

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/27450463>

How language changed the genes: toward an explicit account of the evolution of language

Article · January 2001

Source: OAI

CITATIONS

11

READS

137

2 authors:



Daniel Dor

Tel Aviv University

28 PUBLICATIONS 641 CITATIONS

SEE PROFILE



Eva Jablonka

Tel Aviv University

128 PUBLICATIONS 7,903 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



The Social Origins of Language [View project](#)

How language changed the genes: toward an explicit account of the evolution of language

Daniel Dor and Eva Jablonka

1. Introduction

This paper represents an attempt to construct a programmatic framework for the evolution of human language. The theory we suggest is interdisciplinary in its essence. It is based on what we believe are noteworthy recent advances in linguistics and evolutionary biology. The theory aims to resolve three fundamental paradoxes frequently encountered in the burgeoning literature on language evolution. These paradoxes leave many of the traditional conceptions of language evolution at a theoretical dead end and inhibit interdisciplinary cross-pollination among linguistics, evolutionary biology, anthropology, and the neurosciences.

1.1. The functional paradox

Two radically opposite claims have traditionally been made in the linguistic literature concerning the functional nature of linguistic knowledge. As far as language evolution is concerned, both claims seem to lead to a theoretical impasse. Generative Grammar, by far the most influential linguistic theory of the last forty years, is famous for making the claim that the aspects of language that form the universal, innate basis of linguistic knowledge (Universal Grammar) are essentially structural, formal, and autonomous from notions of meaning. As such, they are explicitly non-functional. According to this view, what we can do with language—like communicating meanings to each other—cannot tell us anything significant about the linguistic system itself. This implies that a neo-Darwinian theory of language evolution is unattainable because such a mode of evolutionary explanation requires a satisfactory functional characterisation of the relevant evolved trait (here: the linguistic capacity) as a descriptive platform. This is why Noam Chomsky insists that attempting to say anything meaningful about language evolution is a waste of time. Different writers, like Pinker and Bloom (1990), Newmeyer (1998), Berwick (1998), and Jackendoff

(1999), have tried to demonstrate that Generative Grammar does make evolutionary sense, but we believe that Chomsky is right: from the evolutionary perspective, his innateness claim cannot be reconciled with his specific characterisation of language as a non-functional cognitive apparatus.

The opposing view, according to which grammatical complexities are reducible to principles of general cognition, at first seems more tenable, since it implies a functional characterisation of linguistic knowledge. But on closer examination it does not do much better than the Generative view, for at least three complementary reasons. First, although the specific functional theories that attempt to reduce linguistic complexities to general cognitive principles (Givón 1995, Langacker 1987, and Wierzbicka 1988) propose preliminary approximations of the cognitive precursors of linguistic complexities, they do not usually provide rigorous explanations for specific grammatical facts. Therefore, they do not offer a real alternative to the Chomskian characterisation of the object of evolution. Second, attempting to reduce language to general cognition is problematic. There is ample evidence—from language acquisition, language breakdown, and the formation of *de novo* languages (like the sign language developed by deaf Nicaraguan children)—that language is a unique and highly specialised cognitive capacity. Whereas Chomsky's theory captures the uniqueness of language at the expense of its functionality, cognitive theories attempt to salvage language's functional aspects at the expense of its uniqueness. Third, the cognitivists' characterisation of language as a general-purpose communication tool is too broad. Asserting that we use language to communicate is similar to asserting that we use our vision to see. As Marr (1982) has convincingly argued, a functional characterisation of any cognitive system must be much more specific. A real functional characterisation of language should be both empirically viable and functionally specific.

1.2. The paradox of domain-specificity

Linguists unanimously concur that linguistic knowledge is extremely domain-specific; that is, it manifests properties not found in other behavioural systems. The usual extension of this understanding to the domain of language evolution goes as follows: if language is domain-specific—and if at least some of our linguistic ability is innate—, then some of this domain-specificity should be reflected in brain structures. Ironically, the one clear assertion neuroscientists seem to agree on is that the brain is an organ of

extreme plasticity and generality (see Elman et al. 1996 and Deacon 1997), which means that the chances are slim of finding explicit representations of linguistic specificities innately encoded in brain tissue prior to acquisition. Characteristically, linguists either subscribe to domain-specificity or adopt the neurophysiological position of non-specificity and resort to a general learning theory to account for the acquisition of linguistic specificities from external input. This is paradoxical, since both linguistic specificity and brain plasticity seem to be irrefragable, as is the notion that language acquisition cannot be based solely on external input. So what we need is a theory of the relationship between language, genes, and the brain that a) reconciles language's domain-specificity with the brain's high level of plasticity and b) relates them to each other in non-contradictory ways, both in phylogeny and ontogeny.

1.3. *The paradox of the dynamic and variable nature of language*

Most scholars who believe in linguistic innateness adhere to a static and universalistic conception of language. The generative theory of *principles and parameters* is the most famous such conception: even when some variability between languages is conceded, the theory encodes the variability in the genes. Children come to the world with a few parameters for each linguistic principle, and choose the right one for the language they encounter. This notion faces what we consider to be insurmountable difficulties. First, it cannot account for the fact that languages are dynamic entities that constantly change and evolve in their social contexts. Second, it cannot easily account for the considerable diversity among different languages. Third, it must posit a genetic mutation that enabled an individual in a hominid community to use the full range of future languages at a stage when no languages existed, and it must then assume that this property spread across the entire community. This is hardly a plausible evolutionary scenario.

It is our goal in this paper to suggest a framework that successfully resolves these paradoxes, explicates the theoretical prerequisites for an empirically viable theory of language evolution, and describes what we think are the fundamental properties of the evolutionary process. Our argument runs as follows:

A. Our point of departure is the functional paradox. We will claim that a series of empirical results, accumulated in the last two decades in the field of linguistic semantics, strongly indicate that Chomsky's long-standing

hypothesis of the autonomy of syntactic structures from meaning regularities should be abandoned in favour of an explicit, semantically based, and empirically oriented theory of transparent meaning-form relations.

B. Crucially, the specific semantic categories which turn out to determine syntactic regularities in languages are *not* reducible to general cognition or to conceptual structure. Instead, they seem to manifest uniquely linguistic properties. As such, they determine the expressive envelope of language. That is, they determine which meanings—and which meaning combinations—are expressible by means of natural language structures.

C. This state of affairs allows for a novel characterisation of language as a cognitive capacity: neither a non-functional, formal system (the generative characterisation), nor a functional, general-purpose communication system (the functionalists' characterisation), but a unique and highly constrained communication system dedicated to the communication of a specific set of meanings. According to this characterisation, language is not just functional and unique. Its uniqueness is *in* the specificity of its function. In line with traditional cognitive psychology, we characterise language as a functional, unique, and transparent *mapping system* between the representational level of *linguistic meaning* (the set of constitutive semantic categories) and the representational level of *linguistic form* (the set of grammatical markers for these meanings).

D. This characterisation of language enables us to reframe the evolutionary question as the question of the gradual *expansion* and *sophistication* of the linguistic mapping system; that is, the expansion and sophistication of the set of constitutive semantic notions, their interactions, and their modes of mapping onto the speech channel. Note that this is not the same as the question of the origin of language, that is, the question of the evolution of language's cognitive precursors, and the first "leap" from these to what we think of as natural language. We will concentrate on the question of evolution (as formulated above) and leave the question of origin for further research.

E. The above reframing of the evolutionary question naturally calls for an answer in terms of the interaction between cultural and genetic evolution. Here, a distinction should be made between three closely related questions. How did language evolve? How did speakers (and linguistic ability) evolve? How did these two processes interact?

F. According to our theory, cultural evolution played a major role in the evolution of language. The process of cultural linguistic evolution consisted of the selection of, the social agreement on, and the cultural evolution of the

constitutive semantic categories for communication, as well as the gradual sophistication of the mapping system for these categories. In this long, gradual, and complex process, a social group gradually isolates certain aspects of its epistemology, sharpens and develops them, reaches social agreement about them, and develops sophisticated structural means for communicating about them within the community. Needless to say, this is a permanent process, which continues today.

G. Throughout most of evolution, this process was made possible by homonids' great behavioural plasticity, which includes all aspects of linguistic behaviour: production, comprehension, acquisition, and transmission.

H. Crucially, however, at various points in evolution, cultural evolution stretched speakers' behavioural plasticity to the point where differences in the ability to learn language became selectively important. At these points, genetic assimilation occurred—on all cognitive fronts. At every step, linguistic culture constituted the selective environment for genes that contributed to linguistic performance, acquisition, and transmission. The interaction between continuous, directional cultural evolution and partial genetic assimilation resulted in a consecutive set of evolutionary stages in which language's expressive envelope was expanded and sophisticated and in which speakers were selected on the basis of their linguistic performance. This process of cultural evolution and genetic assimilation gradually created what we think of as a *linguistically biased* cognition: a cognitive makeup which, without encoding linguistic specificities on a genetic basis, is still biased toward rapid learning of the linguistic mapping system.

I. As we will claim, our approach successfully resolves the three above-mentioned paradoxes, and has several additional, non-trivial consequences. It transforms the traditional dichotomy between innateness and learning into the question of *how much, when, and what type of learning* is necessary at each evolutionary stage. It renders the *continuity-discontinuity* debate obsolete by rejecting the idea that there is a relevant distinction between a syntactic and a presyntactic language. And it reconciles socially oriented as well as structurally oriented approaches to language evolution and incorporates significant insights from linguistic anthropology into the universalistic discourse on language evolution.

2. On the nature of grammars

Our point of departure is a theoretical reappraisal of Chomsky's long-standing hypothesis of the autonomy of syntactic structures from meaning considerations. As we have already indicated, recent empirical research on the interface between syntactic and semantic representations consistently demonstrates the dramatic extent to which syntactic phenomena are determined by semantic regularities. These studies strongly suggest that Chomsky's hypothesis should be abandoned in favour of an explicit, semantically based, and empirically oriented theory of transparent meaning-form relations. Before rehearsing one such demonstration, we need to make a few general observations about Chomsky's autonomy hypothesis.

First, Chomsky's hypothesis is not a mere existence claim. It does not claim that natural languages display purely structural syntactic phenomena. This assertion would be self-evident because such structural phenomena—one need only think of word order—obviously exist. The autonomy hypothesis goes much further. It claims that a significantly large set of core syntactic phenomena in natural languages cannot be theoretically correlated with a corresponding set of functions formulated in terms of meaning. The autonomy hypothesis takes the non-functionality of syntactic structures to be a necessary and fundamental property of human language, not a peripheral, accidental property of this or that construction in a specific language. From an evolutionary perspective this is crucial because according to Chomsky the set of non-functional syntactic generalities is innately given.

Second, the autonomy hypothesis is ultimately irrefutable. It is a negative claim: the claim that a significantly large set of core structural phenomena is immune to explanation in terms of a meaning-based theory. Any attempt to empirically refute such a hypothesis should, in principle, provide for total coverage of the complete set of structural-syntactic phenomena of an entire language—and preferably all languages. It should demonstrate that every non-accidental syntactic phenomenon can be correlated with a coherent functional theory of meaning. Because no linguistic theory, regardless of its ideological stamp, comes close to providing a fully explicit description of a single language (let alone of the universal parameters of language as such), the autonomy hypothesis is in no danger of refutation.

This means that the fate of the autonomy hypothesis should be decided on the basis of empirical demonstration. We should take some complex, syntactic phenomena—which everybody agrees are significant and non-accidental and which seem to be divorced from any considerations of meaning—and

show that they can be given an explicit meaning-based explanation. Then we should do it again. And again. To the extent that the analyses turn out to be empirically sound, they demonstrate that an explanation of the relevant structural phenomena in terms of meaning is not just possible, but necessary. To the extent that the phenomena constitute uncontroversial, core domains of linguistic knowledge, such meaning-based analyses gradually weaken the autonomy hypothesis as a default assumption. They strengthen the suspicion that the autonomy hypothesis reflects a stage in the history of the misunderstanding of the nature of meaning rather than a profound insight into the nature of language.

In the last two decades, semanticists and lexical semanticists have generated empirical demonstrations of exactly this type (see Alsina, Bresnan, and Sells 1997; Butt and Geuder 1998; Dor 1996, 1999, and 2000a; Dowty 1979 and 1991; Frawley 1992; Goldberg 1995; Grimshaw 1990; Jackendoff 1983 and 1990; Levin 1993; Levin and Rappaport 1991 and 1995; Parsons 1990; Pustejovsky 1991 and 1995; and Van Valin 1997). Some scholars, ourselves included, believe that enough results have accumulated for us to entertain the idea that grammars are not autonomous from meaning. Below we briefly present one empirical demonstration regarding what is probably the most famous set of syntactic facts discussed in the generative literature, namely *island constraints*. This cursory presentation is meant to provide a general sense of the proposed explanation and its functional nature (see Dor 2000a for details).

2.1. *Island constraints*

Island constraints were discovered by John Ross (1967) and have since acquired a uniquely prominent status in syntactic theory. The different syntactic mechanisms postulated through the years to (partially) grasp the phenomena have been extremely complex, and island constraints have come to be known as the standard example of syntax's innately given complexity and of its autonomy from meaning. The basic facts are demonstrated in the following examples:

Languages like English allow long-distance transformations of the type that we see in (1) below. Moving the *wh*-word from its "natural position" (after the verb *meet*) to sentence-initial position seems to be possible regardless of the distance between the original position and the final one:

154 Daniel Dor and Eva Jablonka

- (1) a. *Who did John meet at the supermarket?*
 b. *Who did Bill say that John met at the supermarket?*
 c. *Who did Mary think that Bill said that John met at the super market?*
 d. *Who did George realise that Mary thought that Bill said that John met at the supermarket?*

Crucially, however, some transformations are not allowed. For example, it seems that moving an NP out of some constructions is ungrammatical—at least in most cases. Below are two of the most famous constraints. First, it seems that NPs cannot be moved out of complex NPs:

- (2) The complex NP constraint:
 a. *John kissed* [_{np} *the girl who delivered* [_{np} *the pizza*]].
 b. **What did John kiss the girl who delivered?*

Second, it seems that NPs cannot be moved out of subjects:

- (3) The subject constraint:
 a. [_{subj} *Your obsession with* [_{np} *Madonna*]] *annoyed your father.*
 b. **Who did your obsession with annoy your father?*

Within the generative framework, facts like those in 2) and 3) have traditionally been explained by structural means. Informally, the different versions of the theory have taken syntactic constituents to be structural barriers to extraction. Obviously, it is difficult to imagine the functional significance of such structural barriers. To the extent that constraints on extraction are taken to be innately given—and they certainly are within the generative framework—then the question of their evolution becomes a prime example of the functional paradox.

Dor (2000a), however, explains the relevant set of facts on a semantic basis. According to his theory, the grammatical extractions meet a semantic condition that the ungrammatical ones do not. The semantic condition has to do with the interaction between two semantic domains: *event structure* and *epistemic licensing*.

Informally, when a speaker performs an assertive speech act—when she tells an interlocutor, for example, that *John kissed the girl who delivered the pizza*—the speaker performs two types of speech acts at once. She provides the interlocutor with the *description of an event* (like the event in which *John kissed the girl who delivered the pizza*) and she makes the epi-

stemic claim that this event took place in the world. The theory of event structure addresses the first part of the speaker's speech act. It specifies the semantic properties of events described by natural language sentences and the semantic properties of the participants in these events. Thus, informally, our speaker's sentence describes an event with the following representation:

- (4) There is an event *e*,
 e culminates before now
 e is an event of the type kissing
 the agent of *e* is John
 the patient of *e* is the girl who delivered the pizza

The agent and the patient in (4) are the thematic constituents of the described event (there are of course many other possible thematic constituents).

The theory of epistemic licensing has to do with the second part of the speaker's speech act: her claim that the event actually took place. Making the factual claim that the event actually took place in the real world is just one of the speaker's many epistemic options. She may, for example, say any of the following sentences:

- (5) a. *I believe that John kissed the girl who delivered the pizza.*
 b. *Bill told me that John kissed the girl who delivered the pizza.*
 c. *Mary saw John kissing the girl who delivered the pizza.*

In (5) a, the speaker tells the interlocutor that she has some reason to believe that the event actually took place. In (5) b, she tells the interlocutor that her epistemic claim is based on Bill's epistemic claim regarding the kissing event. In (5) c, the speaker tells the interlocutor that Mary actually witnessed the event as it occurred and that this forms the basis for her epistemic claim. In all these cases (and in many others) the kissing event is epistemically licensed. The speaker specifies who is responsible for the claim that the event took place and the type and strength of the claim.

Dor (2000a) proposes that the syntactic behavioural patterns we saw in examples (1) to (3) are best captured by the interaction of these two pieces of semantic theory. This is Dor's transformation rule: a transformation can apply to a syntactic constituent if and only if it refers to a *thematic constituent of an epistemically licensed event*.

We have already seen that the thematic constituents of our kissing event are the agent (*John*) and the patient (*the girl who delivered the pizza*). The sentence is epistemically licensed by the speaker's factual claim, which means that a question transformation can apply to the syntactic constituents which refer to the agent and the patient of the event.

- (7) a. *Who kissed the girl who delivered the pizza?*
 b. *Who did John kiss?*

Moreover, because the kissing event is also epistemically licensed in all the examples in (5), transformations can also apply to its thematic constituents:

- (8) a. *Who did you believe kissed the girl who delivered the pizza?*
 b. *Who did Bill tell you that John kissed?*
 c. *Who did Mary see John kissing?*

Crucially, however, the pizza does not participate in the epistemically licensed event. It plays a thematic role in the delivering event, but not in the kissing event. The NP *the pizza* thus cannot be transformed:

- (9) **What did John kiss the girl who delivered?*

A similar type of explanation applies to the examples in (3): The participants in the *annoying* event are the experiencer (*your father*) and the source (*your obsession with Madonna*). Madonna herself is not a participant, and the question transformation in (3) b is thus ungrammatical.

Despite its brevity, the above presentation is, we hope, specific enough to provide a sense of the theory's meaning-based nature. The relevant set of constraints on syntactic transformations in English is captured by a semantic constraint on the performance of the interrogative speech act. It seems that the interrogative speech act can only target the thematic constituents of an event that is epistemically licensed by the speaker. This is a functional theory in at least two complementary ways. First, it correlates the syntactic complexities with a well-defined set of semantic complexities, which we need to independently assume for the sake of semantic and pragmatic analysis. Second, the specific semantic categories that determine the relevant syntactic complexities are categories of linguistic communication. They are categories relating to the communication of information about events. When we tell each other about events in the world, this is exactly what we do.

We provide descriptions of the events and their participants, and we evaluate the descriptions' epistemic validity. If island constraints can be explained on the basis of a set of semantic categories directly related to linguistic communication, then they are far from evidencing the autonomy of syntax from meaning. Indeed, they demonstrate the functional nature of grammatical structures.

Does this refute the autonomy hypothesis? Strictly speaking, it does not. There remains a long list of syntactic phenomena that has not yet been explained in terms of meaning. Nevertheless, island constraints are often invoked as the prototypical example of the autonomy of syntax. And together with numerous other semantically based explanations of other major syntactic complexities, the above theory significantly weakens the autonomy hypothesis. Every advance in this field strips yet another syntactic phenomenon of its autonomous appearance. We think enough demonstrations have accumulated for us to reject Chomsky's hypothesis and to see where a transparent view of the syntax-semantics interface will take us.

3. On language as a unique mapping-system

Event structure and epistemic status are not the only two semantic domains that determine grammatical structures in natural languages. Logical categories like negation and conditionals determine some patterns of structural behaviour (such as negative polarity items); spatial categories (categories of spatial relations between physical entities) determine some behavioural patterns of verbs and prepositions; categories of time interact with event structure to determine aspects of the syntax and morphology of tense and aspect; classifier categories determine morphological patterns in many languages; and so on. Some pragmatic categories, such as topic and focus, seem to play a similar role, as does the inventory of speech acts. Some of these semantic and pragmatic categories seem to be universally relevant and to determine aspects of grammar in all languages. Others seem to be particular to a subset of languages.

Where do the semantic (and pragmatic) categories that determine syntactic regularities in natural languages belong? What is their cognitive status? Some scholars, most notably Jackendoff (1983 and 1990), assign them to the conceptual domain and view them as conceptual categories. To us, the evidence suggests that these categories belong to a *uniquely linguistic* level of meaning representation, one closely related to, and yet crucially different from, conceptual structure (see Dor 2000b). Why?

First, the semantic categories that play a constitutive role in determining syntactic generalisations in natural languages seem to belong to a *small subset* of all the semantic categories we use to think about and conceptualise the world. Even more significantly, these semantic categories cut across conceptual structures in ways that seem arbitrary from a conceptual viewpoint. Some semantic categories are grammatically relevant across many languages, whereas others are not. To take just one example, Frawley (1992) compares two important and robust distinctions of meaning: the distinction between *natural* and *nominal* kinds and the distinction between *animate* and *inanimate* objects. The former distinction relates to the difference between common nouns that denote inherently and those that denote compositionally. It is visible throughout the lexicon. *Tiger, gold, hepatitis, heat, pain*, and *red* are natural kinds. *Car, wheel, coat, wedding, divorce*, and *president* are nominal kinds. This distinction plays a central role in our cultural conceptualisation, but no language appears to mark the distinction by any sort of formal device. The distinction between animate and inanimate objects, on the other hand, is extremely relevant for linguistic structure: "the linguistic evidence shows that in every language there appears to be some grammatical reflex of the difference between animate and inanimate objects" (Frawley 1992: 9). Frawley concludes that the "fundamental question of philosophical semantics—what kinds of meaning are possible—contributes to the identification of a variety of potential meanings that language may encode. But only some of the results of an inquiry driven by this question are relevant. Not all possible meanings are grammaticalised; not all have empirical status" (Frawley 1992: 10).

Second, constitutive semantic categories always determine the behavioural patterns of *word classes* and not of *individual words*. Verbs, for example, belong to verb classes. The members of each class denote different versions of the same event type. The members of the class of surface-contact verbs, for instance, denote different versions of the same event type: movement in contact with a surface. They include *sweep, wipe, scrape, scratch*, and *scrub*. The class of change-of-state verbs includes *break, smash, crash, fracture*, and *shatter*. They all denote different versions of an activity resulting in a change to the physical state of the patient (Levin 1993). As many researchers have demonstrated, all members of the same semantic class display the same pattern of syntactic behaviour. The meaning distinctions between the different members of the same verb class—between, say, *sweep* and *wipe* or between *break* and *smash*—are syntactically irrelevant. That is, there does not seem to be a syntactic generalisation

that is sensitive to these meaning distinctions. It is as if all such verbs are grammatically identical. Conceptually, of course, we do distinguish between sweep and wipe, and between break and smash. But as far as the structure of language is concerned, *sweep* and *wipe* are indistinguishable. Grammar only isolates their event type, which is thus a constitutive determinant of their structure. We therefore must adopt two lexical representations of the verbs' meanings: a fully detailed conceptual representation and a skeletal linguistic-semantic representation.

Third, the semantic categories which turn out to determine grammaticality judgements do not constrain our ability to assign *conceptual interpretations* to ungrammatical sentences. We are fully capable of understanding what the non-grammatical transformations in (2) and (3) were supposed to mean. There is nothing inherently illogical or conceptually impossible about asking a question concerning the pizza the girl delivered or Madonna as the cause of your father's annoyance. To take another example, Rappaport-Hovav and Levin (1998) use the semantic properties of a large number of verb classes (including change-of-state verbs and surface-contact verbs) to account for intriguing phenomena like the contrast between the grammaticality of *Terry swept the crumbs off the table* and the ungrammaticality of **Terry broke the vase off the table*. The point is that we can readily imagine a situation where the vase was glued to the table and Terry removed it from the table by breaking it. The problem here is not with the ability to understand a certain event conceptually, but with the ability to describe it linguistically. This is a telling fact about the semantic categories of language, not about the conceptual categories in which we think.

Finally, the semantic categories that determine syntactic generalisations seem to manifest discrete (or digital) properties, whereas conceptual categories seem to manifest continuous (or analogue) properties. Take, for example, the category of animacy. Making a conceptual decision about the animacy of some entity is not always easy. Conceptually, animacy is a continuous category. Some entities are prototypically animate, some are prototypically inanimate, and others are somewhere in between. But the linguistic category of animacy is not continuous. As far as grammar is concerned, an entity should be either animate or inanimate. This discreteness is manifested in the structural markers of animacy in languages. Languages that mark animacy in their morphology force an obligatory, discrete choice on speakers, a choice that speakers do not necessarily have to make when they conceptualise.

All the above observations seem to point to the same conclusion. Linguistic-semantic categories comprise a constrained subset of all possible conceptual categories, a subset that is systematically highlighted, foregrounded, isolated, and digitalised for the purposes of linguistic communication. This set of linguistic-semantic categories determines language's *expressive envelope*. It determines which meaning combinations can be expressed in language, and which meaning combinations cannot. This conclusion echoes a major insight discussed by Aitchison (1996), Levinson (1997), and others: language is not the best tool to communicate all meanings. Some meanings are better communicated by means of visual imagery, music, body language, and mime. Other meanings, especially narrative ones, are best communicated by language. They constitute its expressive envelope.

All this allows for the construction of a novel characterisation of language as a unique and functional system: a transparent mapping-system between the set of constitutive semantic notions (that determine its expressive envelope) and the set of linguistic markers used to express these meanings. The following is a schematic representation of this view of language:

[conceptual representations] ↔ || [linguistic meaning] ↔ [linguistic form] || ↔ [phonetic representations]

The above schematic representation characterises language as a transparent mapping-system between the levels of linguistic meaning and linguistic form. The level of linguistic form, which we have not yet discussed, includes all the structural tools which are visible on the speech-channel and are used to mark linguistic meanings in natural languages. Besides phonology, these include morphological markers, linear order, adjacency, and so on. What is missing from this picture is abstract, autonomous, invisible syntax. As we have seen before, we have good reasons to assume that such syntactic representations are no longer needed. Moreover, the above schematic characterisation allows for both linguistic variability and universality—on all fronts. Different languages may, in principle, occupy the level of linguistic meaning with different semantic categories and different categorical combinations (this is the problem of linguistic relativity: see Gumperz and Levinson 1996), but some major categories—like event structure and animacy—seem to be universal. Different languages use different subsets of linguistic form to mark semantic categories. Some markers, however, are universal. Finally, different languages may map different semantic categories onto different markers; this is an essential property of

the system. Thematic roles, for example, may be marked by linear order in some languages and by morphological case markers in others.

This novel characterisation of language immediately reframes the question of language evolution. It is now neither the question of the evolution of a formal, non-functional system, nor of the evolution of a general-purpose communication system. It is the question of the evolution of a specific communication system dedicated to the communication of a constrained set of meanings by means of sound concatenation. In cognitive terms, it is the question of the evolution of a mapping system: of the gradual expansion and sophistication of the representational levels of linguistic meaning and linguistic form, and their transparent mapping onto each other. The answer to this novel evolutionary question lies in the interaction between two different evolutionary processes: cultural evolution and genetic evolution.

4. Cultural evolution and genetic assimilation

Genetic evolution involves a change in the nature and frequency of genes in a population. Similarly, cultural evolution involves a change in the nature and frequency of socially learned and transmitted behaviours in a population. Both cultural and genetic evolution clearly played important roles in the evolution of hominids. Early hominids must have had cultural traditions, which are the consequence of cultural evolution, and modern humans and early hominids are certainly genetically different.

Traditions are ubiquitous in higher animals and encompass every aspect of their lives: modes of foraging, mate selection, avoiding predators, criteria for choosing a habitat, practices of parental care, and so on (Avital and Jablonka, 2000). Traditions are particularly well studied in primates. Thirty nine different cultural traditions were recently described in seven populations of the common chimpanzee, and five of these traditions have something to do with communicative-social functions (Whiten et al. 1999). There is much more to be learned here, since researchers have only recently started to study animals' traditions systematically and comprehensively. That some form of culture and much cultural variation (some of which was associated with communication) existed in hominids, can be taken for granted. In hominids, cultural evolution is often cumulative and often leads to the gradual sophistication of a cultural practice. A good example is socially learned and transmitted improvements in tool-making techniques, which could have gradually led to the more elaborate fashioning of stone-tools.

That said, genetic evolution was certainly also involved in language evolution. After all, the linguistic differences between humans and chimpanzees seem to be at least partially genetically based. The interesting question concerns the nature of this genetic difference, and the relationship between linguistic cultural evolution and the evolution of the genetic difference. Can the cultural evolution of languages be related to the genetic evolution of hominids? We claim that it can and that this interaction is particularly important for understanding language evolution. We suggest that the evolutionary process involved the co-evolution of genes and culture through a dynamic process of genetic assimilation. What was genetically assimilated was the increasing capacity to acquire language—a process that resulted in a cognition biased towards the acquisition of language.

4.1. Genetic assimilation and simple "instinct" evolution

Genetic assimilation is the transition, through Darwinian selection, from an acquired (learned or induced) response to a more genetically fixed or "instinctive" response. Also known as the Baldwin effect (Simpson 1953), this process involves selection for the ability to respond rapidly and efficiently to the new stimulus. When individuals face a new environmental challenge, they usually first adapt to it by learning. If the selective pressure is ongoing (learning takes time and is costly), there will be selection for the best and fastest learners. This may culminate in a population of individuals who learn very quickly and even of individuals for whom a single exposure to the stimulus is sufficient to elicit the adaptive response. The learned response becomes an "instinct".

Conrad H. Waddington, the British geneticist and embryologist who coined the term "genetic assimilation", showed experimentally how an acquired trait could be transformed, through Darwinian selection, into a trait that is partially or completely independent of environmental induction. He focused on induced physiological responses in fruit flies.

Fruit flies normally have two wings. Waddington induced the development of four wings by treating fertilised fly eggs with ether. Not all treated flies developed four wings. But Waddington selected those that did, bred from them, treated their fertilised eggs with ether, and allowed them to become adults. He then selected again and repeated the whole procedure. In each generation he kept some eggs unexposed to ether and checked whether they developed into four-winged flies. In the first 20 generations, none did.

But after 20 generations of systematic selection, a few flies with four wings started appearing in the selected line—even without ether treatment. The trait whose development was at first dependent on external induction by ether became genetically fixed and independent of the ether treatment (Waddington 1953).

How did this work? The ability to develop four wings as a result of ether induction has a genetic basis. This genetic basis was exposed by the ether treatment and was then selected. By gradually selecting the gene combinations that produced an ether-induced, four-winged phenotype most effectively, a threshold was eventually crossed, and a particular combination of genes that enabled the development of four wings now appeared without the external inducement.

R. F. Ewer and John B. S. Haldane used Waddington's analysis to explain the evolution of behavioural instincts (Ewer 1956, Haldane 1959). Haldane suggested that the innate, excited response of sheepdog puppies to the smell of sheep may have evolved through the genetic assimilation of an initially learned response. For hundreds of years, shepherds selected for dogs that performed their task effectively. Many sheepdog properties were selected for, including the ability to react excitedly to the smell of sheep. The combinations of genes that contributed to this excited response were gradually selected. A response that was initially learned by reward and punishment became an almost entirely automatic response.

Another example, this time involving natural and not artificial selection, may be the innate avoidance response of spotted hyenas to the smell of lions, and the avoidance response of many small mammal and bird species to hissing, snake-like noises (Kruuk 1972, Edmunds 1974). Individuals that learn quickly and remember the sound or smell that should be avoided have a better chance of surviving, and the genetic constitution of these fast learners will be passed on to the next generation. After generations of selection for fast association between a certain sense impression and the danger—and for evasive action—the avoidance response will become "innate". Its expression will depend on very few or even a single exposure to the danger stimulus. It is important to note that although in these extreme cases a particular response may become independent of learning, it is in fact the *ability to learn* that was selected. Learning became increasingly efficient and rapid until it was ultimately internalised.

In some cases, there is an additional factor: the learned response changes the individuals' environment. For example, assume that instead of merely learning to detect a predator's smell or sounds more quickly some orga-

nisms learn to hide by digging burrows. Here, there will be selection for learning to dig more effectively. However, because burrows amount to a new environment, individuals are now also selected for their ability to live in them. Those individuals that are both efficient diggers and efficient burrow-dwellers are positively selected. In this case, the pressure to avoid predators led to what is called *niche construction*. Organisms actually construct the environment in which they and their offspring are selected (Lewontin 1978, Odling-Smee et al. 1996).

The process of genetic assimilation, which involves the exposure of new genotypes to selection and often also the construction of the selective environment, explains how effectively blind genetic variations can simulate an acquired response within a brief period of evolutionary time. For the process to work we must assume that: a) populations have abundant genetic variation that is relevant to individuals' ability to respond to stimuli; b) different sets of genes become selectively relevant under new circumstances; c) phenotypically visible genetic variation can be recruited and organised into new adaptive genotypes via sexual reshuffling and selection; and d) selection for the adaptive genotypes (genotypes that enable more adaptive responses) is maintained for several generations. What we know of the nervous system and of the abundance of genetic variation in animals not only allows us to make these assumptions, but also suggests that such processes must have been common during evolutionary history.

4.2. *Stretch-assimilate: the sophistication of behaviour*

Avital and Jablonka (2000) discuss an important consequence of genetic assimilation (which enables the lengthening of a sequence of learned behaviours) by making a portion of it partially innate. Imagine, for example, a bird capable of reliably learning a sequence of four consecutive acts that culminate in a simple nest. Additional learning is difficult. Assume, for the sake of simplicity, that there is a constraint on the learning capacity of this species of bird. Improved learning ability is unlikely to evolve (perhaps because a large brain requires more energy, or there may be some developmental constraint on brain growth). However, if there is consistent selection for the efficient and reliable performance of the nest-building behavioural sequence so that one of the steps becomes genetically assimilated: it becomes innate. The bird now needs to learn only three steps and will construct its simple nest more efficiently. Yet part of its unchanged learning

capacity is now freed up. If selection for building good nests continues, the bird can now learn an additional nest-improving skill. For example, it may learn to use plant strips to secure the nest, thus enhancing the nest's stability in windy conditions. There are now five consecutive acts, one of which is innate. If building nests rapidly and efficiently continues to be advantageous, another previously learned act can become assimilated and another newly learned one can be added, extending the behavioural sequence by a further step. It is thus possible to gradually lengthen the sequence of acts without changing the capacity to learn. Genetically assimilating previously learned behaviours frees the individual to learn additional acts without extending the limits set by its learning capacity.

It is not necessary to assume that any of these acts is completely assimilated. It is sufficient to assume that the number of trials required for the effective performance of the behavioural act is significantly reduced. This is, in fact, the most likely effect of genetic assimilation (Hinton and Nowlan 1987, Behera and Nanjundiah 1995). Reduction in the amount of time and energy spent learning one activity allows for time and energy to be spent on another activity. The stretch-assimilate process may underlie the evolution of complex behavioural sequences that comprise both learned and partially or fully innate components. It could explain the evolution of complex behaviours such as nest building, bird singing—and human linguistic communication.

5. The linguistic spiral

We are now in a position to propose an answer to our evolutionary question, the question of the gradual and directional evolution of the linguistic mapping system. Think about the process of language evolution as comprising an arbitrarily long number of stages, and concentrate on two early, consecutive stages: stage N and stage N+1. Assume that at stage N, a community of hominids shares the necessary precursors for linguistic communication: a certain minimal theory of the mind, a certain level of conceptualisation, a motivation for information sharing, a certain implicit understanding of social relations and hierarchies, and so on. Assume also that this community uses a preliminary, culturally transmitted system of linguistic or quasilinguistic communication: a system that maps a set of meanings onto a set of phonetic markers. The meanings and their markers need not be recognisable in present-day languages. They may include, for example, ritualised calls,

social-emotional vocalisations, and so forth. Whatever stage N's specifics are, one thing is certain. The system's expressive envelope is much more limited than the expressive envelopes of our languages. But, more importantly, the system's expressive envelope is much more limited than the individuals' conceptual envelopes. Although their conceptual envelope is much narrower than ours, they still—like chimpanzees—can think and feel much more than they can say. Assume further that the individuals in this community use their quasilinguistic system comfortably and naturally and that their children comfortably acquire it. Finally, assume that this community has a particular genetic constitution that allows them to acquire and use the system—with the necessary amount of variability: some individuals are better at acquiring and using the system than others.

Now assume that on the way to stage N+1 at least two changes occur. First, the communication system developed and its expressive envelope was expanded and became significantly more sophisticated. Second, the genetic constitution of the individuals in the community changes so that they comfortably acquire and use the more sophisticated system.

Our explanation starts with cultural evolution. Assume that at different times throughout stage N, individuals or groups of individuals make linguistic innovations. The driving force behind such innovations must have been associated with a growing pressure for better communication within the group. It may have been related to increases in group size (Dunbar 1996), significant changes in ecological conditions, changes in tool usage, changes in the need for social cooperation, or changes in interactional patterns between different hominid populations. The range of linguistic innovations during the evolutionary process must have been wide, and most of them must have occurred repeatedly: new lexical items for specific referential meanings; new abstract markers for existing and novel conceptual distinctions; new epistemic markers and speech-act markers; new pragmatic conventions for linguistic communication; more sophisticated morphological and phonological structures; a more sophisticated usage of linear order and adjacency to mark meanings; and so on. We will not speculate on the order of these innovations (but see Jackendoff 1999 for some interesting hypotheses). Let us assume that some of them occurred at stage N. Like any other type of cultural innovation (a new tool-making technique, say), these linguistic innovations may have been accidental or the result of conscious efforts by clever, inquisitive, or just lucky individuals who happened to be in the right social context at the right time. Many of these individuals were probably juveniles, who are particularly inclined to explore and innovate. In

many cases the innovation—like the sign language developed by Nicaraguan children—may have been the result of group effort. The important point is that we do not need to invoke a genetic explanation for any of these cultural innovations. The cultural-linguistic innovations of stage N were within the genetically based capacity of these talented—or serendipitous—linguistic innovators.

Now, although only a small minority in any community is capable of real innovation, a much larger group of individuals is capable of learning to understand and use the innovation once it exists. This also does not require a genetic explanation. Research on chimpanzees and on human children has demonstrated the considerable plasticity of primate cognition. Because comprehension typically precedes production, and because social learning takes advantage of the system's considerable plasticity, the innovator has a good chance of being understood (for example, most of us can understand Newton's theories, though few of us have Newton's genius). This is especially true of the innovator's family members and close friends. Different individuals, however, will differ in their ability to understand and use the innovation. At least some of this variability will stem from variability in their genetic makeup. Some will grasp the innovation better than their peers, some will learn to use it themselves, some will manage to passively comprehend it, and others won't understand it at all.

What happens to a cultural innovation once it is learned by a few members of the community? Its fate depends to a significant degree on its propagation and dissemination across the population. This is the real bottleneck. In the innovator's own generation, the propagation of the new linguistic tool may be unstable and uncertain. Many innovations, including some very adaptive ones, will probably disappear at this stage because their significance can sometimes be fully appreciated only when they are used by a significantly large and cohesive group of communicators (there is—up to a point at least—positive frequency-dependent selection). We may assume that innovations have a better chance of establishing themselves after the first learners transmit them to their offspring. Cultural evolution in primates and other animals (like the spread of food washing by Japanese macaques on Koshima island) demonstrates that children play an important role in establishing cultural traditions. We may also assume that for a long time after their invention, innovations undergo further correlated cultural evolution. They may be improved in all sorts of ways as well as become conventionalised and streamlined. The semantic categories, marked by the innovations, will go through a gradual process of differentiation, amplifica-

tion, and sharpening. They will gradually acquire their discrete character. Moreover, categories will gradually dissociate themselves from emotional connotations, from specific prototypical contexts, and so on, because of their application to a wide variety of social contexts.

An innovation's chances of establishing itself do not, however, depend solely on its pattern of propagation. They are crucially dependent on its adaptive value as a tool of social communication. For a linguistic innovation to survive, its usage should benefit those who adopt it. In general, a linguistic innovation's adaptive value is a direct function of its *information potential* and an inverse function of its *processing effort* (Sperber and Wilson 1985). An innovation carries high information potential to the extent that it allows for the transfer of more information which is relevant to the community, to the extent that it adds relevant elements to the system's expressive envelope, and to the extent that it enables more precise production and interpretation. An innovation requires low processing effort by being relatively easy to acquire and use in contexts of social communication.

The idea that the survival of a linguistic innovation depends on its adaptive value requires a few additional remarks. First, an innovation's information potential is not perforce related to the practical considerations usually discussed in the literature, such as cooperating efficiently in hunting or fighting as well as sharing information about the natural environment. Although these considerations are important, it seems to us that an innovation's information potential is also a social issue. It is related to the sharing of social information (social relations, social events, and social hierarchies), to the sharing of social narratives and myths, and to the construction of social epistemology (see Knight 1998 and Heeschen 2000). As we know from the recent literature on the relativity problem (see Gumperz and Levinson 1996), linguistic markers help determine the extent to which a community of speakers isolates and foregrounds some aspects of its environment, and thus establishes its epistemological perspective. This epistemological establishment, in turn, plays an important role in linguistically based social identity, which further strengthens the innovation's adaptive value. The categories we discussed in sections 2 and 3—event structure, epistemic status, and animacy—seem to be especially relevant in this respect (see Dunbar 1996). Second, for a linguistic innovation to survive and propagate it must be adaptive for a sufficient amount of time and preferably in a wide variety of changing circumstances. This is especially true for categorical markers. New lexical items may come and go, but categorical markers that manage to survive probably remain adaptive throughout numerous social changes over

a long time span. Third, some types of information—emotional messages, manual instructions—are effectively communicated non-linguistically through body language, facial expressions, mime, song, and dance. Linguistic innovations directed at these types of information may not survive (or may not be invented in the first place) because other means of communication render them unnecessary. Division of labour among different communication systems may thus play a significant role in the cultural evolution of language's highly constrained expressive envelope. Fourth, for linguistic innovations to survive they have to meet the conditions set by *system constraints*. These fall into at least two types. First, linguistic innovations have to comply with psychological constraints. Those innovations that correspond to pre-existing cognitive or developmental biases will probably be selected, since they are the easiest to learn, remember, and transmit (see Sperber 1996). Second, as the linguistic system evolves, it sets its own constraints on new innovations. They have to comply with the already established system. This means that, at least after a certain point in the evolution of language, the system itself dictates the direction of its own future evolution.

Let us assume, on the basis of the above considerations, that some of the adaptive linguistic innovations of stage N manage to spread and establish themselves in the community. They endure because they are both dependent on, and constitutive of, the social structure and because social traditions are by their very nature self-perpetuating. This cultural change enhances the capacity of the members of the community to communicate. Yet it also raises the demands for social learning imposed on individuals in the community. They not only have to acquire the new innovations in order to participate in social communication, they also have to learn to look at the world in new ways, direct their attention to new aspects of reality, process and remember new types of information, and so forth. In short, the linguistic innovations that establish themselves in the community change the social niche, and the inhabitants of this new niche have to adapt to it. In adapting themselves to the niche, the individuals will probably be able to count on their mind's built-in residual plasticity. Individuals and cohesive groups of individuals that make better use of the innovations for efficient communication (for whatever cultural or social reason) will benefit. They will probably be reproductively more successful than others—and more likely to thrive.

Gradually, however, the increasing cognitive demands set by the evolving linguistic niche will start to expose hidden genetic variation.

Individuals will find the accumulating linguistic innovations more and more demanding. Residual plasticity will gradually be stretched. After a long period of consistent, directional cultural selection, genetic assimilation will occur. Some individuals will drop out of the race; others will survive. The population will not become genetically homogenous, but the frequencies of the gene combinations that contribute to easier language acquisition and use will increase via the continuous processes of genetic reshuffling and selection. Note that genotypes may contribute to acquisition and use in a wider variety of ways. Eventually, at stage N+1, we will find a community whose general genetic makeup enables individuals to use the system comfortably and children to acquire it comfortably. Now the whole process can start over again. Assimilation frees individuals to make further use of their cognitive plasticity.

What might be genetically assimilated in our transition from stage N to stage N+1? All the relevant aspects of general cognition were assimilated, to a certain degree at least, according to suggestions by Lieberman (1991), Donald (1991), Jablonka and Rechav (1996), Deacon (1997), and others. Individuals at stage N+1 probably were more intelligent, had better memories, had better voluntary control of their sound production mechanisms, and were smarter social agents. We believe, however, that individuals at stage N+1 had a cognitive constitution that was marginally more biased toward acquiring and using language than the cognitive constitution of individuals at stage N. In other words, after the long period of cultural evolution in which the community became increasingly dependent on linguistic communication and individuals' survival depended increasingly on their linguistic performance, the process of genetic assimilation must have targeted the cognitive capacities most useful for this specific type of behaviour. Some examples are the capacity to recognise discrete conceptual categories, to rapidly process the speech channel, to recognise linguistic-communicative intent, and to expand lexical memory. These are language-specific and must have been targeted by linguistically driven genetic assimilation.

Moreover, the genetic assimilation of these capacities was, for two complementary reasons, probably partial rather than complete. First, genetic assimilation of specific linguistic behaviours cannot lead to a completely innate response because of the variability of the situations in which the linguistic behaviours are adaptive. This is particularly true in our case, where cultural evolution is an integral part of the process. The adaptive value of the linguistic behaviours has to track a constantly changing cultural environment. Second, partial assimilation makes learning easier and faster and reduces the

selection pressure for additional assimilation. All this has straightforward linguistic implications. Very specific innovations, such as specific words or specific morphological markers, are not assimilated. They are too variable and context-dependent and change too rapidly throughout cultural evolution. But certain semantic categories can be partially assimilated because they remain adaptive for a long time across many cultural environments. But even these categories will not be *completely* assimilated—they will not create a semantic "instinct"—because cultural change puts a high premium on epistemic flexibility.

The process of linguistically based genetic assimilation may be related to the general evolution of human culture and human conceptualisation. As we have indicated, genetic assimilation also targets general intelligence. We know, after all, that there was no significant constraint on the evolution of the hominids' general intelligence. Hominid brains doubled in size in 2.5 million years. As the process of cultural and linguistic evolution constantly leads to an extension of the environment as perceived by the community, individuals are constantly faced with more information about the world. They can learn more about more aspects of the world because they can think and communicate more effectively. This creates a process of positive feedback. The more individuals learn about the world, the more they can communicate; and the more they communicate, the more they can learn. On the one hand, individuals and the whole community are now in a position to evolve their conceptual structures with the aid of a more complex communication tool: language. On the other hand, the evolution of conceptual structures, of general cognitive tools for learning, and of remembering aids the concomitant evolution of the linguistic system. The linguistic system thus spirals together with the conceptual system (and with the motor control system, which we have not discussed in this paper). This wider spiral also includes a wide variety of non-linguistic, culturally based evolutionary processes that interact with each other in complex ways. The process resulted both in the expansion of hominids' conceptual capacities and in the construction and expansion of their linguistic expressive envelope.

This conception of the process renders theoretically unnecessary the traditional distinction between the syntactic nature of present-day languages, and the supposedly presyntactic nature of so-called protolanguage (see Bickerton 1984 and 1990). Because we conceive of the gradual increase in grammatical complexity as a reflection of the gradual increase in the complexity of the expressive envelope, and because the capacity to acquire this complexity was gradually and partially assimilated on the basis of the

system's cultural evolution, we conceive of the entire evolutionary process as a gradual and continuous one. The simple and crude structural properties of the mapping system in its first evolutionary stages were a reflection of the system's expressive envelope, just as the complex and sophisticated properties of present-day linguistic mapping systems reflect their elaborated expressive envelopes. There were no island constraints, for example, prior to the point at which the interaction between event structure, epistemic status, and the structure of speech acts became a constitutive part of linguistic semantics.

6. Conclusion

The framework developed in this paper takes us a long way toward resolving the three constitutive paradoxes presented in the introduction. We started out by characterising language as a cognitive system which is both functional and unique: a transparent mapping system dedicated to the expression of a constrained subset of meanings by means of sound concatenation. As we claimed, recent advances in linguistic research support this conception—in direct opposition to Chomsky's traditional hypothesis of the autonomy of syntax. We then characterised the evolution of this system—and the evolution of its social users—as the interaction between cultural and genetic evolution. We discussed the evolution of the linguistic system in cultural terms as the social process of innovation, production, comprehension, transmission, and propagation of linguistic conventions, in which a community isolates and foregrounds certain aspects of its epistemology and develops social agreement about the means of their expression. This process results in a functional (rather than formal) and highly constrained (rather than general-purpose) communication system because it is founded on a *selected* subset of semantic categories. At each stage of this long and continuous process, the system's expressive envelope expands, and the structural means of expression become more sophisticated. We then discussed the genetic evolution of this system's users in terms of partial genetic assimilation. Partial assimilation resulted in linguistically biased cognition, that enabled easier and more effective language acquisition and use. This conception allows for the resolution of the paradox of domain-specificity. Partial genetic assimilation does not copy linguistic specificities into brain structures, and it does not result in genes for linguistic specificities. Instead, it constructs a genetic make-up that supports the development of a cognition

biased towards acquiring and using linguistic specificities (though a significant amount of learning remains mandatory). The question of innateness thus becomes the question of how much, when, and what type of learning are necessary at each stage of the evolutionary spiral. In line with the approach suggested by Elman et al. (1996), our model avoids the nature-nurture dichotomy.

Our view of the continuous interaction between cultural and genetic evolution is not only consistent with the dynamic nature of languages and with the attested variability among different languages, it actually considers these properties fundamental to the evolutionary process. Languages are constantly changing in their social contexts, and a certain degree of linguistic universality is accompanied by a certain degree of linguistic variability. This is exactly what one would expect as the result of this process.

Finally, the framework developed in this paper reconciles the two major approaches to language evolution, one focusing on the evolution of language as a system of social communication, and the other focusing on the evolution of the structurally unique properties of language. As the exchange between Bickerton (1996) and Dunbar (1998) makes clear, scholars adhering to the two approaches have instituted an artificial division of labour. Socially oriented researchers have concentrated on the adaptive value of language as a communication system and largely ignored its formal properties. Structurally oriented scholars have focused on formal specificities and largely ignored social communication. According to our theory, language's formal properties are a reflection of meaning relations. These, in turn, have been selected throughout the evolution of language on the basis of their adaptive value in terms of social communication. The formal question and the social question are one and the same.