. Is redescription a process that acts upon the vehicles of representation or on the semantic contents of those vehicles? Some domain-specific knowledge will be encoded in proprietary formats: for example, in analog formats especially suited to the computational requirements of specific sense modalities or specific motor domains. For that knowledge to be made available for general inferential use, it must be translated into the common, presumably symbolic format used in central cognition. This translation is a change in the vehicle of representation. Provided we leave aside general philosophical claims about the semantic indeterminacy of all translations, redescription need not be associated with any change to the content of the representation. If redescription is also generally accompanied by a loss of information (Karmiloff-Smith 1994a, p. 700), that is an interesting fact, perhaps telling us something about the costs of symbolic representation. But it is not essential to the redescriptive process. The main point is that encoding into a symbolic format is necessary for the inferential integration of knowledge.

4. Levels of knowledge. What is the relationship between the idea of knowledge that is articulated and inferentially integrated, and Karmiloff-Smith's idea of the several levels of explicit knowledge? The most interesting question is whether a significant measure of integration and articulation in the representation of knowledge can be achieved apart from the achievement of natural language competence. If it cannot, then we should expect to discover a lawful correlation between the possession of a capacity for cognitive flexibility and the possession of linguistic (or other symbolic) competence. What this would indicate is a constitutive connection between language competence and higher cognition.

Consciousness does not have the same explanatory value, however. To be avoided at all costs is the thought that availability to consciousness *explains* the cognitive flexibility that mature human beings display. This thought presupposes that we can help ourselves to the idea of a conscious, executive agent. On the contrary, by explaining how flexibility emerges from below, we might hope (in part) to explain conscious agency.

Beyond representational redescription

Fiona Spensley

Institute of Educational Technology, The Open University, Milton Keynes, MK7 6AA, England. m.f.spensley@open.ac.uk

Abstract: There are a number of elements in the representational redescription (RR) theory which elude definition, including behavioural success, implicit information, endogenous metaprocesses, and the detail of the representational levels. This commentary proposes an information processing approach to the development of cognitive flexibility – the Recursive Re-Representation (3Rs) model (Spensley 1995) – which re-defines the developmental process and thereby eliminates these problematic concepts.

Karmiloff-Smith (1994r) acknowledges that the lack of definition in the representational redescription (RR) model which has been identified by previous commentators, but expresses disappointment that these commentators have not offered any concrete suggestions to help specify the problematic concepts. This commentator will be even less helpful, by suggesting that the problems are insuperable, and concurring with Dartnall's (1994) concern

“that we have misconceived the problem and that we need to redescribe representational redescription” (p. 712). More positively, a possible reformulation of the problem is suggested: the nascent Recursive Re-Representation (3Rs) model (Spensley 1995) – an information processing approach to the development of cognitive flexibility. In this commentary the 3Rs approach is introduced in conjunction with a discussion of some fundamental problems with RR theory.

Implicit information and endogenous metaprocesses. In the RR model, the transition from level-I to level-E1 representations involves redescribing the information which is implicit in the initial opaque procedures into a more explicit form, by endogenous metaprocesses. However, as Campbell (1994) suggests, the concept of information being *represented* implicitly seems implausible. Karmiloff-Smith's block balancing task (1992, pp. 84–87) provides a clear illustration. Behavioural success in block balancing is achieved by purely proprioceptive feedback procedures. However, the phase 2 “theory” that blocks balance in the middle involves completely different concepts, that is, it involves representing the relationship between the elements “middle of the block” and the “point of balance.” The middle-of-the-block constituent would never be part of a proprioceptive feedback procedure. So, the centre-theory could never be extracted by endogenous processes operating over the original, successful, proprioceptive feedback procedures. The “theory” must have come from the *observation* of balancing behaviour and the generation of a novel representation of the observation that centered placement is apparently crucial to success. As Campbell observes, endogenous RR processes cannot create new information.

Schultz (1994) provides an example of a connectionist balancing program which supports Karmiloff-Smith's behavioural progression: the program tending to try balancing blocks in the middle after experience. However, this “central tendency” would not be *represented* in the network, and therefore would not be accessible to endogenous processes. Rather it would be implicit in the program's *behaviour*. To represent this explicitly, the re-representation must come from an analysis of *behaviour*, not of the network which generates the behaviour. It is this interactive approach that the 3R's model adopts – new representations being generated from the observation and analysis of the behaviour generated by previous representations.

Behavioural success. An important insight of Karmiloff-Smith's was that developmental change does not end with “success,” although developmental theories do not account for this. She is concerned to account only for the period of development beyond “behavioural success,” maintaining in the face of criticism (Goldin-Meadow & Alibali 1994) that “one researcher cannot do everything” (p. 737). However, Karmiloff-Smith acknowledges that the concept of behavioural success is extremely difficult to define. However, it may be problematic because it is not a useful concept, and it does not actually play a critical role in the development of cognitive flexibility.

There are two implicit claims behind the centrality of “behavioural success” in the RR model which need to be – but have not been – justified. The first is that the existence of post-success development requires the operation of a completely different developmental mechanism following behavioural success from that which precedes it. The second claim is that it is development beyond “success” which leads to our uniquely human flexibility and creativity. It is argued in a later section that cognitive flexibility is independent of behavioural success, and argued here that a single developmental re-representation process could take the child to success and beyond, eliminating the need to define the concept of behavioural mastery.

Karmiloff-Smith has clearly demonstrated that theories which stress the role of negative feedback cannot account for development beyond success. However, it does not follow that the negative feedback accounts are the correct explanation of *pre*-success development, which would then necessitate a different post-success theory. Boden (1982) has argued, in general, against the

notion that negative feedback is important for development at any stage. If negative feedback is not required to enable the child to achieve successful performance then the same processes which precede behavioural success in any domain could also take the child beyond it. Following Dartnall's (1994) concern with evolutionary validity, it does not seem plausible that a completely new mechanism evolved to modify representations following success. A more parsimonious account which covers both pre- and post-success representational development mechanisms must be *in principle* more convincing than two separate systems. The difficult problem of defining behavioural success is then not solved, but removed.

Dropping the concept of behavioural success would allow representational redescription to develop into a more generally applicable developmental theory. Although the suggestion, of course, changes the theory fundamentally and has implications for other theoretical concepts.

Phases and levels. Dropping the notion of success leads to theoretical problems with the three phases and four levels of the RR model. However, the empirical justification for the phases and levels has, indeed, been questioned by previous commentators. Most of the empirical evidence for RR theory relates to the transition from level-I to level-E1. Zelazo (1994) argues that the evidence for level-E1 representations is far from compelling and Vinter and Perruchet (1994) argue that there are many exceptions to the association of behavioural mastery with opaque procedures. This commentator has also failed to find supporting evidence for opaque procedural representations in replications of Karmiloff-Smith's drawing and block balancing tasks (Spensley 1990; 1995; in preparation; submitted). Children's behaviour seems to involve more flexibility and reflection than Karmiloff-Smith has suggested.

In theoretical terms Karmiloff-Smith has replaced Piaget's age-related stages with domain-dependent, age-independent phases. However, the idea of a sequence of structurally different representations has remained – and is very limiting. Karmiloff-Smith acknowledges the basic problem that the representational formats are underspecified – a point made by Bodor and Pléh (1994). However, there are additional problems related to the delimitation of the domains over which redescrptions are made. There may be a progression to flexibility within a micro-domain, and then again in a superordinate domain or across micro-domains. It is not clear at which stage the RR sequence occur, or whether it is repeated. Karmiloff-Smith (1992, p. 18) has prefixed her most recent description of the three phases with the word “recurrent,” although it is not clear to this commentator how repeated applications of the three phases would operate. It seems to require the passage from opaque representation to flexible representations and then back to opacity before a subsequent redescription sequence can be initiated. This seems unduly cumbersome.

The standard cognitive science solution to the problem of a limited sequence of levels would be to propose a *recursive* redescription process. This removes the specific problem of defining the starting conditions (i.e., behavioural mastery), although it creates others. A recursive process requires a more generally applicable redescription mechanism than the highly specific meta-processes alluded to in the RR model. It also necessitates the dropping of the sequence of *qualitatively* different representational formats. The 3Rs model suggests a developmental sequence of representations differing in content and compactness of representation rather than in structural format. A general representational mechanism which operates recursively could be the same one that creates the earliest representations, regardless of success or failure. The 3Rs model conceives of a process similar to Mandler's (1988; 1992) “perceptual analysis” mechanism.

Cognitive flexibility. As Karmiloff-Smith states, humans are special because they are able to use their knowledge flexibly. However, cognitive flexibility is not necessarily a consequence of development beyond behavioural mastery. The alternative approach in the 3Rs model is to conceive of cognitive flexibility as the

result of the availability of cognitive capacity. This approach satisfies Zelazo's (1994) concern for a broader conception of cognitive flexibility and inflexibility. Flexibility may be the result of having the cognitive space to reflect, rather than of having knowledge represented in a flexible, accessible structure. Development may proceed by knowledge being recursively re-represented into a more concise and compacted form, thus liberating processing capacity. These, of course, are the standard information processing concepts of a limited capacity central processor and the (ill-defined) notion of “chunking.” Redescription or re-representation may occur, as Zelazo points out, whilst the child is engaged in problem solving which requires the development of a new representation to achieve the solution. This certainly seems a more compelling motivation for redescription than a state of “stable success.”

Recursive re-representation. The Recursive Re-Representation (3Rs) model (Spensley 1995) is being developed on the basis of considerations outlined in this commentary. It offers a new perspective on the problem, but is not currently specified in the detail that the RR model has been. The major contentions are that cognitive flexibility is determined by the interaction of the capacity limitations of the central processor and the “compactness,” in terms of “chunks” of the representation. New representations are continually developed recursively through some kind of perceptual analysis process (Mandler 1988; 1992) utilising the same mechanism which would be responsible for the earliest representations. The basis of perceptual analysis is the recognition of analogies, which involves interpreting new information from the environment in terms of pre-existing representations (within or between domains) to re-represent the problem. This would be compatible with Olson's (1994) conception of representational redescription in terms of relating a new task to a model, and would be consistent with the restructuring observed by Bloom and Wynn (1994) in the domain of number.

The 3Rs model emerged from a cognitive scientist taking a developmental approach to the problem of cognitive flexibility – as advocated by Karmiloff-Smith in *Beyond modularity*. It has evolved through a detailed analysis (Spensley 1995) of Karmiloff-Smith's RR model – the *only* model of cognitive flexibility in the literature. The 3Rs model, whilst clearly not a variant of RR theory, has nonetheless evolved from it, and owes its existence to Karmiloff-Smith's pioneering work. The 3Rs model, at this stage in its development, seems to replace RR's ill-defined concepts with a further set. However, it is hoped that the information processing approach may ultimately offer a more parsimonious account of the development of cognitive flexibility.

ACKNOWLEDGMENT

The author would like to thank Richard Joiner, Mike Scaife, and Josie Taylor for useful discussions of this work.

Author's Response

Promissory notes, genetic clocks, and epigenetic outcomes

Annette Karmiloff-Smith

MRC Cognitive Development Unit and University College, London WC1H 0BT, England. annette@edu.ucl.ac.uk

Abstract: I respond to three continuing commentaries on *Beyond modularity*, two concerning the representational redescription (RR) framework and its attempts to account for the growing flexibility of human intelligence, and one relating to the putative mysteries of developmental timing. I discuss misunderstandings about the RR framework as well as some of its shortcomings. I

strongly reject the notion of a genetic clock and go on to argue for epigenetic outcomes in which genes and environment interact during the protracted period of postnatal brain development.

The three commentaries touch on two different but inter-related aspects of *Beyond modularity*. For me, the most important notion developed in the book is that of a gradual process of modularisation/specialisation where modular-like structures are argued to be the *outcome* of the developmental process, not its starting point. A second focus of the book was the notion of representational redescription (henceforth RR) with which I attempted to account for the growing flexibility of human intelligence that develops in parallel with the growing specialisation. I will first deal with the RR issue on which two of the commentaries touch (**Browne** and **Spensley**), and finally discuss at some length the putative genetic clock that **Bonatti** offers in preference to gene expression and epigenetic outcomes.

Much of **Browne's** commentary restates in somewhat different terminology the basic tenets of the RR framework developed in *Beyond modularity*. His notion of "articulation forming finer-grained partitions. . . and translated into more abstract formats" seems to correspond to the RR notion of extracting component parts, recoding them into new formats that can then be manipulated and transported elsewhere to form new intra- and inter-representational links which are not directly detectable in the environment. In Browne's view the overlap between redescribed representations is semantic. In my view, however, the cognitive system capitalises on multiple forms of overlap: semantic, indeed, but also syntactic, morphological, phonological, spatial, temporal, and structural overlaps.

Browne argues that the notion of consciousness does little explanatory work in the RR model. This is correct, but as I stressed in the book, I consider conscious access to internal representations to be the *outcome* of the process of RR, not as generating it. Consciousness was never meant to contribute directly to the actual process of RR; a form of conscious access to the products of internal processes was argued to emerge from it.

Browne's concern is how the knowledge created by RR is put to work. He worries about the status of implicit knowledge at the I-level. Note that I draw a distinction between *information in a system* (I-level), versus *knowledge to that system* (E-levels). Information at I-level is represented, but it is special-purpose, bracketed information and therefore its component parts are not accessible separately for other purposes. But Browne is right that the terms "procedural" and "implicit" can be misleading because of their uses in other theories. The term "implicit" has been bandied about without definition in the developmental literature for several decades. I believe that the RR framework provided the first definition of "implicit" in the developmental literature (see p. 20, *Beyond modularity*), and that it can help to explain why children sometimes seem to have knowledge that they cannot use outside the special-purpose processes in which it is embedded.

I deeply regret the use of the term "procedure," however, because it has given rise to so many misunderstandings. Throughout the formulation of the RR framework, I was grappling with a notion of information represented in a system whose component parts were not accessible because bracketed within a closed process which I called a "proce-

dure." The procedure contained the set of actions required to produce an output in language or problem solving, and (contra **Spensley's** criticisms about representations containing only proprioceptive information) obviously contained also the result of that output ("the sensitivity to information emanating from observable data" p. 84, *ibid.*; "from redescrptions of their stored level-I representations of objects that balance, children extract a common feature" p. 86, *ibid.*).

Spensley argues that actions are not represented. I disagree. The mind represents both the actions it performs and the product of those actions, that is, the resulting balanced state. My argument was that the child might use a similar set of actions leading to the same resulting state but that while these were still at I-level, the potential representational links between resulting states could not be capitalized on. Subsequently, new knowledge can emerge from transformations of and new links between what we already know; we can enrich our minds from within. New information does not come solely from observing things happening in the world. This would be pure behaviourism. Children go beyond collecting facts about the world. They develop theories on the basis of their internal representations through a process which is basically abductive, not inductive (Karmiloff-Smith 1988; O'Loughlin 1995). New representations allow one to notice things in the world that went completely unnoticed before. Children (and adults) will ignore clearcut information in the world until their cognitive systems are in the right state for the information to become relevant. They will also create observables that are not actually in the world to meet their theoretical commitments.

I of course agree with **Browne** that the simple procedural/declarative dichotomy is restrictive. I make that point repeatedly, and it is why I hypothesized several levels of redescription. My focus was on the status of internal representations, rather than the uses to which they are subsequently put. Yet Browne is correct that explicitly defined representations which are not tied to special-purpose procedures would become available for inferential integration. I agree that loss of information is not essential to the redescriptive process. But recoding often cannot carry over all the details of modality-specific information, for example, the recoding of spatial information into linguistic form. That is why the resulting re-representations often emerge as more abstract. It may well be that the system is not driven towards more abstractness, but that abstractness is the mere product of constraints on what can be carried over from one representational format to another.

Finally, **Browne** raises questions about the relation between natural language competence and higher cognition. This of course depends on how one defines "higher cognition," for it is clear that language does indeed enhance our cognitive capacities, a point that Piaget willingly conceded. However, I argue against the idea that language is essential to representational redescription which makes higher cognition possible in the first place. First, as reported in *Beyond Modularity*, there are cases of relatively elaborate language – syntactically and lexically – that are not accompanied by higher levels of cognition, as well as cases of very impaired language coexisting with higher levels of cognition. So language and cognition can display relative dissociations during development. Second, some of the concep-

tual achievements found in the prelinguistic infant would also mitigate against a constitutive relation between language and cognition. No other species has anything like language in the human sense. However, no species other than the human seems to display any signs of RR and, as I argued and as shown by Mandler's (1992) research, RR does occur in prelinguistic infants. Language is not a precondition for redescription, but I have consistently stated that later in development both spoken and written language greatly enhance the redescriptive possibilities of the cognitive system (Karmiloff-Smith 1979; 1992, Ch. 6; Lee & Karmiloff-Smith 1996a; 1996b).

Spensley raises a number of other issues in her important commentary. I disagree with her with respect to her rejection of connectionist network's capacity to represent information implicitly. In fact, they are ideally suited to do so in contrast to classical symbolic models. This can be demonstrated by carrying out principal components analyses which show how the networks represent common features (e.g., centrality, Shultz 1994; nounhood, plurality, etc., Elman 1993) along a series of representational trajectories which are different from information on other trajectories (e.g., those representing right/left side, verbhood, singularity, etc.). However, the representations are a function of the inputs that they process over time in a particular network. As yet networks do not re-represent that information in a format that could be used directly by other networks. Hence my criticism that something like RR is necessary to make the potential representational relationships on the common trajectories available for other networks in a way that is not constrained by the particular inputs being processed. Take an abstract concept like "opposite." This could emerge from different networks processing different kinds of input, but to be seen as a common concept across the networks, it would have to be represented in a format that was not constrained by the processing of local input vectors.

Spensley next questions the usefulness of the notion of behavioural mastery. In the book I mention the fact that there is an important difference between *overt* behaviour (which until the recent availability of brain imaging techniques was all that was available to the experimental psychologist) and internal processes. Connectionist networks indeed show that the hidden layer starts to represent, say, nounhood, before this is apparent in the behaviour at the output level. The need to qualify the concept of behavioural mastery was already recognised in the book. My more recent studies on children with learning difficulties show that they can reach behavioural mastery without ever going beyond success. So behavioural mastery is clearly not sufficient for representational change and, as Spensley suggests, it may not even be necessary. Yet I fail to see how dropping the constraint of behavioural mastery changes in any fundamental way the basic notion of representational redescription, particularly as Spensley seems to accept success-based developmental change.

I have never denied that exogenous constraints can also generate RR. The original idea sprang from the repeated discovery that children go beyond success; this needed to be explained in the light of the then prevailing developmental models. The rejection of failure-based models of change was first developed in my work on language (Karmiloff-Smith 1979a) and on microdevelopmental details of the use of other representational systems (Karmiloff-Smith 1979b).

In both of these cases, I challenged the prevalent focus on negative feedback for learning and showed how positive feedback can also generate change. However, unlike Boden's (1982) subsequent arguments, cited by **Spensley**, that positive feedback alone explains all change, my argument has been that negative feedback plays a more important role at the earlier stages of learning about a new domain than at the later stages of consolidating and changing that knowledge. So my view is that *both* negative and positive feedback are essential *at different moments in developmental cycles*. I do not argue for a new mechanism arriving at a late age in development. This is a confusion between stage-like ontogeny and reiterative processes throughout ontogeny (data-driven processes followed by RR during microdevelopment or over longer periods of macrodevelopment, depending on the domain).

Recall that I have shown the processes to exist after success in microdevelopment also (Karmiloff-Smith 1979b). With Mandler, I also claimed that there was no reason why RR might not occur from early infancy onwards. A close comparison of Mandler 1988 (two parallel processes) and Mandler 1992 (a redescriptive process) suggests that she also subsequently opted for the process of perceptual analysis to operate on analogies extracted from previous *representations*, not directly from external actions on the world. When I stated that the process was "recurrent" I certainly did not intend to imply that the child's same *representations* return to opacity in a particular microdomain. It is the *process* of RR that is deemed recurrent (perhaps the more correct term is indeed recursive), not the resulting representations of course.

Like **Spensley**, a number of researchers have now shown that flexibility in drawing occurs at younger ages than I originally claimed (see, for example, Zhi et al. 1997; the references to her own work that Spensley cites are unpublished or "in preparation," so I cannot comment on their empirical value here). But this does not detract from the general importance of pinpointing the mechanisms of the growing flexibility in human cognitive development. While the sequential constraints of the drawing task were replicated with some but not all children (Zhi et al. 1997), the growing intra- and inter-representational flexibility has held. My focus was not on drawing, but on the general endeavour to account for two apparently contradictory facts about human development. If a system became progressively more specialised and automatised, how could it simultaneously become more flexible? My response was to invoke a gradual process of modularisation accompanied by a parallel process of representational redescription.

Spensley assures us that her own recursive representational redescription model (the 3-Rs model) obviates a number of the problems that she raises with respect to the RR model. There are sparse details of the 3-Rs model in her commentary and she herself states that it is still less developed than the RR framework, so at this stage it is impossible for me to comment theoretically either. The little that is said suggest that the 3-Rs model is a combination of RR and neo-Piagetian theories (Case 1987; Fischer 1980; Halford 1987; Pascual-Leone 1972).

Although **Spensley** challenges the necessity for behavioural mastery which I myself questioned in the book, she retains the central RR notion of redescription of representations into more a compact form, but wishes to do away with the constraints that *different* representational formats

might have on what information can be brought into relation with others. I welcome the opportunity to assess the new perspectives that Spensley's 3-Rs model promises, and am delighted that my own work has been a springboard for these endeavours. Spensley is clearly addressing a number of crucial developmental questions, but for the purposes of this commentary, she offers but a tantalising promissory note.

To turn now to the commentary by **Bonatti**. His major point is about timing: that the putative developmental clock is a biological clock with the environment acting as a mere trigger. This is a truly nativist position, implying that the infant brain contains at birth a large set of innately specified, domain-specific patterns of connectivity available prior to any form of learning or experience. And since his concern is language, one of the higher human cognitive processes, he must mean that the prespecified patterns of connectivity are cortical in nature.

Bonatti claims that normal neonates distinguish linguistic from acoustic nonlinguistic signals. Correct, but the implication that this is proof that the discrimination is based on an innately specified language module does not follow. First and foremost, infants have the capacity to learn very rapidly and such learning starts in the womb (see Karmiloff-Smith 1994; 1995). Infants are sensitive to the *abstract* structure of visual and auditory patterns in the environment and not solely to surface features. This holds across several modalities after birth, but even for the foetus in the auditory modality.

In the final trimester of interuterine life, infants show sensitivity to different kinds of auditory input to which they are exposed. In utero, they distinguish language from other auditory stimuli (Lecanuet et al. 1988; Shahidullah & Hepper 1993; Wilkin 1991). Such data could of course be given a nativist interpretation, but it is a far from necessary conclusion in the light of other data. For instance, foetuses can be trained to recognise and gradually habituate to a repeatedly heard story or a piece of Mozart and will prefer to listen to the familiar story/music ex utero (Hepper et al. 1993). Other experiments have shown that, despite the profound differences between the mother's voice heard in utero filtered through the amniotic fluid and her voice as heard ex utero once the infant is born, recognition of the maternal voice is learned prenatally. Although young infants prefer to listen to mother's voice as it sounded in utero compared to its sound ex utero (Fifer & Moon 1988), they can nonetheless discriminate at birth mother's voice heard ex utero from that of other females (Hepper et al. 1993). In other words, prior to any ex utero experience, newborns show that they have already extracted information about some of the basic features of mother's voice during their period in the womb. The abstract components of a *specific* mother's voice obviously cannot be prespecified; they have been acquired by an active, fast-learning foetus.

Bonatti makes a great deal of the fact that neonates show categorical perception of speech sounds. But this is irrelevant to the innateness of language argument. First, let us not forget that we are talking about *speech* processing capacities, not *language*. One cannot jump directly from categorical perception to semantics and syntax! Yet the nativist literature constantly slips between the two. Second, the infant capacity for speech discrimination is not species-specific, because as Kuhl (1991) has shown, chinchillas and other species also display categorical perception of human

speech sounds. But clearly the chinchilla has no modular (or even nonmodular) capacity for language! Categorical perception of speech sounds might be innate, but, if so, it is a cross-species biological capacity for acoustic discrimination on which human languages subsequently capitalised for the distinctions in their sound systems. Perhaps this well-developed auditory capacity makes infants particularly attentive to linguistic input, but the jump from categorical speech perception to language is a huge one.

The fact that children tend to lose the sensitivity to some contrasts irrelevant to their natural language at somewhere between 10 and 12 months is used by **Bonatti** as evidence of innately specified timing and modularity. By contrast, I see it as a particularly nice example of the *gradual postnatal specialisation* of the brain (a gradual process of modularization), as a function of the particular inputs it processes over the early months of life. First, the loss of sensitivity to non-native sounds co-occurs with a substantial increase in sensitivity to the particular phonemic boundaries of the infant's own native language, suggesting that language-specific specialisation and non-native loss are two sides of the same learning process. Second, recent work in our London lab (Rivera-Gaxiola 1996) has shown that although adults lose the *overt behavioural* capacity to discriminate non-native speech contrasts, ERP recordings make it clear that their brains continue to process them. By extension we hypothesise that infant brains also continue to process non-native contrasts beyond 10 months. It is simply that their progressive specialisation with their mother tongue stops them from demonstrating this at the behavioural level.

Bonatti's next claim is about the purported lack of individual variation. First, it is important to recall that *variation between individuals* is fundamental to the dynamics of any system's survival and development. Evolution crucially depends on variation. Yet Bonatti's argument hangs in part on his strong claim that there is no significant variation across individuals in acquiring language and other skills. That belief is not consistent with the empirical data. Large-sample studies have yielded evidence for huge variability in perfectly healthy children.

Staggering individual variance was found in a study by Fenson et al. (1994) of early language acquisition in 1800 children between 8 and 30 months. As an example, at 12 months there are children who understand hundreds of words, while others still show no signs of understanding more than "no" and their own name. At 24 months there are children who can produce more than 600 words and speak in paragraphs, while others of the same age, who are perfectly healthy and understand a lot but can still produce only a handful of well-formed single words. So substantial individual variation is the norm rather than the exception, and this holds for both lexical and syntactic development. Furthermore, behaviour genetic studies (e.g., Reznick et al. 1997) of expressive language, receptive language, and non-verbal skills across the second year of life in over 400 monozygotic and dizygotic twin pairs have also shown that there is a large and significant amount of variation that can only be explained by environmental factors. In fact, Reznick reports that the environmental effects are much stronger for the language measures than for nonverbal skills. Note, also, that by adulthood even the brains of monozygotic twins are not identical.

I am not suggesting that nothing is innate. Clearly there are some predispositions that channel learning of certain

types of input and constrain the representations that can subsequently emerge. But these predispositions are rarely representational, that is, there seem to be no fine-grained patterns of cortical connectivity laid out in advance, at least not at the synaptic or dendritic levels (Elman et al. 1996). Where and in what form, for instance, is the “theta principle” stored in the neonate brain? Surely it must be a cortical constraint? Yet no nativist has addressed such questions in the type of detail required to make their claims go beyond sheer belief. The protracted period of postnatal brain growth in humans allows for plasticity and for multiple influences from the interaction of a progressively reconfigured brain and the environmental input. The nativist seems to ignore what we now know about embryonic and postnatal brain development. Plasticity is the rule, not simply an exceptional response to pathology.

Rather than jump to strong nativist conclusions, a better scientific strategy in my view is to attempt to uncover the *minimal* predispositions that enable the infant to learn rapidly and progressively structure its own cortex, since the protracted period of human postnatal brain development allows it to do just that. Even in the motor domain of learning to walk, it has been shown that nothing is uniquely prespecified, but is always the result of interactions at multiple levels within the system itself and with the external environment (Thelen & Smith 1994). Obviously there is a universal sequence: children sit before they crawl, and crawl or cruise before they walk. And they walk with their legs wide apart before they walk like older children and adults. But these sequences have no need to be genetically specified; they are the natural solution to the similarity of the structure of the problem space and the outcome of multiple interactions as muscular strength progressively increases from the activity of the infant herself (Thelen & Smith 1994).

It is crucial to draw a distinction between dissimilarity of environmental conditions, on the one hand, and similarity of the structure of the problem space, on the other. **Bonatti** notes the striking similarities between the acquisition of sign language and of spoken language despite very different environmental input. But is the situation so different? The modalities differ indeed, but the problem of mapping multi-dimensional meanings onto a fast-fading, linear signal in real time is almost identical in the two cases. It is therefore not surprising that young children come to similar solutions at similar times due to factors that are intrinsic to the problem space of language and also influenced by growing short term memory capacities and the like. Furthermore, the general problem of many-to-one dimension reduction is a general one, also involved in mapping 3-dimensional visual inputs onto 2-dimensional retina-like representations (Churchland & Sejnowski 1995). Solving such mapping problems are the product of multiple factors, not triggered by a fixed genetic clock.

Another important indication of the gradual specialisation of the brain for language comes from recent work (Mills et al. 1993) showing that the left hemisphere bias for known words does not appear until well into the second year of life. Increasing language ability, and not chronological age, turned out to be highly associated with increasing cerebral specialisation for language processing over the temporal and parietal regions of the left hemisphere. Prior to this, young infants start with a bilateral response and only gradually specialise to the left hemisphere. Both of these

findings are precisely what the notion of progressive modularisation would predict.

Finally **Bonatti** claims that all normal children pass “the (sic) false belief task” (which single task does Bonatti have in mind?) at around 4 years of age, but not a few months before. It must again be recalled that these are *group data* with a normal distribution showing that some children of 3 do indeed pass but others fail until about 5 years of age. Since I wrote *Beyond modularity* an abundance of theory-of-mind tasks have been devised. Age of success has now been shown to depend on the way in which the task is presented, and also on whether the child has siblings or not (Perner et al. 1994). More recent research has uncovered children’s ability to give correct responses in terms of their looking behaviour some 6 months earlier (again group data), even when their verbal responses are not yet correct (Clements & Perner 1994). Such findings have none of the hallmarks of an innately specified module or a genetically timed clock, but lend themselves to an interpretation in terms of a combination of both endogenous and exogenous factors.

Of course I do not claim that there are no cognitive changes due to maturational changes in brain structure (see, for example, discussions in Johnson 1994; 1995), but to my knowledge, there are no neuroscientific data showing that there are prespecified representations in the cortex allowing for success on theory-of-mind tasks. Rather, it is possible that until the frontal region of the brain has become linked to other regions of the brain, certain types of computation (e.g., holding two competing representations in mind) are not possible. But such changes do not appear to be theory-of-mind specific. They affect computations across several areas of cognition.

In sum, the case for innately specified modules and a genetically timed clock has not been made. I believe that the data – psychological and neuroscientific – make a better case for two parallel processes in human development: progressive epigenetic outcomes, that is, a *gradual* process of modularization and progressive cognitive flexibility via some process of representational redescription.

References

- Bertoncini, J., Bijelac-Babic, R., McAdams, S., Peretz, I. & Mehler, J. (1989) Dichotic perception of laterality in neonates. *Brain and Language* 37:591–605. [LB]
- Bertoncini, J., Floccia, C., Nazzi, T., Miyagishima, K. & Mehler, J. (1995) Morae and syllables: Rhythmical basis of speech representation in neonates. Unpublished paper. [LB]
- Best, C. T. (1988) The emergence of cerebral asymmetries in early human development: A literature review and a neuroembryological model. In: *Brain lateralization in children*, ed. D. L. Molfese & S. J. Segalowitz. Guilford Press. [LB]
- Bijelac, R., Bertoncini, J. & Mehler, J. (1993) How do four-day-old infants categorize multisyllabic utterances? *Developmental Psychology* 29:253–69. [LB]
- Bloom, P. & Wynn, K. (1994) The real problem with constructivism. *Behavioral and Brain Sciences* 17(4):707–8. [FS]
- Boden, M. (1982) *Failure is not the spur (CSRP 015)*. University of Sussex. [FS, AKS]
- Bodor, P. & Pléh, C. A. (1994) A Fodorian guide to Switzerland: Jung and Piaget combined? *Behavioral and Brain Sciences* 17(4):709–16. [FS]
- Campbell, R. L. (1994) What’s getting redescribed? *Behavioral and Brain Sciences* 17(4):710–11. [FS]
- Case, R. (1987) The structure and process of intellectual development. *International Journal of Psychology* (special issue: The neo-Piagetian

- theories of cognitive development: Toward an integration) 22:571–607. [AKS]
- Churchland, P. & Sejnowski, T. (1995) *The computational brain*. MIT Press. [AKS]
- Clements, W. A. & Perner, J. (1994) Implicit understanding of belief. *Cognitive Development* 9(4):377–96. [AKS]
- Dartnall, T. (1994) Redescribing redescription. *Behavioral and Brain Sciences* 17(4):712–13. [LB, FS]
- De Gelder, B. (1994) The risks of rationalizing cognitive development. *Behavioral and Brain Sciences* 17:713–14. [LB]
- Elman, J. L. (1993) Learning and development in neural networks: The importance of starting small. *Cognition* 48:71–99. [AKS]
- Elman, J., Bates, E., Johnson, M. H., Karmiloff-Smith, A., Parisi, D. & Plunkett, K. (1996) *Rethinking innateness: Connectionism in a developmental framework*. MIT Press. [AKS]
- Estes, D. (1994) Developmental psychology for the twenty-first century. *Behavioral and Brain Sciences* 17:715–16. [LB]
- Fenson, L., Dale, P. S., REznick, S. J. & Bates, E. (1994) Variability in early communicative development. *Monographs of the Society for Research in Child Development* 59(5):173. [AKS]
- Fifer, W. P. & Moon, C. (1988) Auditory experience in the fetus. In: *Behavior of the fetus*, ed. W. P. Smotherman & Scott R. Robinson. Telford Press. [AKS]
- Fischer, K. W. (1980) A theory of cognitive development: The control and construction of hierarchies of skills. *Psychological Review* 87:477–531. [AKS]
- Fodor, J. A. (1983) *The modularity of mind*. Bradford/MIT Press. [DB, LB]
- Foster-Cohen, S. H. (1994) Arguments against linguistic “modularization.” *Behavioral and Brain Sciences* 17:716–17. [LB]
- Goldin-Meadow, S. & Alibali, M. W. (1994) Do you have to be right to redescribe? *Behavioral and Brain Sciences* 17(4):718–19. [FS]
- Halford, G. S. (1987) A structure-mapping approach to cognitive development. *International Journal of Psychology* (special issue: The neo-Piagetian theories of cognitive development: Toward an integration) 22:609–42. [AKS]
- Hepper, P. G., Scott, D. & Shahidullah, S. (1993) Newborn and fetal response to maternal voice. *Journal of Reproductive and Infant Psychology* 11:147–53. [AKS]
- Johnson, M. H. (1994) Brain and cognitive development in infancy. *Current Opinion in Neurobiology* 4:218–25. [AKS]
- (1995) The development of attention: A cognitive neuroscience approach. In: *The cognitive neurosciences*, ed. M. S. Gazzaniga. MIT Press. [AKS]
- Karmiloff-Smith, A. (1979a) *A functional approach to child language*. Cambridge University Press. [AKS]
- (1979b) Micro- and macrodevelopmental changes in language acquisition and other representational systems. *Cognitive Science* 3:91–117. [AKS]
- (1987) The child is a theoretician, not an inductivist. *Mind and Language* 3(3):183–95. [AKS]
- Karmiloff-Smith, A. (1992) *Beyond modularity: A developmental perspective on cognitive science*. MIT Press. [DB, LB, FS, AKS]
- (1994) *Baby it's you: A unique insight into the first three years of the developing baby*. Ebury Press/Random House. [AKS]
- (1994r) Author's response to commentary: Transforming a partially structured brain into a creative mind. *Behavioral and Brain Sciences* 17:732–40. [LB, FS]
- (1994t) Précis of *Beyond modularity: A developmental perspective on cognitive science*. *Behavioral and Brain Sciences* 17:693–746. [DB, LB]
- (1995) Annotation: The extraordinary cognitive journey from foetus through infancy. *Journal of Child Psychology and Child Psychiatry* 36(8):293–313. [AKS]
- Kuhl, P. (1991) Ontogeny and phylogeny of speech. In: *Plasticity of development*, ed. S. Brauth, W. Hall & R. Dooling. MIT Press. [AKS]
- Lecanuet, J. P., Granier-Deferre, C. & Busnel, M. (1988) Fetal cardiac and motor responses to octave-band noises as a function of central frequency, intensity and heart rate variability. *Early Human Development* 13:269–83. [AKS]
- Lee, K. & Karmiloff-Smith, A. (1996a) The development of cognitive constraints on notations. *Archives de Psychologie* 64:3–26. [AKS]
- (1996b) The development of external symbol systems: The child as a notator. In: *Cognitive development*, ed. R. Gelman & T. Kit-Fong Au. Academic Press. [AKS]
- Losonsky, M. (1994) Beyond methodological solipsism? *Behavioral and Brain Sciences* 17:723. [LB]
- Mandler, J. M. (1988) How to build a baby: On the development of an accessible representational system. *Cognitive Development* 3:113–36. [FS, AKS]
- (1992) How to build a baby: 2. Conceptual primitives. *Psychological Review* 99(4):587–604. [FS, AKS]
- Mehler, J., Jusczyk, P. W., Lambertz, G., Halsted, N., Bertoncini, J. & Amiel-Tison, C. (1988) A precursor of language acquisition in young infants. *Cognition* 29:143–78. [LB]
- Mills, D. L., Coffey-Corina, S. A. & Neville, H. J. (1993) Language acquisition and cerebral specialization in 20-month-old infants. *Journal of Cognitive Neuroscience* 5(3):317–34. [AKS]
- O'Loughlin, C. G. (1995) The child as a realistic scientist: Developing the analogy between science and commonsense. MA thesis, University of Canterbury. [AKS]
- Ohlsson, S. (1994) Representational change, generality versus specificity, and nature versus nurture: Perennial issues in cognitive research. *Behavioral and Brain Sciences* 17:724–25. [LB]
- Olson, D. R. (1994) Where redescriptions come from. *Behavioral and Brain Sciences* 17(4):725–26. [FS]
- Pascual-Leone, J. (1972) A theory of constructive operators, a neo-Piagetian model of conservation, and the problem of horizontal decalages. Unpublished manuscript, York University, Toronto. [AKS]
- Perner, J., Ruffman, T. & Leekam, S. R. (1994) Theory of mind is contagious: You can catch it from your sibs. *Child Development* 65:1228–38. [AKS]
- Petitto, L. A. & Marentette, P. F. (1991) Babbling in the manual mode: Evidence for the ontogeny of language. *Science* 251:1397–1536. [LB]
- Quartz, S. R. & Sejnowski, T. J. (1994) Beyond modularity: Neural evidence for constructivist principles in development. *Behavioral and Brain Sciences* 17:726–27. [LB]
- Reznick, J. S., Corley, R. & Robinson, J. (1997) A longitudinal twin study of intelligence in the second year. *Monographs of the Society for Research in Child Development*. 62(1)1–162. [AKS]
- Rivera-Gaxiola, M. (1996) *Electrophysiological correlates of the development of speech perception in adults and infants*. Poster presented at the Second European Conference: Development of Sensory, Motor and Cognitive Abilities in Early Infancy, Spain. [AKS]
- Shahidullah, S. & Hepper, P. G. (1993) Frequency discrimination by the fetus. *Early Human Development* 36:13–26. [AKS]
- Shultz, T. R. (1994) The challenge of representational redescription. *Behavioral and Brain Sciences* 17(4):727–28. [FS, AKS]
- Spensley, M. F. (1990) Representational redescription and children's drawings. Paper presented to the BPS Cognitive Psychology Section Conference, Leicester. Published as *C.I.T.E. Report No. 126*, The Open University. [FS]
- (1995) Representational redescription and the development of cognitive flexibility. PhD thesis, The Open University. [FS]
- (in preparation) Representational redescription and the block balancing task. [FS]
- Spensley, M. F. & Taylor, J. (submitted) Representational redescription: Confounding evidence from children's drawings. [FS]
- Thelen, E. & Smith, L. B. (1994) *A dynamic systems approach to the development of cognition and action*. MIT Press. [AKS]
- Vintner, A. & Perruchet, P. (1994) Is there an implicit level of representation? *Behavioral and Brain Sciences* 17(4):730–31. [FS]
- Werker, J. F. & Tees, R. C. (1984) Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7:49–63. [LB]
- Wilkin, P. E. (1991) Prenatal and postnatal responses to music and sound stimuli: A clinical report. *Canadian Music Educator* (research ed.) 33:223–32. [AKS]
- Xu, F. & Carey, S. (1996) Infants' metaphysics: The case of numerical identity. *Cognitive Psychology* 30(2):111–53. [LB]
- Zelazo, P. D. (1994) From the decline of development to the ascent of consciousness. *Behavioral and Brain Sciences* 17(4):731–32. [FS]
- Zhi, Z., Thomas, G. U. & Robinson, E. J. (1997) Constraints on representational change: Drawing a man with two heads. *British Journal of Developmental Psychology*. 15(3)275–90. [AKS]