

Chapter 2

Beyond the Symbolic Species

Terrence W. Deacon

Abstract Confusions about the nature of symbolic reference are at the core of two major challenges to understanding human language. A failure to take into account the complex iconic and indexical infrastructure of symbolic interpretation processes has blocked progress in the study of language structure, language evolution, neural processing of language, and language acquisition. Simplistic notions of symbolic interpretation are critiqued, the semiotic infrastructure of symbolic interpretive processes is described, and some implication for understanding the universals of grammar and syntax are explored. Finally, the evolutionary problem of language origins is re-examined and an unexpected important role for relaxation of selection is demonstrated.

2.1 The Problem with ‘Symbol’

In the years since the publication of *The Symbolic Species* (Deacon, 1997) one consistent source of confusion has persistently been used as a reason to take a critical view of the symbolic threshold as key to the human difference. This is in one sense merely a terminological problem with interpretations of the term ‘symbol,’ and yet it obscures a critical issue that if not resolved will be a roadblock to both the study of language and the further development of semiotic theory itself. The confusion superficially has to do with the concept of arbitrariness of reference, but more deeply it involves a tension between a structural and dynamic conception of the process of semiosis more generally.

I will first address the terminological dispute, which although a source of confusion in the literature, should be resolvable with a bit of care in defining terms and avoiding the attribution of one definition to uses where it does not apply. The conceptual dispute is much more subtle, and I think critical to sort out. Failure to do so will have two serious consequences. First, it will doom semiotic theories to the status of mere taxonomic exercises where different scholars are free to invent their

Parts of this chapter are reprinted from Deacon (2010a)

T.W. Deacon (✉)

Department of Anthropology, University of California, Berkeley, CA 94720-3710, USA
e-mail: deacon@berkeley.edu

own categorical principles without careful reflection on the underlying generative processes and constraints that determine the semiotic differences they hope to distinguish. This often ends up turning semiotic research into a renaming exercise, where commonly studied phenomena are redescribed in semiotic terms and it often devolves into battles over competing naming paradigms from the past. Second, and more serious, it will cut semiotic research off from the sciences of psychology, neurology, and biology due to a failure to come to grips with the *process* of semiosis; the dynamic of interpretive activity by which semiotic relationships emerge from other semiotic relationships and ultimately derive their grounding on the physical phenomena they thereby bring into consideration. The problem here is the tendency to imagine signs as things, or as synchronic relationships, whereas they are instead intrinsically dynamic phases in a generative process, and ultimately something apart from the artifacts being manipulated in this process.

The term ‘symbol’ has come to be used differently in different traditions, and so first we need to be clear what we are talking about. If all that is meant is a mark that need not share any specific quality with its object of reference, then the term has trivial consequences. This gloss of the concept makes it easy to dismiss its importance for evolution, and indeed this simplification has been the motivation for many language origins researchers to imagine that it is only syntax that demands explanation. This assumption about the concept of symbol is also reflected in many critics’ claims that most species are capable of learning arbitrary associations (e.g. see [Chapter 3](#), this volume) so claiming that the symbolic capacity divides humans from other species must be trivially false.

This focus on arbitrary correlation as the defining attribute of symbolic reference is a serious oversimplification that collapses critical distinctions between sign vehicle and referential properties. The common usage of a ‘code’ analogy in describing language reference also reflects this simplification, and for similar reasons leads to serious theoretical misunderstandings. A code does indeed involve an arbitrary mapping or correspondence relationship, but that is precisely why its reference is opaque and is the basis for encryption. A code is a mapping of a parallel set of sign tokens to a language, and typically a token-to-token mapping. So to describe language or any of its attributes, such as the basis for phonology, syntax, or semantics as a code, merely begs the question: what is the basis for this mapping relationship?

It is often argued, for example, that arbitrariness is a property of many animal calls. Consider the case of predator-specific alarm calls (which have been identified in species as diverse as vervet monkeys and chickens). The assumption that these calls ‘mean’ or ‘name’ a particular predator is as, the linguist Derek Bickerton (2010) has also argued, a ‘back-projection of our own language-saturated view of the world.’ Alarm calls are indexical, even though they don’t sound like the predator they indicate and even though they are emitted to many similar types of predators. Their arbitrariness and generic reference are red herrings in this detective story. Their reference depends on and evolved from repeated correlations between the presence of a predator, the production of a call, and an appropriate escape behavior, and merely distinguished from other experiences, vocalizations, or behaviors.

A symbolic sign relationship is, in contrast to an iconic or indexical sign relationship, a doubly conventional form or reference. It involves a conventional sign type that is additionally conventionally-mediated in the way it represents.

Arbitrariness is a negative way of defining symbols. It basically tells us that neither likeness nor correlation are necessary. But this is inadequate, even though it is a common shorthand way of characterizing symbolic reference. All sign relationships include some degree of arbitrariness, because those attributes that are taken as the ground for the sign-object linkage can be chosen from many dimensions. Thus, anything can be treated as iconic or indexical of almost anything else depending on the interpretive process.

For example, with a bit of imagination a face can be discerned on the full moon, or in a cloud formation, and it might even remind you of someone you know. But iconism can also be highly abstract, as in the complex way that a mathematical equation refers iconically, once you know how to discern its symbol-mediated isometry (e.g. between the structure of the equation and a corresponding geometric or dynamical relationship). An equation can be interpreted to be iconic (e.g. of a parabolic trajectory) only, however, if you know how to discern the way that differences in the values or operations directly correspond to differences in the geometric object of reference. So one first needs to be able to interpret the symbolic components before the diagrammatic iconism of the equation can be appreciated.

Indices refer by contiguity in space, time, or substrate. A simple correlation can therefore be the ground for indexical reference. A lipstick smear on a man's shirt collar can be a troublesome indication to his wife, a urine scent on a branch can be a sexual index to a female lemur, and the mobbing call of a small bird can indicate the present of a raptor. What gets correlated and how (accidental, cultural, evolutionary) can be arbitrary, only the fact of correlation is not. Thus, a rat in a Skinner box pressing a bar in response to a bell in order to get a water reward has learned that the bell is an arbitrary index of the state of the apparatus (an indexical legisign). These states are arbitrarily paired in the experimental design, but that doesn't make the one a symbol of the other.

So symbolic reference is not merely a function of arbitrariness, conventionality, and generality, though these features are properties that symbolic reference makes available. First of all, arbitrariness isn't required. For example, many symbols used to designate religious concepts employ obvious iconism and yet this doesn't undermine their potential to symbolize quite complex esoteric abstractions. This also demonstrates that the sign vehicles used for symbolic reference need not be widely understood as conventional. When first encountering an unfamiliar religious symbol it may only require a brief few comments to understand its symbolic import. And of course icons, such as the eye-spots on male peacock tail feathers or faces 'seen' in the clouds often bring to mind general types of objects, not just specific instances. These attributes are not sufficient determinants of symbolic function, either individually or collectively.

As Charles Peirce (1931) pointed out over a century ago, we must distinguish properties of the sign vehicle (which he terms a representamen), which can include being an arbitrarily defined (i.e. conventional) type of sign vehicle, from properties

taken to link it to its object of reference. Thus although current vernacular has habitually termed alphanumeric characters “symbols” this usage ignores any referential relationship. If not used carefully, in recognition of this shorthand, it can lead to all manner of theoretical confusions.

Thus when your computer begins randomly spewing alphanumeric characters onto your screen they are indices of a malfunction, not symbols of anything. And likewise the typographical character combination ;-) does not refer symbolically, even though it is composed of conventional tokens designed for symbolic purposes. Peirce terms conventional sign vehicle types ‘legisigns,’ and argues that symbols must also employ legisigns. However he notes that legisigns can also serve iconic and indexical roles as well. Consider, for example, the conventionalized stick figure icons on restroom doors, or the use of red for traffic lights and road signs to indicate the requirement to stop (i.e. it indicates a convention—an injunction to act according to a rule—but it does not ‘mean’ “stop” in the way that this word does. Because legisigns are often created (or chosen) with a specific type of referential relationship in mind it is the arbitrary choice of the creator which properties are to be used referentially. This is why legisigns created for typographical use to symbolize the parsing and punctuation of written text can also be recruited for their iconic features (as in the case of the smiley face).

Of course communicative intention is also an interpretation, and this also does not fix the referential function of a sign vehicle. Whether something is interpreted iconically, indexically, or symbolically depends on what’s going on in the mind of the beholder.

Recognizing that the same sign vehicle need not always be interpreted as intended, or as referring always in the same way is the first step toward reframing semiosis in diachronic, not synchronic, terms. A sign vehicle can be interpreted in multiple ways not because it is in some way a combination of sign types, a fractional mixture of iconic, indexical, and symbolic features, but because its semiotic significance is not vested in the sign vehicle at all. Although a given interpretation may depend on some feature intrinsic to that artifact for motivating its semiotic function, no semiotic attributes are invested in the sign vehicle itself. They are properties of it being interpreted (whether in its creation or its consideration). So given that the same sign vehicle can be interpreted differently by different individuals, or at different phases of considering it, worrying about whether *it* is a ‘pure’ sign of a given type or a ‘mixed’ sign commits the fallacy of misplaced concreteness.

As we will discuss below, a given sign relation is created by an interpretive process. It is a phase in this process in which the sign vehicle is incorporated in a particular way, but which may be transitory, leading to a different mode of considering that same sign vehicle. And at any given phase of this interpretive process there is no ‘mixture’ of semiotic characteristics. It is only when we attempt to analytically collapse this process into a single synchronic relation that we run the risk of confusing sign vehicle properties with semiotic properties and think of signs as simultaneously exhibiting iconic, indexical, and symbolic features.

Although it is far beyond the scope of this chapter to attempt a reframing of semiotic theory in process terms, carefully dissecting a few examples of interpretive processes can help to illustrate the difference between this and more synchronic forms of semiotic analysis and clear up confusions created by the ‘compositional’ account of symbolic reference presented in *The Symbolic Species* (Deacon, 1997). More importantly, exemplifying the process of hierarchic differentiation of referential form that constitutes an interpretive process allows us to see how semiotic analysis is directly relevant to understanding cognition, and by implication the evolution of symbolic cognition.

As a starting point for exhibiting the hierarchic dependency of the different modes of referential interpretation consider one of the classic examples of a symbolic form: the impression of a signet ring in wax used to seal a note and verify the sender’s identity. Tracing the minute cognitive steps necessary to interpret this simple sign demonstrates that symbolic function depends on more than a simple arbitrary correspondence. First, the formal similarity between the impression and the ring is primary. This is iconic. But without the physical action of the ring-bearer pressing the ring into hot wax to produce this likeness, it would not indicate that this message, thus sealed, was produced by the bearer of that specific ring. The presumed connection between ring and bearer further indicates that a particular individual actually sealed the note. Finally, possession of such a ring is typically a mark of authority, royalty, etc. This status is a social convention. To interpret the wax impression as a symbol of social position, one must also understand these social conventions, because nothing intrinsic to the form or its physical creation supplies this information. The symbolic reference is dependent on already knowing something beyond any features embodied in this sign vehicle.

This dependency on an external system of relations within which the formal similarities and correlative aspects of the wax impression are embedded is a critical property of its symbolic reference. But without familiarity with this entire system of relationships, these non-symbolic components remain merely icons and indices. Indeed, if any link in this chain of referential inferences is broken, symbolic reference fails. So while the features comprising the sign vehicle are not necessarily similar in form or physically linked to what is symbolized, this superficial independence is supported by a less obvious network of other modes of reference, involving both iconism and indexicality.

Notice that the first step in this interpretive analysis involved recognition of an iconism. Only after this recognition was the implicit indexicality relevant and only after that was the social convention able to play a role in providing symbolic significance to the sign vehicle. This hierarchic dependency of symbols on indices on icons was the core semiotic argument of *The Symbolic Species*. But notice that it is not a simple compositional relationship. Indices are not *made of* icons and symbols are not *made of* indices. These are stages in developing and differentiating ever more complex forms of reference. Throughout the interpretive process described above there was only one sign vehicle: the wax impression. At first it is interpreted iconically, then indexically, and finally symbolically. The constructive nature of this interpretive process was what was critical. These semiotic relationships were not

mixed in some fractional sense, they were distinct dependent phases in the process, and most of the relevant detail was supplied by the interpreting process not the wax impression.

This account leaves out many subphases of the interpretive process, but it captures the crucial architectonic that I believe is critical to understanding why there might be a cognitive threshold separating iconic and indexical forms of communicating, common to most mammals and birds, from symbolic communicating that is distinctive of humans. Interpreting something symbolically is simply more complex, and unlike iconic and indexical interpretation there is nothing inherent in the form or physical relationships of the sign vehicle to provide an interpretive clue. This must be supplied entirely by the interpretive process itself, and it is of the nature of a systemic relationship, not some singular object or event.

Before turning to language, it is worth exploring a few other simpler examples of this interpretive differentiation process in order to appreciate the generality of this hierarchic semiotic dependency.

Let me begin with a trivial index: a wind sock that indicates the strength and direction of the wind. What constitutes the interpretive competence to recognize this indexicality? Imagine that it is being seen for the first time through a window. It is iconic of cloth or clothing, and yet it is clearly not clothing or randomly fluttering cloth. Its distinctive shape and careful design, in contrast, indicate that it is likely designed for a purpose. Another iconic feature is its extended fluttering behavior, again iconic of clothing, but of clothing being blown by the wind fluttering on a clothes line. This iconism now brings to mind something that is not directly provided by the sign vehicle: wind. By virtue of developing these iconic interpretations then this sign vehicle is now embedded in a larger context in which something present points to something that it is not: the wind. And a further juxtaposition of iconisms that have involved other windblown experiences can eventually (quickly) lead to interpreting its behavior as an index of both the direction and intensity of the wind. The indexicality is not 'composed of icons' but rather emerges from the comparisons among iconic interpretations. Failing to recognize these iconisms, e.g. because of never having experienced the effects of wind, would make the indexical interpretation impossible to develop.

Next consider the interpretation of the chevron insignia on a military jacket. Initially, it appears just a colored shape, an iconic sinsign in Peircean terminology (a singular instance of something familiar). As similar shapes are seen on other shoulders, it develops from an iconic sinsign to an iconic legisign (shapes of the same type). As it is understood to distinguish the individual wearing it, it becomes interpreted as an indexical legisign (a type of sign vehicle pointing to something about this person). When its particular configuration is understood to designate that person's military rank it becomes interpreted as a symbolic legisign. The same sign vehicle thus is the locus for a sequence of interpretive phases in which both the relationship of the sign vehicle to other sign vehicles and the relationship of the sign vehicle to its reference are progressively developed.

Some of my favorite examples of this hierarchic interpretive dependence are captured in political cartoons and illustrations that make a general statement about

Fig. 2.1 Cartoon from the cover of *New Yorker* Magazine which exemplifies the progressive differentiation of iconic to indexical to symbolic interpretive phases (see text)



things by virtue of the atypical juxtapositions they employ. Consider the cartoon cover from the *New Yorker* Magazine in Fig. 2.1.

On first glance, as soon as the discordant features of the image are appreciated, one's mind jumps to an interpretation that is beyond anything depicted. It is commenting on a somewhat paradoxical aspect of motherhood. But how does it induce us to make this quite abstract interpretation? Seen in isolation an image of a mother and baby or an image of a child playing with a puppet do not 'say' anything, or even provide new information. But the violation of expectation created by the baby controlling the mother puppet is not merely interpreted iconically. Its inversion of expectation is interpreted indexically, pointing to its opposite: mothers control babies. This, in turn, reciprocally points back to the partial truth of the abstract relationship of baby controlling mother, and thereby to the paradox that both abstract relations are true, though the image is absurd. In this example, relationships between icons, one present another brought to mind by it, initiate an indexical interpretation of this relationship that ultimately leads the viewer to interpret this as being about something much more abstract and general. Although this interpretive process involves iconic, indexical, and possibly symbolic interpretive phases (the latter

to the extent that it comments on the conventional cultural assumptions about motherhood), these are not vested in the sign vehicle and are not mixed or additive. They are distinct phases of interpretation in which the same complex sign vehicle is given progressively more differentiated and context embedded interpretations. Failure to initially interpret the iconisms would make it impossible to interpret any indexicality and failure to interpret the indexical relationships would make it impossible to ever assign any symbolic meanings to the image.

The import of these simple examples is this: to generate an indexical interpretation of any sign vehicle requires interpreting it iconically and interpreting this iconicity with respect to other iconic interpretations, and interpreting it symbolically requires interpreting it indexically and interpreting this indexicality in context with other indexical interpretations. A higher order interpretive process must in this way be supported by a lower order interpretive process, and so on down to the most minimal form. Although this analysis only focuses on this representational triad, it in fact captures an enigmatic aspect of Peirce’s 9-part sign categorization system (shown in Fig. 2.2).

In this taxonomic scheme there are three levels of sign vehicle relationship three levels of sign-object relationship and three levels of relationship between a sign and its immediate interpretive semiotic effect (its *interpretant*). One of the strictures that Peirce imposes on the use of this taxonomic triad of triads is that the level of the sign vehicle must be at least as high as the level of the sign-to-object relationship and this must be at least as high as the relationship of the sign to its interpretant. But recognizing that an interpretant is itself, according to Peirce, another sign generation process (what I have above described as a phase of interpretation) we can now see that indeed a sign of a higher type depends on being interpreted (and thus its referential capacity generated) by the generation of lower order signs.

Language competence rests on a quite elaborate system of iconic and indexical relationships that necessarily come into play in the production and interpretation

	sign vehicle itself	sign to object	sign to interpretant
3	legisign	symbol	argument
2	sinsign	index	dicisign
1	qualisign	icon	rheme
	1	2	3

Fig. 2.2 Peirce’s 9-part sign taxonomy. Each sign type is defined by the combination of one property from each column such that no property from a column to the right is at a higher level than that to its left. Thus there can be a rhematic indexical sinsign or a dicent symbolic legisign but not a rhematic symbolic sinsign or a dicent iconic legisign

of linguistic communication. What is remarkable about the semiotic infrastructure supporting the symbolic capacity of language is its incredible size and complexity. Its indexical character is made evident by the web of pointing relationships exemplified by a thesaurus, with its one-to-many reciprocal mapping relationships, or a dictionary in which each word or morpheme is mapped to a particular combinatorial relationship among other words. Indeed, a dictionary suggests that a language is a bit like an organism in which every molecule is created by combinations of other molecules interacting. It is this dependence on an underlying semiotic system of relationships that makes this threshold hard to cross for other species. But not only does this serve as the foundation for language reference, these underlying semiotic supports and requirements are unmasked, so to speak, when symbolic relationships are juxtaposed to form even higher order iconic, indexical, and symbolic complexes. Thus, like a circuit diagram that can only be seen as iconic of a type of electronic circuit when its component features are given correct symbolic interpretations, a sentence or narrative depends on first interpreting its symbolic components and then interpreting the higher order iconic and indexical relationships that their combinatorial relationships offer. These hierarchically embedded and emergent semiotic constraints turn out to be key to understanding the higher order logic of grammar and syntax.

2.2 The Semiosis of Grammar and Syntax

True symbolic communication and grammar are inextricably intertwined. They are hierarchically dependent. It is fundamentally impossible to have grammar without symbolic reference, though grammatical relationships don't automatically come to the fore with all forms of symbolic interpretation. Grammar and syntax are, however, intrinsic symbolic attributes that emerge into relevance as symbols are brought into various semiotic relationships with one another; e.g. in combinatorial referential processes. Once we overcome the tendency to treat symbolic reference as mere synchronic arbitrary correlation we can begin to discern the many contributions of the iconic and indexical supports of symbolic reference that have become incorporated into the constraints that define the grammar of language.

Because symbolic reference involves a complex higher-order interpretive development in order to emerge from more basic iconic and indexical relationships, there are implicit constraints that these supportive semiotic relationships impose on operations involving symbol combinations, such as phrases, sentences, arguments, and narratives. These constraints emerge from below, so to speak, from the semiotic infrastructure that constitutes symbolic representation rather than needing to be imposed from an extrinsic source of grammatical principles. Although this infrastructure is largely invisible, hidden in the details of an internalized system acquired in early experience, using symbol combinations in communicative contexts unmasks the iconic and indexical constraints that are implicit in this infrastructure. These semiotic constraints have the most ubiquitous effect on the regularization of

language structure, but in addition there are sources of weaker less ubiquitous constraints also contributing to cross-linguistic regularities. These include processing constraints due to neurological limitations, requirements of communication, and cognitive biases specific to our primate/ hominid evolutionary heritage. Although none of these sources of constraint play a direct role in generating specific linguistic structures, their persistent influence over the course of thousands of years of language transmission tends to weed out language forms that are less effective at disambiguating reference, harder to acquire at an early age, demand significant cognitive effort and processing time, and are inconsistent with the distinctive ways that primate brains tend to interpret the world.

The list of sources of constraint on language structure can be broken down into four main categories as listed below. They each contribute a number of quasi-universal traits and highly probable language regularities, many of which are listed for a given category of constraint type. These categories and language consequences are listed below:

A. *Semiotic constraints*

1. Recursive structure (only symbols can provide non-destructive [opaque] recursion across logical types)
2. Predication structure (symbols must be bound to indices in order to refer)
3. Transitivity and embedding constraints (indexicality depends on immediate correlation and contiguity, and is transitive)
4. Quantification (symbolized indices need re-specification).
5. Constraints can be discovered pragmatically and ‘guessed’ prior to language feedback (because of analogies to non-linguistic iconic and indexical experiences).

B. *Processing constraints*

6. Chunking-branching architecture (mnemonic constraint)
7. Algorithmic regularization (procedural automatization)
8. Neural substrates will vary on the basis of processing logic, not linguistic categories

C. *Sensorimotor schemas & phylogenetic bias*

9. Standard schema/frame units (via cognitive borrowing)
10. Vocal takeover (an optimal medium for mimicry)

D. *Communication constraints*

11. Pragmatic constraints (communication roles and discourse functions)
12. Culture-specific expectations/prohibitions (e.g. distinctive conventions of indication, ways of marking discourse perspective, prohibitions against certain kinds of expressions, etc.)

2.2.1 *Semiotic Constraints*

The most important and ubiquitous source of constraints on language organization arise neither from nature nor from nurture. That is, they are not the result of biological evolution producing innate predispositions and they are not derived from the demands of discourse or the accidents of cultural history. Semiotic constraints are those that most directly reflect the grammatical categories, syntactic limitations, and phrasal organization of language. They are in a real sense *a priori* constraints, that precede all others. Consequently they are most often confused with innate influences.

In a recent and now well-known theoretical review of the language origins problem (Hauser, Chomsky, & Fitch, 2002) Noam Chomsky appeared to retreat from a number of earlier claims about the innate ‘faculty’ for language, but he repeated his long-term insistence that what makes the human mind unique is an innate capacity to handle recursive relationships. Like many related claims for an innate grammatical faculty, this one too derives from a reductionistic conception of symbolic reference. If we assume, in contrast, that non-human communication is exclusively mediated by iconic and indexical forms of reference and that only human communication is symbolic it becomes clear why recursively structured communication is only present in humans.

Symbolization enables substitutions that cross-logical-type (e.g. part for whole, member for class, word for phrase) levels in linguistic communications. Neither icons nor indices can refer across logical types because of the involvement of sign vehicle properties (e.g. similarity of form, correlation in space or time) in determining reference. But because of the independence of sign vehicle properties from the objects of reference, symbols can represent other symbolic relationships including even combinations of symbols forming higher logical type units (such as phrases, whole sentences, and even narratives). This is exemplified by pronominal reference and also includes recursively operating on iconic and indexical relationships.

In summary, recursion is not an operation that must be added to human cognition over and above symbolic capabilities, it is a combinatorial possibility that comes for free, so to speak, as soon as symbolic reference is available. But it is not possible when restricted to only iconic and indexical forms. So the absence of recursion in animal communication is no more of a mystery than its presence in human communication. The reason that it is not found in the communication of other species is simply due to their lack of symbolic abilities.

Though recursion is made available with symbolic communication, it need not be taken advantage of, and so its paucity in child language and pidgins and its absence in some languages (e.g. Everett, 2005) is not evidence that it is an unimportant feature of language. But it is an important means for optimizing communication. Recursion provides means for condensing symbol strings. By repeated recursive operations it becomes possible to refer to an extensive corpus of prior discourse. This not only optimizes communicative effort, it also reduces working memory load because a large corpus of material can be subsumed into the reference of single symbolic unit

(such as a pronoun). However, recursion also creates new ‘housekeeping’ requirements that demand specialized forms of symbolized iconic and indexical operations (see below).

All languages require at least a dyadic sentential structure, i.e. something like a subject-predicate sentential form or a topic-comment structure. Although holophrastic utterances, commands, and expletives, are not uncommon, they typically are embedded in a pragmatic context in which what they refer to is made salient by non-linguistic means. Previous suggestions that this fundamental structure reflects an action-object, agent-patient, or what-where dichotomy have been easily refuted by demonstrating the ease with which these cognitive categories can be interchanged in their grammatical roles. In any case, this most general feature of language structure requires an additional explanation if language reference is treated as simple arbitrary correspondence.

Since Frege, it has been explicitly recognized that isolated terms express a sense but lack specific reference unless embedded in a combinatorial construction roughly corresponding to a proposition. The assignment of a specific reference to an expression or formula and thus to make an assertion about something is called predication. In logic a well-formed (i.e. referring) expression requires both a symbolic function and an argument (i.e. that to which the function is applied). In addition a complete ‘predication’ requires ‘quantifying’ the argument (unless it is a proper name). This latter requirement and exception is telling. In English, quantifiers include such terms as “a,” “the,” “some,” “this,” “these,” and “all.” Since a proper name refers to an individual thing or person, reference in this case is unambiguous as it is in such mass terms as “water” or abstract properties such as “truth” when speaking generally.

Why is this basic structure necessary and what are the linguistic consequences? Again, I believe that the answer is to be found in the complex structure of symbolic interpretation.

Consider propositional form and argument structure in logic. First order predicate logic is often considered the semantic skeleton for propositional structure in language, though its primary form is seldom explicitly exhibited in natural language. It is characterized by a “predicate(argument)” structure of the form $F(x)$, where F is a function and x is a variable or “argument” operated on by that function. Such an expression is the basic atomic unit of predicate logic. Such an expression may refer to an event, state, or relationship, and there can be one-, two-, three- and no-place predicates determined by how many arguments they take. So for example the function “is green” typically is a one-place predicate, “is next to” is a two place predicate, and “gives” is a three-place predicate.

This suggests the following hypothesis: *Predicate (argument)* structure expresses the dependency of symbolic reference on indexical reference as in *Symbol (index)*.

Once source of evidence for this semiotic dependency is implicit in the way that deictic procedures (e.g. pointing and other indicative gestures) are used to help fix the reference of an ambiguous term or description, and can even be substituted for the subjects and arguments of a sentence. Thus for example, uttering the word “smooth” in a random context only brings attention to an abstract property, but when

uttered while running one's hand along a table top or pointing to the waveless surface of a lake, reference is thereby established. It can also refer even if uttered in isolation of any overt index in a social context where the speaker and listener have their joint attention focused on the same flawless action. In this case, as with holophrastic utterances in general, the symbolic reference is established by implicit indication presupposed in the pragmatics of the communicative interaction. Indeed, where explicit indexing is not provided, it is assumed that the most salient agreeing aspect to the immediate context is to be indicated. In general, then, any symbolic expression must be immediately linked to an indexical operation in order to refer. Without such a link there is sense but no reference.

This is a universal semiotic constraint (though not a universal rule) that is made explicit in logic and is implicit in the necessary diadic structure of sentences and propositions. It is a constraint that must be obeyed in order to achieve the establishment of joint reference, which is critical to communication. Where this immediate link is missing reference is ambiguous and where this constraint is violated (e.g. by combinations that scramble this contiguity between symbolic and indexical operations; so-called word-salad) reference typically fails.

This constraint derives from the unmasking of indexical constraints implicit in the interpretation of symbolic reference. Because symbolic reference is indirect and "virtual," by itself it can determine only ungrounded referential possibility. The subject, topic, or argument (= variable) performs a locative function by symbolizing an indexical relationship; a pointing to something else linked to it in some actual physical capacity (e.g. contiguous pragmatic or textual context). This reference determination cannot be left only in symbolic form because isolated symbols (e.g. words and morphemes) only refer reciprocally to their "position" in the system or network of other symbols.

The importance of immediate contiguity in this relationship reflects the principal defining constraint determining indexical reference. Indexical reference must be mediated by physical correlation, contiguity, containment, causality etc., with its object in some way. Indexicality fails without this immediacy. There are, of course, many ways that this immediacy can be achieved, but without it nothing is indicated. These constraints on indexicality are inherited by the grammatical categories and syntactic organization of sentences, propositions, and logical formulae.

To state this hypothesis in semiotic terms: A symbol must be contiguous with the index that grounds its reference (either to the world or to the immediate agreeing textual context, which is otherwise grounded), or else its reference fails. Contiguity thus has a doubly indexical role to play. Its contiguity (textually or pragmatically) with the symbolizing sign vehicle points to this symbol and their contiguity in turn point to something else. This is an expression of one further feature of indexicality: transitivity of reference.

Simply stated, a pointer pointing to another pointer pointing to some object effectively enables the first pointer to also point to that object. This property is commonly exploited outside of language. Thus the uneven wear on automobile tires indicates that the tires have not been oriented at a precise right angle to the pavement, which may indicate that they are misaligned, which may in turn indicate that the owner

is not particularly attentive to the condition of the vehicle. Similarly the indexical grounding of content words in a sentence can also be indirect, but only so long as no new symbolically functioning word is introduced to break this linear contiguity.

Of course, every word or morpheme in a sentence functions symbolically and a word or phrase may take on a higher order symbolic or indexical role in its combinatorial relationships to other language units at the same level. This flexibility provides a diversity of symbolized indexical relations. So, for example, arguments can be replaced by pronouns, and pronouns can point to other predicates and arguments, or (via quantification) they can point outside the discourse, or if a language employs gender marking of nouns a gender-specified pronoun can refer to the next most contiguous noun with agreeing gender expressed in the prior interaction, even if separated by many non-agreeing nouns and noun phrases. A sentence that lacks inferrable indexical grounding of even one component symbolic element will be judged ungrammatical for this reason. However, the basis for this judgment by non-linguists is not determined with respect to either explicit rules or constraints. It is determined by the fact that the sentence doesn't have an unambiguous reference.

As mentioned above, both natural language and symbolic logic are constrained by the need to quantify nouns and arguments, respectively. This also exemplifies the need to ground symbolic reference via indices. Quantifiers are specifiers of virtual indexing. Words like "a" "the" "some" "many" "most" "all" etc., symbolize the virtual result of various forms of iterated indications or virtual ostensions (pointings). All quantifiers can be thought of as means for specifying the numerosity of potentially redundant forms of indexicality. They are effectively virtual pointings that take advantage of transitive correlation with other indexical relationships, such as proximity information ("this" "that") or possession information ("his" "your") to differentiate indexicality. One can even imagine a collection that is identified by a symbolized property, being pointed to en masse by a contiguous index, and then carrying out the quantificational operation by literally pointing to some, or few, or all members of this collection.

Analogous to the case of implicit presupposed indexicality in holophrastic utterances, there are also contextual conditions where explicit quantification in language may be unnecessary. This is most obvious in cases where the possibility of specifying individuals is inappropriate (as in some mass nouns; e.g. "a water," "all waters," "few waters"). Pronominal reference doesn't require quantification because it is supplied by the text that it indicates (transitivity of indication). But when general terms are substituted for pronouns or other words serving overt indexical functions (e.g. "this" or "that") they inevitably require the addition of quantification. There are also, of course, many other exceptions to the need for quantification. Proper names and numbers do not require quantification when they are used to refer to a type as a singular class because indicating would again be redundant.

The exception that proves the rule, so-to-speak, is exemplified by highly inflected and/or agglutinated languages where indexical marking is incorporated directly into word morphology. In comparison with English, which maintains the indexical grounding of most of its symbolic functions by strict word order constraints, these languages tend to have relatively free word order. This leads to a prediction: the

more completely that indexical functions are incorporated into word morphology the less restrictive the syntax and vice versa.

So approaching this issue semiotically provides a functional account that can unify a wide range of grammatical and syntactic relationships. It also suggests that our naïve intuition about these linguistic regularities may be more accurate than the formal rule-governed approach would suggest. A naïve speaker seldom comments that an ungrammatical sentence breaks a rule, and is generally hard-pressed to articulate such a rule. Rather the usual comment is that it just sounds wrong or that it doesn't make sense said that way. Compare these examples of breaking the contiguity of symbol and index to knowledge of the rules invoked to explain them:

Implicit subject:

* “_ Roundly shining over flowing shimmering.”

“_ Fire!”

“_ Hot!”

Island constraints:

“John found candy and gum in his shoe.”

* “What did John find candy and _ in his shoe?”

Priority in argument structure:

* “John found surprisingly in his shoe some candy.”

In these cases, and many others, naïve speakers know there is something wrong even if they can't articulate it, except to say that the ungrammatical sentences are awkward or difficult to interpret, and require some guesswork to make sense of them. Moreover, in everyday conversational speech, the so-called rules of grammar and syntax are only very loosely adhered to. This is usually because common interests and joint attention as well as culturally regularized interaction frames provide much of the indexical grounding, and so adherence to these strictures tends to be preferentially ignored. Not surprisingly, it was with the widespread increase in literacy that scholarly attention began to be focused on grammar and syntax, and with education in reading and writing these “rules” began to get formalized. With the written word shared immediate context, common pragmatic interests, and implicit presuppositions are minimally if at all available to provide indexical disambiguation and so language-internal maintenance of these constraints becomes more critical.

Finally, this semiotic functional analysis also provides an alternative understanding of the so-called poverty of the stimulus problem that is often invoked to argue that knowledge of grammar must be largely innate. Consistent with the fact that naïve speakers are generally unable to articulate the “rules” that describe their understanding of what is and is not a well-formed sentence, young children learning their first language are seldom corrected for grammatical errors (in contrast to regular correction of pronunciation). Moreover, children do not explore random combinatorial options in their speech testing to find the ones that are approved by others. They make remarkably prescient guesses. It has been assumed, therefore, that they must have some implicit understanding of these rules already available.

But in fact children do have an extensive and ubiquitous source of information for learning to produce and interpret these basic semiotic constraints on predication. First of all, discerning indexicality is a capacity that is basic to all cognition, animal and human. It requires no special training to become adept at the use of correlation, contiguity, etc., to make predictions and thus to understand indexical relationships. Second, although there is little if any correction of the grammar and syntax in children's early speech there is extensive pragmatic information about success or failure to refer or to interpret reference. This is in the form of pragmatic feedback concerning the communication of ambiguous reference. And this source of information attends almost every use of words. So I would argue that children do not "know" grammar innately, nor do they learn rules of grammar, and yet they nevertheless quickly "discover" the semiotic constraints from which grammars derive.

Although it is necessary to learn how a given language implements these constraints, the process is not inductive. It is not necessary for a child to derive general rules from many instances. Young children make good guesses about sentence structure—as though they already know "rules" of grammar—by tapping into more natural analogies to the nonlinguistic constraints and biases of iconicity and indexicality, and by getting pragmatic feedback about confused or ambiguous reference. Evolved predispositions to point or indicate desired objects or engage joint attention also make sense in this context. This universal human indexical predisposition provides the ideal scaffold to support what must be negotiated and must be progressively internalized to language structure. The early experience of communicating with the aid of pointing also provides additional background training in understanding the necessary relationship between symbols and indices.

Semiotic constraints should be agent-independent, species-independent, language-independent, and discourse-independent. They have been mistakenly assumed to be either innate structures or else derived from cognitive schemas or determined by sensorimotor biases and/or social communicative pragmatics. Though they are prior to language experience, and some are prerequisites to successful symbolic communication, they are neither innate nor socially derived. They are emergent from constraints that are implicit in the semiotic infrastructure of symbolic reference and interpretive processes. They are in this way analogous to mathematical universals (e.g. prime numbers) that are "discovered" (not invented) as mathematical representation systems become more powerful. Though each form of symbol manipulation in mathematics has been an invention and thus a convention of culture, we are not free to choose just any form if we want to maintain consistency of quantitative representation.

Assuming that symbolic reference lacks intrinsic structure has tricked linguists into assuming the need to postulate ad hoc rule systems and algorithms to explain the structural constraints of language. Failure to pay attention to the iconic and indexical underpinnings of symbolic reference has additionally exaggerated the complexity of the language acquisition problem. This myopic avoidance of semiotic analysis has led to the doctrine of an innate language faculty that includes some modicum of language-specific knowledge and this seeming logical necessity has supported an almost religious adherence to this assumption despite the biological implausibility

of its evolution and the lack of neurological support for any corresponding brain structures or functions. Unfortunately semiotic theory has not been of much assistance, primarily because it has remained a predominantly structural theory tied to a static taxonomic understanding of semiotic relationships. But when semiosis is understood as a process of interpretive differentiation in which different modes of reference are understood as dynamically and hierarchically constituent of one another these many conundrums dissolve and these once apparently independent aspects of the language mystery turn out to have a common foundation.

These constraints are the most ubiquitous influences on language structure, and indeed they are even more universal than advocates of mentalese could have imagined, because they are not human specific. They are universal in the sense that the constraints of mathematics are universal. They would even be relevant to the evolution of symbolic communication elsewhere in the universe. But they are not exceptionless rules. Different languages, everyday spoken interactions, and artistic forms of expression can diverge from these constraints to varying extents, but at the cost of ambiguity and confusion of reference. In general, these constraints will probably be the most consistent regularities across the world's languages because means to minimize this divergence will be favored by the social evolution-like processes of language transmission from generation to generation.

However, the universality and non-innateness of these constraints does not mean that there aren't human-specific constraints that contribute to many of the nearly universal regularities that characterize the World's languages. These are the result of constraints of a different sort; some deriving from our biology and some from social processes. None determine language organization in a generative sense, but rather along with semiotic constraints they collectively constrain and bias the range of possible language variations.

2.2.2 Processing Constraints

Probably the most critical factor contributing to the structure of natural languages in addition to semiotic constraints is the need to communicate symbolically in real time. Brains are not computers. They are slow and limited in mnemonic and attentional capacity, and symbolic communication is extremely demanding in both of these domains of brain function. Basically, the ability to use language in real time demands the equivalent of computational optimization. The key to Noam Chomsky's original insight into the structure of language can probably be characterized as recognizing that natural language syntax can be modeled as a Turing machine. In abstract form a Turing machine can be understood as a set of rules for writing erasing and rewriting strings of characters, in which these rules are also encoded as character strings that can be treated the same way. Rendering an operation in these terms makes it possible to automate any finite determinate process (from robotic behavior to mathematical calculations), which is why it contributed to the design of modern computers. Because of the power of this methodology, this insight was not only valuable for developing a formalism for modeling syntax, it also became a driving force for the development of the cognitive sciences.

So in one sense it isn't surprising that natural language structure can be modeled by this formalism, however, by using this approach the remarkable systematicity of languages could also be clearly exemplified. Language structure could have been far more haphazard than it is, but what formal approaches have demonstrated is that languages are remarkably internally consistent despite their flexibility. Instead, the syntax of a highly grammaticalized natural language resembles a formal system or Turing machine architecture where all operations are systematically inter-defined and precisely complementary to one another, and where many operations are almost entirely structure-dependent and content-independent.

This requires an additional explanation, since language structures have evolved spontaneously without any attention to their design logic or optimality. As noted above, this regularization is almost certainly the result of a kind of cultural version of natural selection involving language "traits," in which the selection pressures that determine which forms get passed on and which forms go extinct are the various constraints of referential effectiveness and ease of use. Ease of use is determined by what can be described as processing constraints. So what are these constraints?

Because human brains have distinct limitations due to the nature of neural signal processing as well as distinctive cognitive biases that have been inherited from our primate ancestry, the way they must solve the challenge of online real-time symbolic communication has both biological and computational idiosyncracies.

Probably the most ubiquitous processing constraints have to do with the amount of attentional and mnemonic work that must be done to produce and interpret linguistic communications. Linguistic tricks that enable the various symbolic operations to be most efficiently and thoroughly automated will for this reason be highly favored over the history of a language's persistence. Indeed, the demands of making language functions nearly effortless may even be favored at the expense of easy interpretability. Thus there will likely be linguistic selection over time for what might be described as optimal computational design. This is essentially what skill learning is all about. And in many respects skills are predominantly associated with motor functions.

Automatization of behavior is acquired by extensive repetition. As a behavior is repeated again and again in slightly different contexts those features that are least variable from performance to performance become more streamlined. In this sense the behavior becomes increasingly algorithmic. The key to automatization is simplification and specifically a reduction of options.

The challenge to automatization of language is that symbolic relationships are dependent upon determining relative "position" in a vast web of associations. Taking the time to sample this vast search space with each new combinatorial relationship to interpret would result in an impossibly slow rate of communication, likely straining short-term working memory. This is the source of what might be described as costs of symbolic search. Because symbols are nodes in complex systems of multi-dimensional semantic relationships, selecting and interpreting symbol combinations may involve a very high dimensionality search for appropriate "blend" relationships between them. The semantic search space grows exponentially with the number of elements combined, the dimensions to be considered, and the ambiguity of the

selection criteria. Contextual-pragmatic constraints (including extralinguistic indication of salient symbolic relations, and assessment of recipient knowledge/sender intention) may help to disambiguate the selection criteria and to reduce the search dimensions, but this cannot keep pace with the combinatorial explosion of the search space.

Partial automatization of language performance must therefore be achieved by strictly limiting the amount of symbolic search of memory that is required. So despite the power and flexibility of symbolic representation, the processes of selection at work during language transmission will tend to evolve means to reduce the density of symbolic operations per second in speech. Since the mnemonic and attentional demands of such a combinatorial search will depend on the numbers of dimensions of properties being represented, one way this can be reduced is simplification of certain common symbolic operations. Thus, improved automatization can be achieved by spontaneous linguistic evolution for what might be described as desymbolization; a spontaneous degeneration of the semantic dimensions of selected symbolic elements to the point that their reference is reduced to virtual indexicality (i.e. pointing to a single simple symbolic relation). This aids the efficient formation, identification, and parsing of sentential subassemblies with very shallow symbolic search.

This trade-off between processing constraints and symbolic combinatorial analysis is the source of a curious paradox: often languages tend to change (evolve) away from communicative transparency. The historical process of grammaticalization often reduces semantic and functional transparency rather than increasing it. Lexical specificity is often degraded (sometimes described as “bleaching”). Highly grammaticalized language can include phrase fragmentation, multileveled embedding, and non-contiguous syntactic relationships. This structural complexity often results in reduced learnability by undermining interpretive iconicity, such as direct mapping of the temporal order of events and relations to word order.

If linguistic selection favored only clarity of communication this would not make sense. And this is one reason that some (including Noam Chomsky) have argued that natural language grammar did not evolve for communication but rather only for cognition. But this ignores the importance of reducing processing load. Structural relationships are effectively indexical operations and as such they can become automatized by virtue of having highly regularized unambiguously singular functions. To the extent that these operations play an indexical role in disambiguating which semantic dimensions are relevant in a given symbolic combination they also reduce the cost of symbolic analysis.

Increased automatization appears to lead to a minimal set of mutually exclusive, fully reflexive, indexical operations, whether embodied in morphology or syntax. Not surprisingly, the lexicon of most languages tends to be segregated into content words and function words, as well as primary morphemes and affixes, with different balances between these. Where this is mostly achieved by distinct word classes and syntactic relations the content words, like nouns, verbs, adjectives, and adverbs, comprise an open set that can be indefinitely added to, as need requires. They play the symbolic roles within a sentence. The function words, such as pronouns,

determiners, prepositions, conjunctions, and the functional affixes like “-s” and “-ed” comprise a finite closed set and serve more-or-less inter-symbolic indexical functions, determining which semantic dimensions are relevant to consider when blending or differentiating symbolic relations.

In general we can distinguish between requirements for automatization and the various linguistic tricks to aid in meeting these requirements. Automatization requires a small, closed class of operators that are used in a stereotypic way, repeatedly (i.e. hundreds of thousands or millions of times each year). This invariant repetition is essential for developing a nearly unconsciously implemented skill. Moreover, any optional functions need to be reduced to no more than two or three. This may be aided by processes such as semantic bleaching, by agglutination or strict syntactic adjacency, by standardization of common thematic frames, by indexicalizing highly redundant and phylogenetically salient types (e.g. plurality, tense, possession, animacy, status, etc.), and so forth.

In semiotic terms, then, the index-symbol relationship also corresponds to a fundamental distinction between those aspects of language that can be automated and those that cannot, respectively. This has clear neurological implications.

In neuropsychological terms, automatization is characteristic of what is often called procedural memory. Procedural memories are mostly associated with highly regular activities or skills in which a sequence of component actions and assessments is made highly predictable and easily cued. These are effectively behavioral algorithms that have been acquired by constant repetition, to the point that they can be executed with a minimum of conscious monitoring.

Importantly, once a skilled behavior is well-ingrained it can be executed at a rate that is many times more rapid than if each component operation required monitoring. The result is that automated procedures tend to be automatically initiated by stereotypic cues, once initiated “run” autonomously to completion, are modular in the sense that dissecting them back into component actions is difficult if not impossible, and their structures tend to become inaccessible to introspection. These same characteristics have often been cited as evidence that grammatical functions must be innate, modular, and specific to language. Creating and executing procedural memory functions involves a distinct set of brain systems, typically associated with motor control: including particularly frontal cortex, striatal structures, and the cerebellum. A reciprocal connectivity and functional relationship between cerebral cortex and striatal structures is critical to both creation and implementation of such skilled autonomous operations.

In contrast to the procedural memory system that creates memories by constant repetition, the brain establishes memory traces of singular events or experiences using a very different set of interconnected systems. Remembering what you did immediately after breakfast two days ago, the structure of a narrative, or the meaning of a new technical term, cannot rely on extensive repetition to become ingrained. Recalling such one-off events or experiences or novel associations must therefore depend on a very different strategy for consolidation and recall. Instead of redundancy of performance or rehearsal consolidation of these memories must rely on redundancy of associations, i.e. linkage with many other related memories

by innumerable commonalities and correlations. This is sometimes referred to as episodic or associative memory and is critically dependent on relationships between the cerebral cortex and the hippocampus. The associative memory system is thus ideally suited as a substrate for the storage of open-ended associative information and the procedural memory system is ideally suited as a substrate for the storage of a finite corpus of modular automated procedures. An interesting correlate of this segregation of automatized versus associative features of language processing is that brain damage that predominantly involves striatum and spares cerebral cortex has been found to preferentially impair the first language but not the second in some bilingual patients (Fabbro & Paradis, 1995). This is probably because the second language was not nearly so well automatized.

This functional segregation explains why indexical syntactic functions are performed with minimal effort, are largely unavailable to introspection, and have more-or-less modular organization, and why the analysis of more complex combinatorial symbolic relationships takes more mental effort and is generally the focus of attention. But it also suggests a way that language may provide a fundamental restructuring of cognition compared to other species. Not only are syntactic operations subject to automatization, but word-sound memory more generally is acquired in childhood by untold millions of repeated exposures and productions. So, like the relatively automatic processing of syntax, the production and recognition of words is also acquired like a deeply ingrained skill. Very little attention is paid to analyze and minimal effort is required to produce the familiar words of one's language. But although they may thus be generated from procedural memory, they cue associative loci in associative (episodic) memory. In this way language enables procedural memory traces to cue associative memory traces reciprocally, linking mnemonic strategies that in other animals are probably only minimally interdependent, and primarily with respect to external cuing. In contrast, in humans this acquired functional interdependence of memory systems provides an unprecedented internal reciprocal cuing mechanism for organizing experience. This ability to use a repertoire of acquired procedures to reliably access and organize life-episodes and abstract ideas is likely a major factor contributing to the human preoccupation with narrative.

One benefit of developing a functional account of the nature and origins of these language structures is that it leads to explicit predictions about how the brain processes language. One of the disappointments of the last four to five decades of formal linguistic theory is that while it provided unprecedented precision in describing language structures it has not been particularly useful in providing predictions about how language is processed in the brain. Instead, predictions about human-unique brain systems, dedicated language structures, modular isolation of language capacities from other cognitive functions, and the primacy of linguistic function over surface implementation of these functions, have not borne fruit. In contrast, reflecting on these semiotic and processing constraints a number of predictions spring immediately to mind.

- Hypothesis 1. In general, the way that language functions are neurologically distributed and localized does not respect “linguistic logic,” but rather the processing logic determined by what constitutes a functional unit and how this can be manipulated.
- Hypothesis 2. The neural distribution of different classes of linguistic operations develops during language acquisition as different operations become automatized.
- Hypothesis 3. The linguistic evolution of more thoroughly grammaticalized forms aids the efficient distribution of language processing in the brain.
- Hypothesis 4. Pidgins are less able to be automated because they lack semiotic systematicity. Their functions will be more widely distributed in the brain, processing will be slower, and functions will be more transparently iconic or indexical in surface production.

This has implications for brain-language co-evolution. Whereas semiotic constraints do not evolve and have been ubiquitously present throughout hominid evolution, processing constraints have likely been subject to constant change and variations. Brains have been modified in evolution in response to both. But we need to keep them separate in our analysis, and indeed they are functionally and temporally (in evolutionary terms) asymmetrically related. Adaptations that aid processing of symbolic interpretive competence and linguistic communication are secondary to the presence of symbolic communication. The initial evolution of the symbolic capacity created the context in which both the socially evolved grammatical and biologically evolved neurological adaptations took place, though once this process was initiated adaptation to all of these constraints co-evolved as a complex to produce the modern symbolic species.

Finally, before turning to the issue of human evolution itself, we need to consider the last two sources of constraints affecting language structure: phylogenetic sensorimotor biases and the demands of communicative interaction. I will only very superficially outline these influences because they are somewhat more optional and contingent than semiotic and processing constraints, and therefore have less of a universal character and more of a context-dependent role to play in determining similarities shared by most languages.

An influential alternative approach to formal theories of grammar and syntax that takes a more functional perspective often travels under the name of cognitive grammar. Although this term is often used for a restricted programmatic approach to explaining grammar, I will here use it quite generically to describe all theories that explain grammatical operations as reflecting the structure of sensory, motor, behavioral, and social operations in their form, and thus arguing that grammatical and syntactic relationships have been motivated by these systems. This approach is also often used to explain for example the prevalence of visual metaphors and path-progress analogies built into vocabulary and syntax. It has even motivated a theory of the origins of subject and predicate functions of sentences based on the so-called “what” and “where” visual pathways in the brain. The basic idea is that the logic that organizes language structure is derived and abstracted from evolved embodied

cognitive schemas. I think that it is without question that many aspects of language structure, lexical organization, and descriptive schemas have been shaped by these distinctively human cognitive biases, and also that there may even be culture-specific biases of an analogous sort that have served as linguistic selection biases causing parallel and convergent linguistic evolution in diverse historical contexts. While more contingent on human species-peculiarities than semiotic and processing constraints, these biases have inevitably also contributed to certain of the near universal regularities of human languages.

One nearly universal characteristic of language is its oral-vocal medium. The demonstration that the manual languages which evolved in deaf communities are indeed full-blown natural languages exhibiting features common to most spoken languages has undermined the universality of this feature of language. Nevertheless it is taken as an uncontroversial fact that language evolved as a vocal process, though it may have initially originated in a more gestural form.

One very telling piece of evidence supporting this scenario is the highly atypical human facility for skilled vocal behavior that is almost entirely absent in other land mammals and only modestly developed in cetaceans and certain bird groups. In *The Symbolic Species* (Deacon, 1997) I argue that this ability depends upon some quite unprecedented neurological relationships and that such a radical functional change must have been driven by significant selection advantages. But why this unlikely medium? I think that the answer is that it afforded an optimal medium for mimicry, and for a means of communication whose entire repertoire of sign vehicles must be acquired socially ease of mimicry is critical. It turns out that, despite what gets said in folk zoology, consciously learned mimicry is quite uncommon among animals. What monkeys see they seldom do. There is one general exception to this paucity of learned mimicry: singing in some songbird species, sound mimicry in parrots, mynahs, and mockingbirds, and song transmission in humpback whales, and there are probably other examples as well. Why this exception for oral-vocal communication? I think that the answer is that sounds heard can be behaviorally approximated without any need for any mental transformation. In contrast, gestural behaviors that are observed require a mental inversion before being reproduced. One needs to, in effect, imagine being the other producing this behavior. This shift of perspective is apparently not a trivial cognitive transformation.

Because of certain highly conserved phylogenetic limitations of nervous system organization, it seems reasonable to expect that our non-symbolic ancestors had as little control over vocal articulation, as is the case for other primates, and so the early stages of symbolic communication may indeed have involved more of a gestural embodiment. But once symbolic communication became a critical part of human social organization there would almost certainly have been a significant advantage to being able to shift manually signed symbolic communication to the oral-vocal domain. For this reason, I see this particular universal trait to be a relatively late emerging biological adaptation for symbolic communication, but one that set the stage for many processing adaptations, because of the immense advantages it created for rapidly expanding the sign vehicle repertoire and its combinatorial possibilities.

Finally, there are what I would describe as significant communication constraints that have also contributed to the convergences of language features worldwide. These are often not formally considered to be linguistic issues but rather associated with socio-linguistic and anthropological domains. Nevertheless they do play constraining roles that have shaped languages and provided a source of evolved parallelisms. Most significant of these are what I would lump into the category of pragmatic constraints. Language is used to convey information, to affect others' behaviors, to establish and restructure social relationships, to acquire information, and so forth. These functions and many more are universal simply by virtue of the fact of serving ubiquitous human social needs, and so the way they shape the various modes of organizing and interpreting symbolic communications will also exhibit shared attributes. And in addition there will be culture-specific expectations and prohibitions about how communication is to be used and information is to be shared. In these socio-cultural domains these pragmatic needs and customs are probably more variable and less tightly constraining than any of the other factors discussed, but the universe of these possibilities is probably quite limited and so we should expect these pragmatic constraints to contribute some further degree of constraint as well.

In summary, we should expect that many aspects of language come to exhibit near universal properties even despite its superficial arbitrariness of referential correlation, but not because of any innate set of rules or algorithms that generate these features. Language universals are a reflection of the many constraints that derive from the semiotic infrastructure of symbolizing and the processing demands this entails. The semiotic universals should be reflected in the symbolic communication of any species, should it evolve this competence, whether on Earth or elsewhere in the universe. However, the processing constraints that have influenced the structure of human language are less universal. So, for example, were we ever to find a way to engineer symbolizing minds in silicon, using electronic instead of chemical and ionic means of signal processing, we should expect some very different structures to emerge.

2.3 The Evolutionary Conundrum Posed by Language

In my work I use the phrase, symbolic species, quite literally, to argue that symbols have literally changed the kind of biological organism we are. I believe that we think and behave in many ways that are quite odd compared to other species because of the way that language has changed us. In many respects symbolic language has become a major part of the environment to which we have had to adapt in order to flourish. In the same way that our ancestors' bodies evolved in the context of the demands posed by bipedal foraging with stone tools and incorporating meat into the diet, their brains evolved in the context of a rich fabric of symbolic cultural communication. As it became increasingly important to be able to enter into the social web of protolinguistic and other early forms of symbolic social communication in order to survive and reproduce, the demands imposed by this artificial niche would have selectively favored mental capacities that guaranteed successful access to this

essential resource. So rather than merely intelligent or wise (sapient) creatures, we are creatures whose social and mental capacities have been quite literally shaped by the special demands of communicating with symbols. And this doesn't just mean that we are adapted for language use, but also for all the many ancillary mental biases that support reliable access and use of this social resource.

But this claim depends on language-like communication being a long-time feature of hominid evolution. Theories suggesting that human language is a very recent and suddenly evolved phenomenon would not make this prediction. To them language is almost epiphenomenal. This is particularly true if the claim is that language appeared suddenly due to some marvelous accidental mutation that transformed dumb (but large brained) brutes into articulate speakers. This sort of scenario has become commonplace in recent years, though the evidence supporting it is mostly very indirect (e.g. archeological evidence of representational forms and objects for adornment, appearing in the Upper Paleolithic). I think that it is mostly a reflection of a caricatured view of the human/animal distinction and a sort of hero metaphor imposed upon the fossil evidence. The way that modern human brains accommodate language can be used as a clue to how old language is.

If language is a comparatively recent feature of human social interaction, that is if it is only, say, a hundred thousand years old or so, then we should expect that it had little effect on human brains. Any structural tweaks of brain architecture that evolved to support it would have had to be either minimal or else major but dependent on comparatively few genetic changes. A recent origin of language would give it little opportunity to impose selection pressure on human brains, so language function would not be supported by any widespread and well-integrated neurological changes. This would predict that language abilities are essentially an evolutionary after-thought, inserted unsystematically into an otherwise typical (if enlarged) ape brain. With little time for the genetic fixation of many supportive traits to occur, this adaptation would likely depend on only a few key genetic and neurological changes. As a consequence, language function should be poorly integrated with other cognitive functions, relatively fragile if faced with impoverished learning contexts, susceptible to catastrophic breakdown as a result of certain small but critical genetic defects, and severely affected by congenital mental impairment.

None of these seems to be the case.

On the other hand, if language has been around for a good deal of our evolutionary past, say a million years or so, that amount of time would have been adequate for the demands of language to have affected brain evolution more broadly. A large network of subtle gene changes and neurological adjustments would be involved, and as a result it should be a remarkably well-integrated and robust neurological function. Indeed, there is ample evidence to suggest that language is well-integrated into almost every aspect of our cognitive and social lives, that it utilizes a significant fraction of the forebrain, and is acquired robustly under even quite difficult social circumstances and neurological impairments. It is far from fragile.

The co-evolutionary interaction goes both ways. Languages also have to adapt to brains. Since the language one learns has to be passed from generation to generation, the more learnable its structures, and fitted to human limitations, the more effective

its reproduction in each generation. Languages and brains will evolve in tandem, converging towards each other, though not symmetrically. Brain evolution is a ponderously slow and unyielding process in comparison to the more facile evolution of languages. So we should expect that languages are more modified for brains than brains are for language. Nevertheless, if we have been evolving in a symbolic niche for a million years or more, we should expect that human brains will have been tweaked in many different ways to aid life in this virtual world.

The world of symbols is an artificial niche. Its ecology is radically different than the biological niche we also find ourselves in (or at least our ancestors found themselves in). In the same way that beaver dam building has created an aquatic niche to which beaver bodies have adapted over their evolutionary history, our cognitive capacities have adapted to our self-constructed niche: a symbolic niche. This is not a new idea. Indeed the anthropologist Clifford Geertz (1973) suggested something like this many decades ago. I think that today we may be at a point in our evolutionary theorizing and our understanding of brains to begin to explore exactly what this might mean.

The most intense and unusual demands of this niche should be reflected in the ways that human cognition diverges from patterns more typical of other species. Although it has long been popular to think of the human difference in terms of general intelligence, I think this bias may have misled us into ignoring what may be a more important constellation of more subtle differences. These likely included differences in social cognition (e.g. joint attention, empathy, the ability to anticipate another's intended actions), differences in how we learn (e.g. superior transfer learning, a predisposition to assume that associations are bidirectional—known as stimulus equivalence, a comparative ease at mimicking) or even just unusual motor capacities (e.g. unprecedented articulatory and vocal control). These are members of a widely distributed and diverse set of adaptations that fractionally and collectively contribute to our language abilities.

With respect to the brain, we need to confront another mystery. How could these many diverse brain traits have become so functionally intertwined and interdependent as to provide such a novel means of communication? This is particularly challenging to explain because language is in effect an emergent function, not some prior function just requiring fine-tuning. Our various inherited vocalizations, such as laughter, shrieks of fright, and cries of anguish, are comparatively localized in their neurological control (mostly subcortical) as are other modes of communication in animals. In comparison, language depends on a widely dispersed constellation of cortical systems, each of which can be found in other primate brains, but evolved for very different functions. These brain systems have become collectively recruited for language only because their previously evolved functions overlapped significantly with some processing demand necessitated by language, though evolved for quite different functions altogether. Indeed, the neural structures and circuits involved in the production and comprehension of language are homologous to structures found ubiquitously in most monkey and ape brains: old structures performing unprecedented new tricks.

A related mystery concerns the extent to which this dominant form of communication depends on information maintained by social transmission. Even for theories postulating an innate universal grammar, the vast quantity and high fidelity of the information constituting even a typical vocabulary stands out as exceedingly anomalous from a biological point of view. How did such a large fraction of our communicative capacity wind up offloaded onto social transmission? And what explains the remarkable reliability of this process?

Perhaps the most difficult neurological feature to explain, however, is the evolution of the diversity of brain structures involved. The higher-order synergy of systems that contribute to language requires the cooperative functioning of many diverse brain regions. And it appears to paradoxically require that this synergy among diverse systems must already be in place in order for selection to have honed it for language.

The co-evolutionary niche construction scenario sketched above still does not account for the generation of the novel functional synergy between neural systems that language processing requires. The discontinuities between call control systems and speech and language control systems of the brain suggest that a co-evolutionary logic alone is insufficient to explain the shift in substrate. Recent investigation of a parallel shift in both complexity and neural substrate in birdsong may be able to shed some light on this (see also Pepperberg, [Chapter 7](#), this volume).

In a comparative study of a long-domesticated bird, the Bengalese Finch, and its wild cousin, the White-Rump Munia, it was discovered that the domesticated lineage was a far more facile song-learner with a much more complex and flexible song than its wild cousin (discussed in detail in Deacon, 2010b). This was despite the fact that the Bengalese Finch was bred in captivity for coloration, not singing. The domestic/wild difference of song complexity and song learning in these close finch breeds parallels what is found in comparisons between species that are song-learners and non-learners. This difference also correlates with a much more extensive neural control of song in birds that learn a complex and variable song.

The fact that this behavioral and neural complexity can arise spontaneously without specific breeding for singing is a surprising finding since it is generally assumed that song complexity evolves under the influence of intense sexual selection. This was, however, blocked by domestication. One intriguing interpretation is that the relaxation of natural and sexual selection on singing paradoxically was responsible for its elaboration in this species. In brief, with song becoming irrelevant to species identification, territorial defense, mate attraction, predator avoidance, and so on, degrading mutations and existing deleterious alleles affecting the specification of the stereotypic song would not have been weeded out. The result appears to have been the reduction of innate biases controlling song production. The domestic song could thus be described as both less constrained and more variable because it is subject to more kinds of perturbations. But with the specification of song structure no longer strictly controlled by genetically inherited innate auditory and motor biases, other linked brain systems can begin to play a biasing role. With innate song biases weakened, auditory experience, social context, learning biases, and attentional factors could all begin to influence singing. The result is that the domestic

song became more variable, more complicated, and more influenced by social experience. The usual consequence of relaxed selection is genetic drift—increasing the genetic and phenotypic variety of a population by allowing random reassortment of alleles—but neurologically, drift in the genetic control of neural functions should cause constraints to become less specific, generating increased behavioral flexibility and greater conditional sensitivity to other neurological and contextual factors.

This is relevant to the human case, because a number of features of the human language adaptation also appear to involve a relaxation of innate constraints allowing multiple other influences besides fixed links to emotion and immediate context to affect vocalization. Probably the clearest evidence for this is infant babbling. This unprecedented tendency to freely play with vocal sound production occurs with minimal innate constraint on what sound can follow what (except for physical constraints on vocal sound generation). Babbling occurs also in contexts of comparatively low arousal state, whereas laughter, crying, or shrieking are each produced in comparatively specific high arousal states and with specific contextual associations. This reduction of innate arousal and contextual constraint on sound production, opens the door for numerous other influences to begin to play a role. Like the domesticated bird, this allows many more brain systems to influence vocal behavior, including socially acquired auditory experience. In fact, this freedom from constraint is an essential precondition for being able to correlate learned vocal behaviors with the wide diversity of objects, events, properties, and relationships language is capable of referring to. It is also a plausible answer to the combinatorial synergy problem (above) because it demonstrates an evolutionary mechanism that would spontaneously result in the emergence of multi-system coordination of neural control over vocal behavior.

But although an evolutionary de-differentiation process may be a part of the story for human language adaptation, it is clearly not the whole story. This increased flexibility and conditionality likely exposed many previously irrelevant interrelationships between brain systems to selection for the new functional associations that have emerged. Most of these adaptations remain to be identified. However, if such a dedifferentiation effect has been involved in our evolution, then scenarios hypothesizing selection for increased innateness or extrapolation from innate referential calls to words become less plausible.

There is a much larger biological background behind my approach which of necessity has had to go unmentioned. It traces to my work on brain development and evolution, and more broadly it borrows from work that currently runs under the banner of “evodevo” and which has begun to illuminate once problematic issues in evolutionary genetics, molecular cellular biology, and epigenesis. My point is not to discount the contributions of natural selection, which I agree is the final arbiter of functional adaptation, but to bring attention to another unnoticed facet of the evolutionary process. Natural selection is explicitly NOT the generator of the biological phenomena that it prunes in the process that leads to increased adaptation. Not only are the variants of existing organismic subsystems generated irrespective of function (e.g. by genetic “damage”) but the expression of these varieties of structure and dynamics depends on generative processes whose details we tend to hide in generic

concepts like epigenesis and reproduction. New stuff, new structures, and new processes need to be generated so that there is raw material fed to the engine of natural selection. The second law of thermodynamics has to be locally tamed in order for this to be possible. And natural selection theory is so widely applicable precisely because it can be agnostic as to how any of this is achieved, so long as it is.

Surprisingly, despite our many disagreements about innateness, I find some resonance in Noam Chomsky's periodic suggestion that some of the complexity of grammar may have emerged from general laws of physics analogous to the way that the Fibonacci regularities exemplified in the spirals of sunflower seed and pine cone facets emerge. Natural selection has "found a way" to stabilize the conditions that support the generation of this marvelous regularity of growth because it has important functional advantages. But natural selection didn't generate it in the first place, geometric regularities that can become amplified due to center-out growth process are the ultimate source (as has now been demonstrated also in growth-like inorganic processes).

In closing, I would like to reflect on some of the more esoteric features of humanness that may be illuminated by the paired processes of symbolic niche construction effects and relaxed selection.

For example, I think it makes sense to think of ourselves as symbolic savants, unable to suppress the many predispositions evolved to aid in symbol acquisition, use, and transmission. In order to be so accomplished at this strange cognitive task, we almost certainly have evolved a predisposition to see things as symbols, whether they are or not. This is probably manifest in the make-believe of young children, the way we find meaning in coincidental events, see faces in clouds, are fascinated by art, charmed by music, and run our lives with respect to dictates presumed to originate from an invisible spirit world. Like the flight play of birds, the manipulation of objects by monkeys, the attraction of cats to small feathered toys, our special adaptation is the lens through which we see the world. With it comes an irrepressible predisposition to seek for a cryptic meaning hiding beneath the surface of appearances. Almost certainly many of our most distinctive social capacities and biases—e.g. tendencies to conformity and interest in copying the speech we hear as infants—are also reflections of this adaptation to an ecosystem of symbolic relationships. And of course there is literature and theater. How effortlessly we project ourselves into the experiences of someone else, feeling the joys and sorrows almost as intensely as our own.

Relaxation of selection, on the other hand, may have contributed to another suite of distinctively human traits. Widely distributed dedifferentiation at the genetic and epigenetic level would have increased flexibility of a variety of once phylogenetically constrained cognitive and motivational systems. Perhaps the most striking feature of humans is their flexibility and cultural variety. Consider the incredible diversity of marital and kinship organizations. Most species have fairly predictable patterns of sexual association, kin association, and offspring care, and although they are somewhat flexible, this variety is mediated almost entirely by individual motivational systems. In contrast, despite the evolutionary importance of reproduction, human mating and reproduction are largely controlled by symbolically mediated

social negotiations. This offloading of one of the most fundamental biological functions onto social-symbolic mechanisms is perhaps the signature feature of being a symbolic species. Thus, because of symbols and with the aid of symbols, *Homo sapiens* has been self-domesticated and adapted to a niche unlike any other that has ever existed.

We have been made in the image of the word.

References

- Bickerton, D. (2010). Commentary on Terrence Deacon (2010) *On the human: Rethinking the natural selection of human language*. Retrieved August 24th 2011, from <http://onthehuman.org/2010/02/on-the-human-rethinking-the-natural-selection-of-human-language/>
- Deacon, T. (1997). *The symbolic species: The coevolution of language and the brain*. New York: W. W. Norton & Co.
- Deacon (2010a). *On the human: Rethinking the natural selection of human language*. Retrieved August 24th 2011, from <http://onthehuman.org/2010/02/on-the-human-rethinking-the-natural-selection-of-human-language/>
- Deacon (2010b). A role for relaxed selection in the evolution of the language capacity. *Proceedings of the National Academy of Sciences*, 107, 9000–9006.
- Everett, D. (2005). Cultural constraints on grammar and cognition in Pirahã: Another look at the design features of human language. *Current Anthropology*, 46, 621–646.
- Fabbro, F., & Paradis, M. (1995). Differential impairments in four multilingual patients with subcortical lesions. In M. Paradis (Ed.), *Aspects of bilingual aphasia* (pp. 139–176). Oxford: Pergamon Press.
- Geertz, C. (1973). *The interpretation of cultures*. New York: Basic Books.
- Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 29, 1569–1579.
- Peirce, C. S. (1931). *Collected papers of Charles Sander Pierce. Vol. II elements of logic* (C. Hartshorn & P. Weiss, Eds., CP 2.225). Cambridge, MA: Harvard University Press.