

Mememes as Signs in the Dynamic Logic of Semiosis: Molecular Science meets Computation Theory

Terrence W. Deacon

Department of Anthropology and Neurosciences Program
University of California, Berkeley

Abstract The concept of meme misidentifies units of cultural information as active agents, a shorthand similar to what misleads our understanding of genes and obscures the dynamic logic of evolution; but the concept does offer hope by contributing something missing from many semiotic theories. In treating memes as replicators, Dawkins fails to distinguish mere patterns from information (and replication from interpretation), which leads to the problem encountered in all realms of information processing: what counts as information is context dependent. Nothing is intrinsically meaningful, to be so it must be interpreted. In the evolution of both genes and words, replication has always been a multilevel affair in a dynamic system, from which what we conceive as "adapted" or "interpreted" units emerge. Memes are replicas not replicators, and I suggest that the iconic function of signs, as identified by Peirce, is the essence of the meme concept. As in sign function, both gene and meme functions are informed by whatever relationship exists between the physical pattern of the token and the system of processes in which each is embedded (so these are semiotic relationships). I argue that two, not-clearly-articulated aspects of the meme concept could rescue semiotic from mere descriptive taxonomy and lead the way for a general theory of semiosis and a unifying methodology for the semiotic sciences to emerge.

1 Introduction: The Logic of Replicators

Few scientific terms introduced into scientific and popular vernacular have enjoyed the impact on intellectual debate as has the term "meme." The meme concept, introduced by Richard Dawkins in his influential book *The Selfish Gene*, has spawned a vast following of "memetics" enthusiasts and not a few books expanding on Dawkins' original insight at great length. This popularity owes much of its widespread appeal to Dawkins' corollary characterization of genes as "replicators" and his reformulation of evolutionary theory from a "gene's-eye" perspective. His analysis of replicator selection in biology led directly to his proposal that one could also think of social-communicative processes as though there were a gene-like unit of replicator selection driving the evolution of cultures: the meme.

The meteoric rise of molecular biology to preeminent status as the epitome of cutting edge scientific-technical progress of the last half of the 20th century has only been matched by that of the information sciences. The serendipitously parallel explosion of knowledge in these two fields has provided the perfect context for Dawkins' reframing of evolution in molecular-informational terms. Dawkins, following the lead of such pathbreaking biological theorists as W. D. Hamilton, G. C. Williams, and E. O. Wilson, shifted the focus of evolutionary thinking to genes themselves as the ultimate unit of biological information and as the locus of natural selection, as opposed to traits (phenotypes), whole organisms, or populations, as the

then current evolutionary theories assumed. Organisms were characterized by Dawkins as "lumbering robots" designed by the genes to carry out their (the genes') reproductive imperatives, because after reproducing most organisms perish. Even offspring in sexually reproducing species are not copies. In contrast, genes, or more precisely the sequences of information embodied by them, are conserved from generation to generation. They are passed down minimally modified, if at all, across countless generations. In this way gene sequences retain the possibility of near immortality. Evolution thus appears to be about preservation of gene information at the expense of all else.

According to Dawkins, the critical defining feature of a gene that makes evolution possible is its ability to make precise copies of itself. Focusing on this remarkable molecular predisposition, Dawkins described genes as "replicators": units of information capable of copying themselves and also capable of influencing the probability of being copied in other indirect ways. He argued that this propensity is the most fundamental factor in biological evolution. Evolution could thus be reduced to a kind of molecular information processing based upon this one basic operation — molecular biology meets computation theory.

The patterns exhibited in evolution, then, could be described as merely playing out the statistics of replicators competing with one another. Replicators that, besides merely copying themselves, also happened to exert some additional positive influence on the probability of their own replication, thus biasing the statistics, would be favored over those with less positive effect or none at all; hence genes could be said to exhibit a kind of "selfishness." There are many ways that such biasing can occur. Virus genes produce structures that fool other organisms into treating the virus genes as their own and copying them willy-nilly; genes in multicellular bodies produce structures that increase the probability that the organism will reproduce and thus likely pass on copies of these genes; and sometimes these organism behaviors may even increase the probability that kin will reproduce, and pass on copies of both gene sequences inherited from parents. The "genes-eye-view" of evolution, then, was supported by the recognition that in all these cases, making copies of genetic information, not organism structure, was paramount. Everything else tended to be sacrificed before gene fidelity, fecundity, and longevity in evolution.

Framing this core feature of life in such general terms, also led Dawkins to consider (almost as an afterthought) that genes might not be the only potential replicators in the world. Cultural information capable of being passed on to others by mimicry or copying could also qualify. He named this sort of replicator a mimeme (unit of mimesis) or meme for short. Any cultural artifact or habit that could be transmitted by being copied could qualify as a meme. Dawkins offers such examples as popular songs, fads, mannerisms, recipes, tools, etc., and other memeticists have even suggested that languages, religious traditions, foraging techniques, and folk stories (to name just a few of thousands of suggested examples) might be considered memes. By Dawkins' own reckoning over two decades later the term "meme" has indeed exemplified itself in being propagated at an unusual rate for a newly coined jargon term [1: xiv]. In his search of mentionings of the term on the world-wide web, he noted half a million mentions as of August 1998, out-reproducing by orders of magnitude such recently coined jargon words as "sociobiology" and "spin-doctor."

2 Memetic Difficulties

But memes have not won wide acceptance as a useful theoretical construct in most fields where they might be relevant. Instead, with a few exceptions they remain a sort of analogical stand-in for concepts not fully fleshed out and so unsuitable for precise analysis. In a recent overview, Dawkins and Blackmore offer three commonly cited areas of disagreement and unresolved definitional problems regarding memes [1: x-xvi, 53-66].

(1) Memes appear to have insufficient copying fidelity to be evolvable. This problem long troubled geneticists prior to the discovery of the molecular basis of genetic information, because if the units of inherited information could mix or grade into one another evolution would quickly regress toward the mean and grind to a halt. Information would be rapidly eliminated and subtle deviations never accumulated.

(2) Nobody really knows what a meme physically is. Dawkins and others talk about memes as information. And in many contexts writers often substitute the term "idea" for meme, since memes are defined as whatever is passed on when something is copied by one person from another, either directly or indirectly. This is evident when writers describe memes with such phrases as "ideas begin to 'have a life of their own'" [1: 29] and "a unit of information residing in a brain" [2: 47]. And is also made evident by the plethora of near synonymous terms (e.g., cultural traits, culturgens, corpuscles of culture) and the disagreements as to whether only copied information qualifies or whether any transmissible cultural feature qualifies.

(3) There is difficulty agreeing on how large a unit deserves the name "meme." Is it as small as a word or a morpheme (root meaningful unit constituting a word) or a type of handle on a pot, or as large as a language and a religious tradition?

(4) Finally, some have questioned whether there is a memetic parallel to the genotype/phenotype distinction in biology. For example, is a recipe the meme and the cake the memetic phenotype; is a performance of a song the phenotype of the remembered idea of the song?

Both Dawkins and Blackmore summarize and address these problems in a recent discussion [see 1]. They conclude that all are in one way or another unresolved, but that they are either non-problems and can be ignored or else can be shelved for the time being, awaiting some memetic Watson and Crick to sort out their material basis.

I am far from convinced by these responses, but won't attempt to recount them or systematically critique them here, and argue only that we cannot shelve these problems and still have a science of memetics. To do so abandons this concept to remain the subject of pop science and internet discussion groups, but not of technical research. Though some reviewers would recommend this, I think there is a baby somewhere in this cloudy bathwater. To rescue it I wish to focus on a couple of tacit assumptions that I think underlie these persistent problems and to suggest a change of perspective that may help clear up what exactly we are talking about.

3 What To Do

The trouble with memes, as I see it, is that they have been misidentified. I am convinced that there is indeed something that bears social-cultural information and which evolves by means that are roughly analogous to processes of natural selection acting on DNA sequences and the living processes linked to them. That something is not, however, ideas. It is not something residing in brains. And it is not something imbued with the power to affect its own replication (but then again, neither is a naked strand of DNA, except in very special laboratory processes called PCR). It is not some new class of entity hitherto unrecognized, unstudied, and unnamed.

To state my conclusion at the outset: Memes are signs, or more accurately, sign vehicles (or representamens, to use Charles Peirce's clumsy but precise terminology). They are not some intangible essence that is passed from brain to brain when something is copied. Signs are concrete things, or events, or processes and they have been categorized and studied systematically for centuries. Memes are thus the subject of a body of theory that Peirce called Semiotic!

Having said this, one might think that I will conclude that memetics is much ado about something already steeped in a vast literature and a rich theoretical basis, and to which memes have little to offer but an alternative terminology. But I don't think so. The meme concept has generated recent excitement precisely because it seems to offer hope of providing something that other theories of social and semiotic processes have not succeeded in providing. It addresses the process of semiosis, i.e., the dynamical logic of how the symbolic and concrete constituents of culture arise, assume the forms they assume, and evolve and change over time. This ambitious promise can only be met, however, if the problems with memetic theory are remedied and if a synthesis within a broader semiotic theory can be realized [see Note 1]. This is a tall order.

So, beyond recognizing this correspondence, what needs to be done? First, we must use this insight to remedy some of the theoretical problems that plague the meme concept, and second we must use it to discover what is uniquely emphasized by the meme concept that can contribute to a theory of semiosis and which may be overlooked or misunderstood in current semiotic analyses.

The core problem of this theory, I think, is a kind of misplaced agency, that gives the impression that both genes and memes — replicators — can be understood without considering their embeddedness in a dynamic system which imbues them with their function and informational content. This, then, is not just a problem with memes, but a problem with the replicator concept in general, inherited from Dawkins' short-circuited description of information processes in biology. I say "short-circuited" because it is not wrong, it just cuts corners that suggest that certain essential aspects of information processing in biological systems can be treated as merely derivative from the replicator concept. In fact, this inverts the reality.

Though an anthropomorphic shorthand has often been used to describe replicator "behavior" (e.g., it is common to read such phrases as "for the benefit of the gene, not the organism" or "genes promote their own self-interest" or "genes act selfishly," etc.), this phraseology was never presumed to project any purposeful or even animate character on genes. Genetic replicators are just strings of DNA sequence information

that happen to get copied and passed on to the future successfully. Genes are not active automatons, they are not "agents" of evolution; but are just structures, passive information carriers that can become incorporated into biochemical reactions, or not, depending on the context. Referring to genes as though they were active agents promoting their own copying process is a shorthand for a more precise systemic account of how a gene's effect on its molecular, cellular, organismic, social, and ecological context influences the probability of the production of further copies. Taken out of this almost incomprehensibly rich context — placed in a test tube, for example — DNA molecules just sit there in gooey clumps, uselessly tangled up like vast submicroscopic wads of string.

So what is wrong with this shorthand, if everyone agrees that genes are just patterns of DNA sequence information? It is the information content that defines a replicator, not its material or even energetic embodiment. The physical gene (i.e., the string of atoms comprising the DNA molecule) is not what evolution conserves and passes on. It is the information, embodied by the nucleotide sequence that is the point. Well, the problem is that the DNA sequence is not intrinsically information, it is just a pattern. Indeed, much of the sequence information in my genes may never be treated as information in the construction and operation of my cells, even if it may at some point become information for evolution. What makes a pattern information? And how does some information become involved in its own replication?

The problem of defining the difference between mere pattern and information is encountered in all realms of information processing and has the same unhelpful answer: what counts as information is context-dependent. The information some pattern can convey is both a function of its distinguishable physical characteristics with respect to other patterns and of its embeddedness in a living/semiotic system that can incorporate this into its functional organization. *No thing is intrinsically meaningful, it is interpreted to be so.* And interpretation is more than replication — yet these two processes are intimately related.

We are all familiar with the fact that in different interpretive contexts the same typographical words or spoken phrases can function quite differently. This can be true of homophones used in different sentences or of cognates in different languages that have historically diverged in their denotations and connotations while their phonological or typographical forms have remained the same. In genetics this is also evident. Gene sequences (homeobox genes, for example) that contribute to the embryogenetic assembly of heads and brains in flies have almost identical counterparts in humans, so much so that inserting the human gene into mutated flies with the corresponding gene damaged can partially restore normal fly head development. But notice, that it does not produce human heads (as in a science fiction parallel). The cross-species sequence similarity (and evolutionary common ancestry) is essential to this function, but the information is interpreted in a fly context.

In both the evolution of words and of genes, these units of pattern have been preserved over time because of the ways they contributed to the usefulness and replication of the larger systems in which they were embedded. Replication has always been a multilevel affair. It is not a function vested in a gene or in a word, or even in a complex of genes or words, but in a dynamic system in which these are only

one kind of unit. It is only in certain unusual and degenerate cases, such as viruses or their computer counterparts, that we can provisionally speak of the functional interpretation of the pattern being self-replication itself, and then only because we recognize that even this function is parasitic and depends on a short-circuiting of a more complex system's functional-interpretational organization. Viruses contain only self-referential replicative information, only if considered without their host.

The selfish gene perspective views genes on the analogy of viral genes, ignoring their implicit parasitism. The caricature of genes as agents competing to inhabit organisms (and the parallel memetic version: "The haven all memes depend on reaching is the human mind" [3: 207]) brackets out any consideration of the systemic origin of gene (and meme) information, as well as its means of replication. This tacitly suggests that the system in which a replicator is embedded can be treated as a passive vessel and the information that constitutes the replicated unit can be understood independent of this context. It assumes that gene/meme replicators are primary and primordial and that the systems in which they are found are derived solely from their interactions.

Quite the opposite is true in both biologic and semiotic processes. These units actually *emerge* out of the systemic relationships, crystallize out, so to speak, as regularities and bottlenecks in the passage of information become entrenched in the course of these systems' evolution [see 4; Note 2].

Life did not begin with genes. Most biologists now believe that DNA probably came into the process of evolution long after life began, as a more stable, more easily copiable, and functionally compartmentalized molecular storage medium, with proteins and RNA preceding. The sub-sequences along a DNA molecule that subsequently came to be recruited as "genes" are in this sense artifacts that reflect prior regularities of higher-order systemic molecular relationships comprised by different kinds of molecules than DNA. True, once life adopted this sequestered form of molecular information storage it changed the nature of the evolutionary process, but one should not make the mistake of thinking that this makes genes more primary than the functional chunking of the molecular systems to which they evolved to encode.

The same can be seen to occur with word change. Eliminative and deconstructive processes that simplify and subdivide what count as meaningful linguistic units, and synthetic processes that lump formerly separate meaningful units into larger unanalyzed wholes, are both at work in language change, driven by the social-discursive-typographical-pragmatic chunking that determines their selective advantage with respect to other alternatives. (I won't speculate on whether there could have been a protolanguage before words, etc., that they replaced) [see Note ?].

Genes and memes are not the locus of the replication process, nor are they somehow the functional unit of information. They are replicas not replicators. They are rather more like the concretion of information bottlenecks in a system. Like a computer operation that mechanically manipulates patterned elements irrespective of their interpretation as symbols and yet still produces outcomes that correspond to symbolic operations, the advantage of replicas that capture the functional chunking of a larger system is that they too can be operated on (e.g., mutated, recombined, and

selected) irrespective of function and yet have functional consequences when interpreted into a system [see Note 3].

Recognizing this computational parallel leads to an important redefinition of what constitutes a meme. Once the meme concept is divested of virtual agency we find that it is not an idea, nor some "unit of information residing in a brain" [see 2], nor "something however intangible, ... passed on ... when we copy each other" [1: 52], it is not even the information encoding some cultural function, it is merely a physical pattern. It is something concrete, mostly outside of brains (though something like the instructions for making it may be in a brain), that can easily and unambiguously be copied, precisely because it is some tangible thing [see Note 4].

Genes and memes are the physical loci where the replicative and adaptational functions intersect, but these loci do not "contain" the information that constitutes these two functions any more than they determine their own replication. The information for both functions is constituted by whatever relationship exists between the physical pattern of the token and the system of processes in which each is embedded. This helps to resolve a number of the problems determining what constitutes a meme. It is not insubstantial. It is some artifact form or type. Its information content is not intrinsic. It is a physical pattern that can be both copied and mapped into the functional relationships of other systems. It conveys information via replication of this pattern directly or by having the pattern re-expressed (i.e., transcribed) in terms of features of some other higher-order system. Its physical boundaries are also not intrinsic, but rather are a function of the interpretive processes into which it gets recruited (as is the case for DNA sequence information, by the way). Bounded interpretation may nevertheless be aided if the artifact in question is itself physically bounded. Because this medium has its own physical and copying characteristics, irrespective of its information content, there is an effect of these physical constraints on the ways that the information it bears can evolve.

4 A New Synthesis?

This brings us back to the equation that started this critical evaluation. A meme is a sign: some physical thing which, by virtue of some distinctive feature, can be recruited by an interpretive process within a larger system as re-presenting something else, conveying information into that system and reorganizing it with respect to that something else.

So what does the meme concept add to semiotic analyses? The answer is that it could rescue semiotic from being a merely descriptive taxonomic enterprise by leading the way to a theory of semiotic causality and of the generation of representational forms and interpretive systems. The key to this derives from two prominent aspects of meme theory that are not so clearly articulated by most semiotic analyses: (1) a focus on the primacy of replication, or literally "re-presentation" for grounding (evolving) referential processes (functions); and (2) the recognition that representational relationships emerge from competitive/cooperative evolutionary processes which determine the form and persistence of the interpretive systems in which they are embedded. Though in this short essay I cannot hope even to outline

such a theory [see Note 5], I will conclude with a few suggestive comments on the direction such an enterprise might take with respect to these overlapping conceptual issues.

First, let's consider the primacy of replication in semiosis. According to Peirce, "The only way of directly communicating an idea is by means of an icon; and every indirect method of communicating an idea must depend for its establishment upon the use of an icon" [5: 2.278]. Another aspect of this same relationship can be grasped by noticing that when you communicate an idea to someone else, the fact of its being communicated and understood consists in its being replicated, or re-presented in the other so that there is now a kind of implicit iconism between what is in your mind and in their mind. This is also the essence of the meme concept.

Peirce intends to convey more than this, however, by suggesting that communicating via icons or by indirect reference to icons is how all communication is inevitably grounded. This is because all other forms of reference are indirect and dependent on reference to icons. This hierarchic relationship is critical, and is not fully spelled out in meme theory nor is it agreed upon by semioticians. I give a superficial account of it in my book *The Symbolic Species* [see 6: 69-101]. To summarize three chapters of that book in a single sentence: indices are constituted by relationships among icons (and are thus indirect re-presentations) and symbols are constituted by relationships among indices (and are thus doubly indirect re-presentations). In other words, to interpret any sign is ultimately to analyze it to its constituent iconic basis. So at this level of analysis, mimesis can be seen as both the primary mode of information transmission and also the ground of interpretation.

In evolutionary terms, these two relationships — replication of the significate pattern and transcription of it into some other medium and with respect to some systemic adaptation — also overlap when it comes to the explanation of evolved forms. Natural selection favors the replication of genetically encoded information that in some way or other internalizes critical features of the organism's environmental context that affects its replication (e.g., up-regulation of microbial enzyme production in response to the increased concentration of a food substrate). Thus an adaptation is a kind of mapping of internal organization to some relevant environmental feature: a representation of the outside inside. The "goodness" of the representation of this external feature is thus pragmatically assessed with respect to its effects on the probability of replication of the DNA sequence patterns that helped generate it. In this way adaptation is defined with respect to gene replication and vice versa, producing linked replications at two levels, but neither makes sense without the other. Without any (even indirect) adaptation genes have no information; without some substrate to embody and replicate the information implicit in the adaptation there is nothing for it to be an adaptation for.

A parallel argument can be made for semiosis in general. The ultimate significance of a sign is grounded in the consequences of the system of interpretive habits and dispositions it generates and is generated by (i.e., in the larger context of patterns of sign usage of which it is a part). As in biology, then, the relative prevalence and probability of a sign's use will be a function of the correspondences between its significant consequences and some relevant aspects of the larger context in which it is embedded. Peirce himself appeared to recognize this in his later

writings, hinting that pragmatism might naturally follow from a universal evolutionary logic.

Until now, classic semiotic theories have not had much to say about why certain signs persist and others do not, or why certain semiotic systems evolved the forms they now exhibit. Indeed they have been mostly synchronic and descriptive in their intent. But if the evolutionary analogy is appropriate here, then we may be able to understand fully the structure of semiotic systems and the types of sign relationships and interpretive processes that constitute them only if we analyze them as products of competing replicative processes: signs competing with signs. However, rather than competing for space in human memory, signs must be understood as physical markers competing for representation in the physical media of the semiotic systems that surround us. Their incorporation into brains is not the measure of signs (memes) evolutionary "success," as though signs were parasites. Rather, like the interpretation of genetic sequence information when it is transcribed into the machinery of the cell, a sign's replication depends upon its effect on what brains do, via the behavioral adaptations the signs promote. To the extent that this behavioral or cognitive phenotype increases the probability that the signs will be produced and used again, especially by others emulating similar functions that also make use of them, they will be spread and stabilized in the semiotic corpus.

Semiotic theories since Peirce have distinguished between the replica and the functional aspects of sign relationships. Specifically, Peirce recognized that one could restrict analysis to consider only the physical qualities that serve as sign vehicles without explicit consideration of their relationships to their referent objects or interpretation (an enterprise Charles Morris called Syntactics). Computation theory and mathematical information theory make use of this analytic compartmentalization. And it is to some extent what meme theory attempts as well. However, in these cases it is tacitly assumed that the items under study are already involved in interpretive processes and bear significance to a larger system. Provisionally "forgetting" this antecedent assumption (as is the case here) has often led theorists in all three domains to claim primacy of this lowest level of analysis. So, for example, the idea that cognition is nothing but "computation" — the rule-governed mechanical manipulation of markers according to their shapes — is a parallel fallacy to the memetic fallacy [see Note 6].

Peirce also argued that one could analyze referential relationships without explicit consideration of their relationships to the pragmatic consequences that are their ultimate basis of interpretation (a realm Morris identified with Semantics). Though this hierarchic analytic strategy is implicit in the biological analogy from which the meme concept is drawn, it was obscured by the apparent autonomy of the replicator concept. But the systemic embeddedness cannot be totally ignored, even if a coherent analysis of only the "syntactic" aspect is possible. Both living processes and other semiotic processes imbue the patterns of their constituents (genes or signs) with informative power by virtue of their roles in a larger functional system. Something is only a sign (a meme, a gene) when interpreted, and interpretation is ultimately a physical process involving the production of additional significant patterns: replicating more signs (or "interpretants" as Peirce called them). In other words, sign interpretation is ultimately mediated by sign production (i.e., replication),

as gene function (i.e., interpretation into an adaptation) is ultimately assessed with respect to gene replication. As the biological parallel suggests, this is basis for an evolution-like process that can lead to ever more complex systems of interpretation (adaptation).

5 Some Concluding Remarks

The power of recognizing that memes are signs is that it offers a bridge to help us recognize that biological evolution and life in general are semiotic processes, and conversely that semiotic processes are intrinsically dynamic physical evolutionary processes. Signs evolve, and they have pragmatic consequences, by virtue of which they are selectively favored to remain in circulation or become eliminated over time. It is by virtue of the memetic analogy to genetic evolution that we may discover the dynamical logic still required for a complete theory of semiosis, and not just a semiotic taxonomy.

Biology is a science concerned with function, information, adaptation, representation, and self-other relationships —semiotic relationships, not merely physical relationships. Biologists tend to focus on the physical relationships among the information-bearing units in living processes whereas semioticians tend to focus on the significate relationships of the units of human communication, but ultimately these apparently unrelated analyses in different domains are entirely complementary. Taking this correspondence between biological information processes and social information processes more seriously may provide some new insights into the dynamical logic by which representational relationships form and evolve. I believe that this new synthesis will be the crucible from which a general theory of semiosis will eventually emerge. The theory of memetics is not the answer to a theory of social and psychological evolution, but reinterpreted it may suggest some bridging concepts that can lead to a unifying methodology for the semiotic sciences.

References

1. Dawkins, R. [1999]. Blackmore Susan (Ed.). *The Meme Machine*. Oxford U. Press.
2. Dawkins, R. [1982]. *The Extended Phenotype*. Freeman.
3. Dennett, D. [1991]. *Consciousness Explained*. Little Brown.
4. Deacon, T. [2003]. "Multilevel selection in a complex adaptive system: the problem of language origins." In B. Weber & D. Depew (eds.), *Evolution and Learning: The Baldwin Effect Reconsidered*. MIT Press.
5. Peirce, C. S. [1931-53]. *Collected Papers of Charles Sanders Peirce*, 8 vols., ed. Arthur W. Burks, Charles Hartshorne, and Paul Weiss (Cambridge: Harvard University Press).
6. Deacon, T. [1997]. *The Symbolic Species: The Coevolution of Language and the Brain*. W. W. Norton, and Co.

7. Deacon, T.W. [2003]. "The Hierarchical Logic of Emergence: Untangling the interdependence of evolution and self-organization." In: B. Weber and D. Depew (eds.), *Evolution and Learning: The Baldwin Effect Reconsidered*. MIT Press.

Notes

1. Language is an evolved biological phenomenon, and its core architectonic logic must be conceived in evolutionary terms, and our conception of the relevant evolutionary processes may need to be considerably broadened to deal with the complexities involved.

The key concept for understanding this logic is what I will call "masking." I am using this term to refer to shielding or protecting something from natural selection. Selection is often viewed as a negative or subtractive factor in evolution. This is **implicit**, for example, in Baldwin's theory of evolution, and in part motivates his theory as a means to introduce a "positive factor" (Baldwin, 1896). But this is an oversimplification. Selection is generated by any bias affecting reproduction of genes or traits, positively or negatively. What links all these multilevel co-evolutionary theories together is that they involve differential exposure to selection: either as something is shielded from natural selection or else is unshielded and newly exposed to selection by virtue of some action of the organism itself.

By unpacking the causal logic of the competing evolution theories and recognizing the central role played by masking functions, and the way these distribute, collect, and reorganize selection factors, we have begun to glimpse the generality of these processes, and how they may actually contribute to such complex multilevel adaptations as language behavior. This can liberate us from the tendency to think in terms of single evolutionary causes and simple accounts of adaptive structures. It also creates a biological context within which to begin to consider the contributory roles of systemic self-organizing processes, without succumbing to the tendency to attribute Lamarckian functions to them, or in any way envisioning that they undermine the constraints and logic of natural selection. Natural selection is simply far more complex than is often appreciated. But it would be a mistake to imagine that we could reduce the logic of these processes to nothing but natural selection. The very topology of the process is fundamentally altered in this case. The simple feed-forward of phenotypic effects on reproductive success has become embedded in another feed-forward loop in which other features of the phenotypic effect come to be amplified and to play a significant role in the direction that evolutionary change takes. This kind of directional-biasing effect (although it does not imply that evolution is "directed" in any vague intentional sense) cannot be addressed by natural selection theory itself.

Approaching evolutionary processes in a way that respects the complexity and self-organizing character of the substrates in which they occur can offer hints of a higher-order logic concerning nested levels of evolutionary processes within evolutionary processes. In the language evolution example, we can begin to see the essential interplay that must exist between self-organizing and selection processes that constitutes evolution. Not only must we recognize that evolutionary processes include self-organizing processes, but also that they may include nested levels of

other evolutionary dynamics. To understand the details of language origins (including both the origins of biological language adaptations and the origins of linguistic structures) we need to understand how one evolutionary dynamic can emerge from another, i.e., how linguistic evolution can emerge from molecular evolution, and how this can feed back to alter the very nature of the antecedent process.

Language evolution was shown to be analogous to niche construction in some regards. Niche construction introduces a complicated compound interest effect into evolution in which products of behavioral intervention in the environment can feed forward to become selection pressures in future generations. What we might call the "linguistic environment" is a "niche" that is both a complex behavioral habit shared by a population of humans and yet also a set of highly specific cognitive constraints that places special demands on neural functioning. Though created by human activity the constraints that define it are not derived from human brains, but have emergent features derived from the demands of inter-individual communication and the demands of symbolic reference itself. These niche-like effects have likely shaped human brain evolution as much or more than any ecological demands in the past 2 million years.

In many ways, then, the niche construction model is far too narrowly conceived to capture the sort of phenotypic feed-forward effects that language entails. This analogy needs to be extended to include the special dynamism of linguistic evolution if it is to model brain-language coevolution accurately. Language and symbolic culture are not merely extrinsic inanimate products of human evolution. They are emergent, parasitic (or rather symbiotic), evolutionary processes whose independent dynamics (driven in part by semiotic factors never before relevant to biological evolution) have essentially become the tail that wagged the dog. These semiotic processes are active evolving systems that have their own self-organizing principles and evolutionary dynamics. These have been imposed on human evolutionary biology from the top down producing a totally unprecedented evolutionary dynamic. This evolutionary dynamic linking genes, brains, and symbolic culture is, in effect, a third-order evolutionary process. It is a level above Baldwinian or niche construction processes, both of which may be recruited in its service. In it, complex feed-forward effects predominate, amplifying their influence through an independent selection loop, making human bio-cultural evolution far more convoluted and unpredictable than anything biology has previously produced. There is, literally, no telling where it will lead! [For more detail, see refs. 4 and 7]

2. The explosion of interest in emergentism ("something more from nothing but") among both natural scientists and philosophers has not surprisingly generated considerable ambiguity in what is meant by the term emergence. I offer an inventory of emergent phenomena in the natural world, proposing that emergence takes three forms.

First-Order Emergence: Properties emerge as a consequence of shape interactions.
Example: The interaction of water molecules (nothing but) generates a new property, surface tension (something more).

Second-Order Emergence: Properties emerge as a consequence of shape interactions played out over time, where what happens next is highly influenced by what has happened before. Example: The formation of a snowflake, where initial and boundary conditions become amplified in effect over time. In general, complex or "self-organizing" systems display second-order emergence.

Third-Order Emergence: Properties emerge as a consequence of shape, time, and "remembering how to do it." Example: Biology, where genetic and epigenetic instructions place constraints on second-order systems and thereby specify particular outcomes called biological traits. These traits then become substrates for natural selection by virtue of the fact that 1) their instructions are encoded and 2) they endow organisms with adaptive properties.

3. Though the special demands of communicating with symbols are in one sense abstract and formal, human neural biology has been made sensitive to them for the simple reason **that** when these requirements are respected and embodied in culture they can produce incredible changes in the physical conditions of life and reproduction. This view of language evolution can be seen as a natural extension of developmental systems theory, but it is much more than this. It is no more than a glimpse into a much larger domain that has yet to be exposed: evolutionary semiotics. Until a more principled analysis of the logic of these sorts of processes is understood, we will remain in a prescientific phase of cognitive and social science, and lack a considerable part of the story of human brain evolution.

4. A vaguely similar reversal of Dawkins' view that memes are ideas and their external exemplars are something like the phenotype is presented by W. Benzon in his article "Culture as an evolutionary arena" (Journal of Social and Evolutionary Systems, Vol. 19, p. 321-362), though his justification of this is somewhat different.

5. Language, for example, didn't just evolve in the standard sense of that term. Language is an "emergent" biological phenomenon in ways unlike almost any other biological product [see ref. 7]. Many biologists have been forced to conclude, as I have, that language is as unprecedented in biology as was the very origins of the molecular genetic code. Imagining that explaining it would not require significant augmentation of existing evolutionary theory is, in my opinion, naive. Attempting to understand its emergent character by reducing it to a mere change in anatomy, an increase in intelligence, or a special case of computation inevitably obfuscates this central emergent nature. Because of this special status among evolved phenomena, the mystery of language forces us to explain it on different terms than say the evolution of upright posture or color vision. It is a problem analogous to explaining the evolution of an evolutionary process.

There are three additional considerations that I believe are indispensable to devising an adequate evolutionary approach to this mystery: (1) an appreciation of the role of self-organizing dynamics in the production of the complex systematicity recognized in language structures; (2) an evolutionary analysis that takes into account the multilevel interactions between biological evolutionary processes and non-

biological linguistic evolutionary processes; and (3) an analysis of the semiotic infrastructure of language that is sufficient to articulate its ultimate functional constraints. Ultimately, the emergence of language offers a puzzle that requires us to rethink the concept of evolution itself in many significant ways, and demonstrates that it is intimately linked with processes of self-organization and semiosis [see 7]. One of the most important elements of this is the emergence of a new level of evolutionary dynamic out of biological evolution. With the emergence of language a second-order level of evolutionary process begins at the level of symbols themselves: linguistic evolution. This changes everything.

6. This is, of course, a bald statement with enormous implications that I cannot unfold here, but do deal with in depth in two upcoming books: *Homunculus: How Brains make Experience* and *Golem: Making Things Think*.