

Language as a Communication Technology: A Proposal for a New General Linguistic Theory

The Structure of the Argument

Daniel Dor

The theory of Language as a Communication Technology (LCT) is a large-scale attempt to develop a unified model of language – a model that captures the essence of language as a general phenomenon, and provides the necessary common ground for the description, analysis and explanation of its different facets. This article presents a skeletal description of the entire line of argumentation of the theory. It begins with a new characterization of language as a communication technology of a very specific functional type, a collectively-constructed technology that allows for communication across the *experiential gaps* between its users. It then shows (very informally) how this characterization allows for the unified treatment of major theoretical issues from the realms of functional and structural analysis, lexical semantics, language production, pragmatic interpretation, the philosophy of language, sociolinguistics and the sociology of language, language and power, the relationship between language and thought, the dialectics of the universality of language and the variability of languages, and the questions of language acquisition and the human capacity for language. The article then shows how the theory meets the single most important explanatory challenge that any general theory of language should be able to meet: An explicit, detailed and multi-dimensional speculative narration of the evolution of language – a narration that meets the requirements of current evolutionary theory.

1. Language as a Communication Technology

1.1 To understand language as a unified phenomenon, we have to abandon the cognitive perspective, and begin to conceptualize about language as a collectively-developed communication technology.

For a very long time, fifty years by now, we have been slowly moving away from a certain general picture of human language. Different parts of the picture had gradually emerged from the 17th century onwards, but the picture as a whole, in all its rationalist glory, only appeared in the middle of the 20th century: Language as a universal, static, autonomous, formal, truth-conditional, innately-given, uniquely-human computational competence. It was an exciting picture. But then, almost immediately, we began to move away. Meaning was brought back into the picture with the development of functional and cognitive grammars (especially of the constructionist type); with the advances made in the understanding of the interface between syntax and semantics; and with the emergence of pragmatic theories of linguistic interpretation. The picture of language that we now have connects structure and meaning in complex systems of interrelations. All of the above, especially the advances in Pragmatics, have dragged our notions of meaning from the realm of truth-conditional semantics into the worlds of general cognition and interactive communication. The sociolinguistic quest for socially-determined variation; the anthropological-linguistic search for the relationships between language and culture; the growing understanding of the intersubjective nature of language (based, among others, on Vygotsky and Wittgenstein); the advances made in the theory of grammaticalization; and the development of usage-based grammars (again, especially of the constructionist type) – all these have contributed to the current positioning of language somewhere between human cognition and human society. All of the above, together with what we now know from Conversational Analysis, has challenged the distinction between the original notion of competence and the attested patterns of linguistic performance. All this, together with the results accumulated in Linguistic Typology - and the simple fact that more and more languages came to be deeply researched - has changed our views on the universality of language: What we see today is a web of similarities and differences between languages - many restricted, implicational

universals; not too many absolute ones: Nothing, definitely, that looks like a set of properties worthy of the name Universal Grammar. The enormous advances we have made in the understanding of the cognitive capacities involved in language acquisition and actual linguistic processing now position language somewhere between its original autonomy and the realm of general human cognition – between innate knowledge of language as such and the general human capacity for learning, especially for social learning. All this, taken together, has renewed our interest in the dynamics of language change – and has positioned language, again, somewhere in between the synchronic and the diachronic. A rich and lively discourse, unimaginable fifty year ago, now attempts to tackle the most difficult question of all: The question of the evolution of language.

We have moved quite a long way from the original picture; we know much more than we ever did. And yet, we do not *understand* language any more than we did in the beginning. Why? Because of two reasons: First, Linguistics has turned in the last fifty years into an extremely fragmented field, in which different explanatory apparatuses, incompatible with each other, serve the theoretical needs of highly specialized sub-domains. We have pieces of the puzzle, but they do not assemble into a theory that makes sense of the entire phenomenon of language. Because of that, and that's the second reason, when we look at our fragmented pieces of language, situated as they are away from the original picture, we still look at them *from the point of view* of that very same picture: We have found new answers, but they are still answers to the same *questions*. Our general theories of language are still about human cognition; they are still answers to the question of knowledge. We now have very different grammars, but we still think about them in the same way, as computational characterizations of levels of cognitive representation. We understand pragmatics much more than ever, but we still think about it as something that follows truth-conditional semantics. We struggle to find a place for variability within a framework that is still universalistic. We look at language as a social entity from the point of view of social-*psychology*.

The theory of Language as a Communication Technology (LCT) is a large-scale attempt to propose a new picture of human language, a new point of view from which to look at the entire phenomenon. The basic idea is quite simple: Language is a social communication technology, not a cognitive capacity; we should thus move beyond the cognitive perspective, and adopt a social-technological approach to the study of language.

This does not imply that the cognitive questions are unimportant. On the contrary: Language is a technology that requires much of its users, a whole array of cognitive capacities, and these deserve our full attention. But in order to understand the capacities (in acquisition, in actual communication) - we must first of all understand the technology itself. We have to come to terms (come *back* to terms) with the fact that the technology as such resides at the social level, a level of organization and complexity that transcends the individual mind - *between* the speakers, not *in* them.¹

In terms of the *computer metaphor*: The hypothesis that language is a cognitive capacity, a piece of software *inside* the mind, has achieved an enormous lot. We know enough today to finally abandon the computer metaphor – and replace it with the updated metaphor of the *web*. Language is indeed a piece of software - a *communication* software. It resides on the net. It is a property of the *community*, of the *social network*, the product of a *collective* process of invention and development. Human individuals are *end-users*: They download the software into their mind-brains, and it allows them to communicate with one another (to the extent that they have the same version of the software). What we understand today about the processes and capacities involved in the downloading and usage of the software by the end-users should teach us exactly this: That the essence of the software itself does not lie *there*. The essence of the software lies in the fact that it resides between the end-users; that it facilitates something that is not within the individual capacity of any of them. Language simply cannot be explained in terms of the cognitive dynamics taking place within the individual mind-brain – for the exact same reason that social networks on the net cannot be explained on the basis of whatever is happening in the personal computers of the networks' members. The baroquely-complex theoretical models of language that have emerged in the last half century are a direct result of the

¹ The conviction that language is a social entity (a social *institution*), played a foundational role in much of the linguistic theorizing of the first half of the 20th century – in de Saussure (1916, 1966), Edward Sapir (1921), Antoine Meillet (1921), Sir Alan Gardiner (1933), and others. It has also informed much of the semiotic literature on language (see, for example, Itkonen (2003, 2008) and Zlatev (2009)). Most of the other disciplines interested in language - Communication Studies, Sociology, Anthropology, Literature, Critical Studies and so on – have always worked with a general, intuitive and very often implicit view of language as a social entity. The theory of LCT corresponds with all these in a web of complex ways that I will not describe here. It parts way with all of them, however, in the foundational sense that it begins with an investigation of language as a *technology*.

attempt to do just this – to achieve an impossible goal. Language is a much simpler, much more reasonable entity, but to see that, we have to look at it at the right ontological level, where it belongs.

1.2 The conception of language as a communication technology immediately and necessarily positions a single question at the center of our theoretical quest: The question of the *functional specificity* of language.

The question itself thus rejects both the Generative conception of language as a non-functional, formal entity, and the opposite conception of language, that of Functional and Cognitive Linguistics, as a general-purpose functional system. On the one hand, the essence of a functioning technology simply cannot lie in the non-functional specificity of its architecture: Non-functionality may explain *useless* technology, but it is not an explanatory option with such an enormously successful technology as language. On the other hand, no technology is ever *general-purpose*. Had there been such a thing, we would never need to invent anything else. Language simply cannot be a general-purpose system of communication. Just like any other technology, its essence must lie in the fact that its function is specific. This is where we should start.

It is important to see that this formulation of the question actually accepts the *positively*-formulated *observations* of both camps, and rejects their *negatively*-formulated *explanations*: No, language is not non-functional or non-specific. As a technology, it *cannot* be any of those. Yes, language is functional and specific. As a technology, it *must* be both.²

² A systematic comparison of LCT with the many variations of Generative and Functionalist-Cognitivist theories that inform the current discourse in Linguistics is beyond the scope of this programmatic article. where theoretical contextualization requires, I will highlight points of similarity and difference with the different theories. At this stage, I will only say this: Functionalism and Cognitivism have emerged as a series of attempts to provide new answers to Chomsky's original question concerning the nature of the universal human capacity for language. They have thus played against Chomsky, so to speak, in his own field. LCT attempts to move away from this field - both from the meta-theoretical presuppositions shared by both sides (such as the presupposition of cognitive universality, specific to language or general-cognitive), and from the often artificial distinctions that it came to reify (such as the distinction between functionality and specificity). This is why the theory can build on insights accumulated by both sides: In

1.3 To understand the functional specificity of language, we have to re-position language in the overall context within which it operates – the context of experience. Our individual experiences are always private, different from the experiences of the others, and inaccessible to them. We are forever separated from each other by *experiential gaps*.

The process of experiencing is the foundation of our lives as mental creatures: Like all the other animals with a nervous system, we live *in* experience. Experiencing includes everything that we call feeling, thinking, understanding, seeing, hearing, imagining, wishing, and so on - but also, importantly, everything that happens in our nervous system when we *do*, move, touch, react, try, succeed and fail. As our experiences accumulate in our minds (leaving their traces in our nervous system), we detect similarities and analogies between them, and construct generalizations – experiential generalizations, always analogue, holistic, fuzzy and context-dependent – which then color, shape and sometimes determine the way we further experience. This is how we *learn*. All this does not deny the possibility that our nervous systems might be innately biased, in different ways, towards certain ways of experiencing and accumulating experiences. What it does deny, and very strongly so, is the idea that our general cognition can be described – let alone explained – in terms of formal computational operations on abstract symbol-like entities. It thus denies the existence of a Fodorian *language of thought* and the existence of *conceptual structure* as a level of cognitive representation. The theoretical move from abstract representationalism to experiencing (very often under the title of Embodied Cognition) has been inspired, among others, by the works of William James, John Dewey and Martin Heidegger – and has been substantially argued for in the last two decades on philosophical, cognitive, computational, experimental, phenomenological, biological and evolutionary grounds: I will not attempt to review the literature here.³ The important

many ways, it is deeply inspired by Functionalism and Cognitivism. In other ways, it ends up siding with the Generative view.

³ A few references: Shapiro (2010) and M. L. Anderson (2003) are good introductions to the major issues discussed in the context of embodied cognition (or embodied cognitive science). Prinz and Barslau (2000) is an attempt to re-think the notion of representation within the framework of embodied cognition. Zlatev (2008) is a phenomenologically-based discussion of the relationship between language and experience. Ginzburg and Jablonka (2007a,b) deal with the evolutionary emergence of the process of experiencing in lower organisms.

point for our purposes is only this: The assumption of abstract representationalism (in and of itself a part of the rationalist picture we have been moving away from) has kept us away from the essence of language, because it created an illusion of *sameness* between the way we know and the way we speak, between the categorizations of our experiences and the categorizations of our languages. This illusion of sameness masks the foundational problem that language, as a communication technology, has to overcome, and thus misses out on its functional specificity.

Experiencing is a complex process. There is much that we do not yet understand about it. For our purposes, however, we only need to highlight three trivial properties of experiencing - so trivial, as a matter of fact, that they are very often overlooked in the literature. First, experiencing is *private*. We experience on our own, within ourselves, even when we experience together. (This does not deny that important things happen at the intersubjective level when we experience together, but even then, we still experience intersubjectivity from our own private perspective.) Second, the private essence of experience implies that every human individual, each of us, experiences the world in *different* ways. Every individual looks at the world from his or her own egocentric perspective; every one carries a totally different baggage of memories, different private histories of interaction with different worlds. Every individual comes to rely on different strategies for understanding the world; each is by nature (and by instruction) more deeply attuned to certain aspects of the world than to others. Different individuals have different attention spans, different perceptual capacities, a different talent for detail, a different eye for distinction, a different capacity and interest in generalization, different conceptions of the order of things, of what belongs with what, what is similar to what, what is the cause or effect of what. They have different ideas of what is beautiful, frightening, useful, important or interesting. They are different in gender, age, social status, physical strength, emotional development, curiosity, patience and anxiety. Every human individual lives in a radically different experiential world (which, again, does not deny that there are similarities.) Third, no human individual has access to the experiential processes taking place inside the others' minds. We are forever separated from each other by *experiential gaps*.⁴

⁴ There is a certain theoretical complexity here that requires explication. The discourse on human experiencing overlooks these three foundational properties of experience - for a very good theoretical

In order to understand language, we have to abandon both the Kantian dictum, the foundational presupposition of the cognitive sciences, that all human experiences comply with a universal interpretative scheme, and the neo-Kantian conviction, the foundational presupposition of most of the social sciences, that the members of every culture and sub-culture experience the world in the same ways. We have to begin with the acknowledgment that each human individual lives in a private, experiential world which is different from that of the others, and is inaccessible to them. This is a foundational fact about our cognitive nature, and it is the foundational *obstacle* to communication which language, as a social invention, set out to circumvent. Human cognition participates in the story of language not as the origin, but as the *original problem* that had to be solved by a social technology.

1.4 All the active systems of communication used by humans and other animals, apart from language, are systems of *experiential* communication. They allow for the *sharing* of experience. They only become functional where the experiential gap between the communicators is reduced, if only momentarily, when they experience together.

reason. Much of the discourse has emerged as a counter-reaction to the solipsistic view of human experiences based on Cartesian philosophy, and has thus systematically highlighted the intersubjective nature of human experiencing – the primacy of the interpersonal over the intrapersonal. Much of the discourse thus concentrates on the human facts of experience sharing, social learning, mimesis and imitation, mind-reading, joint attention and empathy, and on the consequent human capacity for co-operation, collective thought and cultural production. All these are undisputable and extremely important. I agree with Herb Clark (1996), Michael Tomasello (1999, 2008, 2009), Merlin Donald (1991), Jordan Zlatev (2008) and many others that language and intersubjectivity are very closely interrelated, and that early humans must have gone through an entire stage of evolutionary development, *before* language, that resulted in the emergence of a co-operative, mimetic, intersubjective species. I do not agree with the implicit assumption that the evolution of human experiencing into the intersubjective dimension somehow erased whatever there was there before. The unparalleled human capacity for intersubjectivity evolved from an experiential foundation much closer to solipsism, and this foundation is still there, at the very core of our experiential lives: Intersubjectivity, like everything else, is always experienced privately. Failing to see this leads to another illusion of sameness between the experiences of different individuals, which, again, masks the challenge that language has to overcome. Language is indeed made possible by intersubjectivity, both ontogenetically and phylogenetically, but what it *does* goes beyond intersubjectivity. To understand what language does, we have to begin with human individuals as private experiencers.

All the active systems of communication used by humans and other biological species are dedicated to the communication of experiences. All of them are thus, by definition, attempts to overcome the obstacle of the experiential gap. This is what communication is all about. All the other systems apart from language, however, with the possible exception of the dancing rituals of the bees, share a functional strategy that is radically different from that of language. They communicate experiences in ways which are in and of themselves *experiential*. They are about the *sharing* of experience. There are two general types of experiential systems of communication: I will call them *presentational* and *re-presentational*.

The great majority of communication systems used by humans and other species are presentational. In systems of presentational communication, the communicating individual, the sender, behaves in a way that is, at the very same time, determined by his or her experience and made available for experiencing by the receiver. Presentational communication takes place *within* experience, in the here-and-now which is *experientially shared* by both sides to the communication event. Its very essence lies in the fact that the internal experience of the sender, whatever it is that gives rise to the entire event, is not hidden away, but is made *present*, revealed, mirrored in perceptible behavior, which immediately and directly becomes part of the experience of the receiver. A cry of pain, a frowning expression, a smile, a hug, a kick, a threatening posture, demonstrations of physical strength, different forms of mimesis (and teaching through mimesis), manual gestures, grunts and screams, music and dance – all these employ the strategy of presentation. There are, to be sure, many important differences between these systems: Some involve an explicit intention to communicate, others do not; some imply a certain level of dialogical coordination (sometimes turned into ritualized convention), others do not; some require a certain amount of learning (on both sides), some seem to be totally instinctive. The point that unites them, however, is that they all take place within the immediate context that is experientially shared by the sender and receiver at the moment of communication. They do not attempt to bridge the experiential gap between the two sides: They become functional only where the gap itself is reduced, if only momentarily, by the actual sharing of experience.

Many other systems of communication, used by humans, employ a different but related strategy: They are *re-presentational*. In such systems – drawings, paintings, maps,

musical recordings, photographs and movies – the sender's experiential intent is recorded, frozen in time, delivered from the here-and-now of the sender's experience into the receiver's here-and-now, where it is 'melted back', so to speak, brought back to life, in his or her mind. Much of the original complexity of the experiential intent is lost in the process. No system of re-presentational communication is ever capable of bringing the entire original experience into the present. But systems of re-presentation are nevertheless capable of restoring different types of experiences to different degrees. They follow the holistic and analogue logic of presentational communication, transferring (dragging, if you will) the experiential intent of the sender, across time and space, into the experiential world of the receiver. This is why the products of re-presentational communication acquire their significance from their similarity with their presentational counterparts, the corresponding directly-presented experiences which they are supposed to *stand for*. In semiotic terms, they are *iconic*: They provide their receivers with echoes, or silhouettes, of what they would perceive had the original experience been directly presented to them. Just like presentational systems, re-presentational systems do not attempt to bridge the experiential gap between the sender and the receiver: They use various technical means to allow for the sharing of experience across time and space, without actually bridging the experiential gap.

1.5 The functional specificity of language lies in the fact that it allows for communication across the experiential gaps between speaker and listener. I will call this *instructive communication*: The speaker provides the listener with a skeletal list of the basic co-ordinates of the experience – and these instruct the listener in the process of *imagining* the communicated experience. Instructive communication allows speakers to communicate experiences that cannot be directly experienced by their listeners.

In instructive communication, the very attempt to share an experience, whether directly or through mediation, is abandoned *in principle*: The sender does not try to make some of his or her experience perceptibly present to the receiver. Rather than doing that, the sender provides the receiver with a code, a plan, a skeletal list of the basic co-ordinates of the experience – which the receiver is then expected to use as *a scaffold for experiential imagination*. The set of co-ordinates does not stand for the original experience. It instructs the receiver in the process of the mental creation, through imagination, of a totally

independent experience – an experience which is supposed to reflect the original experience not because it is perceptually based on it, but because it is (approximately) of the same *type*.

In experiential communication (presentational and re-presentational), the sender communicates: '*This* is my experience'. In instructive communication, the sender communicates: 'My experience is of *this type* - try to *imagine*'.


1.6 The technology of language consists of two components: The *symbolic landscape*, a collectively-developed model of the world, and the *communication protocol* – a set of normative rules for the regulation of instructive communication. The technology allows speakers to channel their experiences to their listeners through the symbolic landscape. Linguistic communication is a process of iterative translation between four levels: The level of private-experiential meaning, the level of socially-constructed meaning, the level of socially-constructed structure and the level of physical articulation.

The symbolic landscape is what we usually think of as the lexicon, but it is much more than just a list of words and constructions. It is a huge semantic web, a radically-simplified *model* of the world of experience, whose primitives, or building blocks, are stripped-down stereotypes, digitized points extracted from the continuum of private experience. The model is a socially-constructed compromise: It reflects the entire history of negotiation and struggle, within the linguistic community, over what should be properly thought of as a *normative* worldview. The communication protocol is a set of socially-negotiated *procedures* for the process of linguistic communication. Just like the signs of the symbolic landscape, the procedures emerge from the struggle over norms – this time, the norms of communication. The two components allow speakers to channel skeletal descriptions of their private experiences - which the listeners then imagine into experiential interpretations - through the socially-constructed model of the symbolic landscape.

Linguistic communication is a process of *iterative translation*. It begins, in the speaker, with an *experiential intent*. An experiential intent is anything that is, for any given speaker, at any given time, the object of the *will to communicate*. The will to

communicate, and its object, are experiential, private, and prior to language. They are there, whenever we speak, before we start speaking: We will to communicate before we actually do. They are there, in the minds of children, before they acquire language. And they *were* there, in the minds of our ancestors, before language began to evolve.

The speaker's intent, then, which is experiential and private, is translated in the speaker's mind into a *message*, a simplified, digitized, stereotyped mental representation, which obeys the internal logic of the symbolic landscape, and is thus already socialized, and consequently communicable. The message is then copied into an *utterance*, a formal structure, still mentally represented in the speaker's mind, whose building blocks are the signifiers corresponding to the signifieds of the message. The utterance is phonologically and phonetically processed, to produce the sound waves which are then physically transmitted to the listener. The listener constructs a mental representation of the *utterance* on the basis of the phonetic information received from the speaker, which is then translated into a representation of a *message*, which is subsequently translated (complexified, analogized, re-contextualized, imagined) to produce a private, experiential *interpretation*:



1.7 The key to the strategy of instructive communication and its implementation lies in the fact that it requires a huge amount of collective effort to make it work, prior to actual communication - an effort of *mutual-identification for language*. It is precisely in this sense that language is a technology: It has to be built before it can be used.

In instructive communication, the listener is not invited to share an experience with the speaker, but to create an independent experience, on the basis of the skeletal formulation of the message, within his or her own experiential world – in isolation from the experiential world of the speaker. In the creative activity of imagination, the listener may in principle imagine in an infinite number of ways, all of which would always follow the analogue complexities of his or her own experiential world, never that of the speaker. The message should thus be able to instruct the listener in a process in which he or she not only has to create a more or less focused image (an object of imagination, not necessarily a visual representation) – but also a focused image that more or less corresponds to the original experience of the speaker: An image of the same *type*. This is a very ambitious goal. The strategy of instructive communication does this through the coordinated

investment of enormous social energies in the never-ending process of careful mapping and marking of those points in experience, and those ways of speaking, which the different speakers within the community may, more or less reliably, count on in the process: 'When I use *this* word, imagine a thing of *this* type (not *that*)'; 'when I use this word together with *this* one, imagine *this* type of experiential relationship (not *that*)'; 'when I arrange the words in my sentence *this* way, imagine you look at the whole experience from *this* type of perspective (not *that*)'; and so on.

The question, then, is this: What actually happens in the process of mutual-identification for language? How is language built for instructive communication? Note that this formulation of the question moves the focal point of our interest from the construction of language inside the individual brain to the construction of language in the community, and thus turns a question about language acquisition into a question about social invention, development and propagation. Later on, I will claim that this change of perspective allows for a major re-formulation of the original question of acquisition – and opens the way for a very different set of answers.

2. The Symbolic Landscape

2.1 The process of the mutual-identification of the symbolic landscape does not result in the replacement of private experience with the shared model of the world. It creates another level of meaning – a digital level of semantics – that is represented in our minds side by side with the level of experience. The two levels of meaning are only partially correlated with each other. In the course of linguistic communication, we translate back and forth between them, always engaging in a process of meaning approximation between the private and the social.

Let us imagine a community of four human individuals, A, B, C and D, who, for one reason or another, have to live together and communicate on a regular basis. Assume, however, that they have absolutely no language in common. How would they go about building a symbolic landscape for themselves that would meet their communicative needs? The challenge that our four individuals face is usually thought of as that of the invention of structural means (words, expressions, constructions) for the communication

of meanings that are self-evidently shared (or at least shareable). But the challenge is actually much deeper. What stands between the four individuals and their language is the foundational fact of the experiential gap. They do not see the world in the same ways, and their worldviews are inaccessible to each other. The challenge they face is that of the mutual identification of *similar-enough experiences* to talk about in the first place.

Imagine A pointing at a chair and then uttering the word **chair**. This is the paradigmatic naming event. What is happening here? Well, the pointing means: 'Let *this thing* be a common experience'. The pointing acknowledges the experiential gap, and calls upon the others to gather around, concentrate their attention on *that* particular object, and make a mental note of the fact that the experience of looking at the object, of identifying the object as a focal point of attention, was something that was, at that moment, shared by all the others. What A means by pointing is: 'If, at this moment, you all follow my finger, we shall all, for just this instance, have a common experience'. The naming adds: 'In the future, to refer to an experience *of this type*, let us use the signifier **chair**.'

The linguistic sign signifies a point of mutual-identification of the experience of its referent. When members of a community reach an agreement about a sign, what they say to each other is: 'Here we talk about the same thing'.

But they never really do. Assume that B, C and D do indeed look in the right direction, and manage to associate the experience of the right physical object (the chair, not one of its legs, or the upholstery, or the color) with the sign chosen by A to designate it. Assume, in other words, that all three of them manage to meet Quine's *Gavagai* challenge. Have they reached an agreement on the sign **chair**? Well, they have reached an agreement on what we may call the sign's *experiential-anchor*: The specific chair A was pointing at. This is no trivial matter. The experience of the chair is now marked by all four as an experience shared with the others, and the sharing is marked by all four by the association with the same name. This would allow, from now on, for communication of a kind that was not available before: If one of them comes across an identical chair somewhere else, where the *others are not present* and pointing is thus impossible, he or she will be able to indicate to the others that a certain chair exists, and, together with some hand waves, may actually be able to get the others to come along and pick it up. This is already a significant achievement.

It is, however, also a very *partial* achievement. The crucial question is how the four individuals might generalize over the meaning of the new sign: What could be other experiences *of the same type*, experiences in which the signifier **chair** might be adequately used? So far, nothing in what the four individuals managed to establish implies that they would come up with the same generalization. The fact that each of them carries an entirely different experiential history, a different way of looking at the world, actually indicates that they would probably zoom in on very different ideas.

Imagine, then, that what A had in mind was a functional, experiential generalization: A chair is an object used for (the physical experience of) sitting. For A, everything used for sitting might be called a **chair**. Assume, however, that B has never seen a chair before. Having no real idea of what the object is (what it is supposed to be), B might register the word **chair** as designating *things that look like this*. For B, the word **chair** would be adequately used to refer to objects whose *visual* experience would be similar (or similar-enough) to the original experience-anchor. C, on the other hand, is mostly impressed by the fact that the object pointed at was a regular, four-legged dinner-table chair, not an armchair, an office chair, a recliner, or a beanbag. C, evidently experienced in the world of furniture, decides that the word **chair** refers specifically to dinner-table chairs. Finally, let us assume that D, exhausted after many hours of walking, happily sees the designated object as A saw it in the first place: Something to sit on.

At this very early stage, then, our four individuals already have everything they need in order to find themselves systematically misunderstanding each other whenever any of them attempts to communicate about office chairs, armchairs, cushions and beanbags, recliners and stools – and everything which might, from this angle or the other, look like the original chair. Assume, for example, that D sees a stool somewhere, and then goes back to tell the others (with the by-now familiar combination of hand-waves and the word **chair**), that there is a chair – something used for sitting – for them to pick up. They go along with him and see the stool. For A, all this would be natural. B and C, however, would probably be confused. Both would think (not for the same reasons): 'What is *this*? This is not a chair'. C might think: 'Had I known that this was a stool, I wouldn't have taken the trouble to come here in the first place'.

Assume, then, that C now decides to *make a point*. He points at the stool and says – **stool**. What does he mean by that? Well, he means to indicate that the experiential extension of the word **chair** should somehow be *limited*. Stools are to be excluded. What he says is: 'I insist that this new experience which we are now sharing is *different* from the one we have named **chair**. In the future, let us refer to experiences of *this* type by using the signifier **stool**.' Note that C does not *have* to do all this – he may very well decide not to – and the others do not have to agree. If they do, however, they will now find themselves connecting the two signs in a way that reflects the history of their linguistic negotiations, as it impressed itself on the background of their entire experiential histories. D would think, for example: 'a chair is anything that might be used for sitting, as long as it is not a stool'. For D, the relationship would be one of set exclusion: A stool is not a chair. A would think: 'A stool is nothing but a type of chair, but if C insists . . .', and he would mark the two signs in terms of set inclusion: Every stool is a chair, but not every chair is a stool. B would think: 'A chair is something that looks like what we saw before; a stool is that thing that we used to sit on when we were children'. And so on and so forth. An endless process.

Throughout the process of mutual-identification, a certain question will loom large over each and every round of linguistic negotiation: When does something in the experience of any of the members of the community count as *different enough* vis-à-vis what has already been established, to deserve its own sign? The question is that of the *threshold of distinction*, and it is, again, a social question: There is nothing objectively necessary in the distinction between a chair and a stool, nothing that exceeds the necessity of dedicating different names to chairs with four and three legs, expensive chairs and cheap chairs, wooden and metal chairs, comfortable and uncomfortable chairs, new and old chairs, and so on – *ad infinitum*. Indeed, the great majority of experiential distinctions which any of our four individuals will notice in his or her lifetime between types of chairs ('this one looks funny'; 'that one is heavier than the one we carried yesterday') will never find its way into their gradually-emerging lexicon. The distinction between the chair and the stool is only there because C decided to make a move, and that the others agreed: The

decision to situate chairs and stools on the two sides of the threshold of distinction was eventually a collective decision.⁵

Reality is a vast continuum of endless distinctions. Everything is always similar in some ways to everything else, and different in some others. We experience the world analogically: The categorical distinctions which emerge in our minds throughout a lifetime of experiencing are continuous, fuzzy, dynamic, transient, constantly variable. They are context-dependent, and the contexts they depend on are exactly those contexts within which we have accumulated our experiences. This (among other reasons) is why we find ourselves categorizing the world in different ways.

The entire rationale of a language, then, is that of *digitalization*. In the process of negotiating their two signs, our individuals picked up two particular distinctions from the continuous landscapes of their experiences – the distinction between chairs and everything else, and the distinction between stools and chairs - isolated and highlighted them, and marked them as mutually-identifiable by all the members of the community. They produced a *digitized landscape* - at the moment including no more than two points – in which isolated points of reference stand for what they have already managed to establish as tentative experiential commonalities. Experiential complexity has been sacrificed for the construction of *common ground*.

Again: The single most important thing to understand about the construction of this common ground is that it does *not* result in a shared worldview. Experiential complexity is not replaced by the shared model. The process of mutual identification creates another level of meaning – a socially-constructed level of *semantics* – that is now represented in our four individuals' minds side by side with the level of experience. The two levels of meaning, the private-experiential and the social-semantic, are only very partially correlated with each other. The symbolic landscape is there to coordinate communication between speakers whose actual experiences are different. In the course of linguistic communication, we translate back and forth between these two levels of meaning, always

⁵ This does not imply that the process of mutual-identification is *egalitarian*: Social power plays a crucial role in the process.

engaging in a process of meaning approximating between the private and the social. The key to the understanding of language lies in this process of approximation.

2.2 The real world is only reflected into the symbolic landscape through the negotiated prism of social consensus. Language does not refer on its own. Reference is achieved by speakers.

Language refers neither to the real world, nor to the world of experience. It mediates between the different experiences of its speakers, which means that it *creates* a world, the symbolic landscape, within which its speakers can communicate. The real world is related to this constructed world not through the experiences of the speakers, but through their recognition of mutually-identifiable experiences.

Reference is accomplished not by language, but by speakers. Speakers use the technology of language to refer their *listeners*, through language, to experiences of the world. The idea that language refers by itself is yet another component of the rationalist picture of language that we must leave behind. It is an attempt to understand language, in Wittgenstein's formulation, "when it goes on holiday". The understanding that language does not refer thus reduces an entire set of reference-related problems in the philosophy of language to the status of non-problems.

2.3 The social negotiation of the symbolic landscape is not just a coordinated effort directed towards the widening of the scope of communication – it is also a struggle over the allocation of expressive power.

The fact that the symbolic landscape highlights and digitizes a certain subset of experiences means that for every speaker, at every point in time, the set of *linguistically-expressible* intents is always a subset of the general set of possible intents. The symbolic landscape licenses certain types of communicative intents, and rejects the others. Thus, the closer the symbolic landscape is to a speaker's experiential world, the better it matches the speaker's communicative intents, the more expressive power the speaker will have when talking. This positions the relationship between language and symbolic power – usually absent from the theoretical discourse on language as such – at the very center of the general theory of language, with implications that I will not discuss here.

2.4 The process of mutual identification begins from scratch with every new generation, and re-opens with every social re-grouping. Every round of mutual identification thus produces new variations, and launches new processes of linguistic change. Variation and change are not peripheral to the essence of language. They are foundational facts about language as a technology: It is always under construction.

The literature on sociolinguistic variation and language change (especially in the last two decades) has deepened our understanding of the web of complexities involved in the propagation, diversification and conventionalization of new linguistic variations.⁶ As many authors admit, however, the most fundamental question – why do variations appear in the first place? What is it about *language* that makes them necessary? – has not yet received a principled answer. In specific cases, where researchers painstakingly analyze the site of innovation, we get very good dialectic descriptions of the social dynamics surrounding innovations – but these only explain the innovations as social facts, not as facts about language. The linguistic question is not why specific innovations appear when they do: The question is why innovation never stops.


There is a very particular reason for the fact that this question has rarely been approached. The emergence of Sociolinguistics and Anthropological Linguistics was motivated from the very beginning by three complementary goals: To establish the scientific fact of linguistic variation within languages (as opposed to the idealizing view of language in Theoretical Linguistics); to analyze the relationships between the attested patterns of variation and their social contexts (in terms of Penny Eckert's third wave of sociolinguistic research, to analyze variation 'as a resource for the construction of social meaning'); and, more generally, as Bucholtz and Hall put it, to 'create a field that would put language at the center of social and cultural life'. Because of all this, no attempt has been made to construct an alternative theoretical model of language - a model that characterizes and explains variation and change as necessary facts about language as such.

⁶ See Eckert (2000, 2005) and Bucholtz and Hall (2008) for two general perspectives on the development of this literature; Johnstone (2000) for a discussion of the role of individual speakers in the processes of change; and Croft (2000) for a constructionist theory of language change.

LCT's starting point is the fact that the technology of language is only functional to the extent that it provides its speakers – in their own present time – with the level of common ground necessary for communication across the experiential gaps between *them*. Languages as we know them, however, are the products of long processes of social negotiation. They have gradually been constructed by other people, in different times, in different sets of circumstance. Thus, they are always obsolete for their speakers. Because of that, the adoption of an existing language by a new community of speakers is always a selective process: When a piece of the language proves to be useful, where it actually provides the necessary level of common-ground for the new group, it is adopted (very often with changes too). But when mutual identification with respect to a certain aspect of the world (or the world of communication) requires signification that the already-existing language cannot provide – because it does not have it, or because what it has to offer already belongs to another community with which the groups does not want to mutually-identify – the community re-starts the process of signification. It adapts the technology to fit its own communicative purposes.

2.5 Every sign on the symbolic landscape (as it is represented in the mind of every speaker of the language) maps three levels of cognitive representation onto each other (three levels, not two): A signifier (vocalic or manual), a signified (a digitized point on the symbolic landscape) and an analogue cluster of experiences, centered around an experience-anchor, in the private world of experience.

The sign **chair**, for example, constructs in the minds of English speakers an association between (i) the signifier **chair**, (ii) the signified *chair*, semantically interconnected on the symbolic landscape with the signifiers *stool*, *armchair*, *upholstery*, *legs*, *table*, *furniture*, *comfortable*, *sit*, and many others, and (iii) a private cluster of chair-related experiences – chairs of different sizes and shapes, their look, smell and feel, various sitting experiences and so on – centered around the experience-anchor of some very familiar type of chair, which, at some point in the past, found its way into the naming event (or events) in which they learned what a chair was:



2.6 Both the definitional and prototypical conceptions of word meaning have assumed a relationship between a signifier and a single level of meaning – semantic or pragmatic. The realization that the signifier is connected to two levels of meaning, that are only partially correlated, resolves a host of traditional problems in lexical semantics.

The entire Saussurian tradition in Linguistics, along with the logical tradition in the philosophy of language, only acknowledged two levels – the level of the signified and the level of the signifier. Experience was rejected as an explanatory concept. The implication was this: If the sign is to have any meaning at all, if the signified is to have any content, its meaning should be computable, either on the basis of its semantic relations with the other signifieds around it, or, alternatively, on the basis of its direct relationship with the real world. The second strategy implied that language must somehow be an objective reflection of reality. The first implied that the relations between a signified and its neighbors must make up a set of necessary and sufficient conditions for the definition of the sign. The signified *chair*, then, should either receive its meaning from the direct

referential relationship it maintains with the set of chairs in the real world, or, alternatively, be definable as, for example, 'a physical object, which has four legs and a back, and is used for sitting'.

The evidence accumulated throughout the 20th century clearly indicates that this does not work. Entities in the real world do not come along with objective name tags, and there is no set of necessary and sufficient conditions capable of defining the word in any rigorous way.⁷ Two arguments are worth mentioning here: Definitions are never *exhaustive* – there are always going to be non-chairs wrongly included in the definition, and genuine chairs left out – and they are always *circular*: We start by defining chairs as physical objects used for sitting, and must then find definitions for signifieds such as *physical*, *object*, and *sit*, and so on, ad infinitum.

The growing disenchantment with the semantic analysis of the sign eventually produced what we may term the Wittgensteinian reaction – in which the semantic relationship between signifieds was replaced by what should be properly described as a pragmatic relationship between the signifier and the cluster of experiences associated with it. The reaction thus maintained the *double nature* of the sign, and replaced one level of meaning with another. The sign now referred to a fuzzy set of experiences, some more *prototypical* than others, which, taken together, produced a sense of familiarity, a sense of *family resemblance*, with which it came to be associated. This move resolved both the problem of exhaustiveness and the problem of circularity. It also created new, and seemingly unresolvable, problems.⁸ Two of these are worth mentioning: First, it seems that speakers sometimes make judgments of prototypicality which are surprisingly unrelated to what they actually think of as the meaning of the word in question. Speakers, for instance, consistently describe prototypical grandmothers as elderly, old-fashioned women who make chicken soup and spoil children with candy, but also consistently agree that the word **grandmother** refers to 'a female parent of a parent': Someone who looks and behaves like a prototypical grandmother may not be a grandmother at all, and someone may be a genuine grandmother without meeting any of the prototypicality conditions mentioned above. The second problem is this: If all there is to a sign is family

⁷ Foundational references are Wittgenstein (1953), Rosch (1975) and Fodor et. al. (1980).

⁸ See Fodor and Lepore (1996), Armstrong, Gleitmann and Gleitman (1983) and Osherson and Smith (1981).

resemblance, and different members of the set designated by the sign may be more or less prototypical, then nothing seems to prevent entities, which clearly do not belong in certain sets, from becoming non-prototypical members of these very sets. Consider *camels* and *mice*. If the set designated by **camel** is identifiable on the basis of family resemblances, what is it that prevents a mouse – evidently a mammal with four legs – from joining the family as a very distant relative?

It is only by acknowledging of the *triple nature* of the sign, as a socially-constructed mediator of meaning across the experiential gap, that these problems may finally be resolved. The problem of exhaustiveness simply disappears, because the set of semantic relationships between signifieds and their neighbors on the symbolic landscape is *not* definitional: It was never meant to be. It emerged, gradually and tentatively, throughout the process of the development of language, as members of communities reached agreements about distinctions and relationships between elements in their experiences which were, for them, worthy of communication. The signifieds **chair** and **sit**, for example, are semantically associated not because they were meant to define each other, but, first, because the members of the relevant communities agreed that chair-experiences and sitting-experiences were to be recognized as things to talk about, and, second, because for all of them the cluster of experiences associated with the sign **chair** was systematically related to the cluster of experiences associated with the sign **sit**. The semantic relationship between the signifieds **chair** and **sit** came to be a conventionalized connection on the symbolic landscape. Naturally, as associations of this type multiplied, the set of conventionalized semantic relations between a sign and its neighbors came to *approximate* a definition. This is a very important fact, to which we shall return, but it is also the result of a highly haphazard, unplanned, evolutionary process, in which signs were gradually added to the symbolic landscape, in order to meet the growing and changing communicative needs of certain communities of speakers, in certain specific social settings. There is no reason to assume that such a process should result in logically well-formed definitions. Indeed, it doesn't.

In other words: The fact that signs refer to focal points of attention in experience – that they are there to mark *social agreement* on these focal points – immediately implies that their meanings do not *have* to be exhaustively definable at the level of the symbolic landscape. The function they are supposed to fulfill on the symbolic landscape is much

more modest than that: They are not there to tell speakers *what* things are, or what the world is like. They are only there to direct speakers towards a certain tentative set of distinctions, to tell them which things should be thought of as different from which (for the sake of conversation), and which things are to be thought of as connected to each other (again, for the sake of communication). What these different things actually *are* is something the speakers have to figure out for themselves, on the basis of their experiences and the experience of the specific event of communication in which the signs are used.⁹ The amount of information required for this categorical function is much smaller than that required for the definitional function. There is no reason to assume that it should ever amount to a definition. This is exactly why, within this conception, circularity is no longer a problem. It *would* be a problem if we had to know what things are on the sole basis of the semantic properties of the signs referring to them. But we do not.

On the other hand, the fact that the symbolic landscape marks differences and connections between things in experience - the fact that analogue experiential clusters are associated with digitized points, the very essence of which lies in their being different from each other - implies that speakers are instructed by the symbolic landscape to focus on certain *types* of distinctions between experiences *which are otherwise not that different from each other*. Grandmothers are indeed similar in all kinds of ways to other human beings: Being old, making soup, handing candy to children - all these are properties which are spread all over the experiential continuum. The speakers of the language, who at some point decided to add the sign **grandmother** to their language, were attempting to abstract away from this experiential continuum, and highlight a single property at the expense of others – because they were specifically interested in talking about mothers-of-parents as such. This was a necessary move, not in spite of, but because of the fact that there are family resemblances between grandmothers and other types of people. Many of the grandmothers who served as experience-anchors in naming events in which the word **grandmother** was learned were indeed busy spoiling children with candy, which means

⁹ Because of this, as the growing literature on lexical pragmatics shows very clearly, the actual process of lexical interpretation involves such processes as the *narrowing* and *broadening* of the conventionally-specified meaning of the sign. For a Relevance-oriented analysis, see Wilson and Carston (2006).

that for many speakers old ladies with candy came to be perceived as prototypical members of the experiential cluster associated with the sign **grandmother**. But what they came to know about the properties of prototypical grandmothers had very little to do with what they *had to know* about the distinction between grandmothers and non-grandmothers in order to use the sign appropriately. Prototypicality cannot tell us how the members of a set are different from non-members. It can only tell us which entities are more exemplary members *of the set itself*. This is exactly why prototypicality theory, on its own, cannot ensure that mice will be excluded from the set of camels.

The meaning of a sign cannot be reduced to the complexities of the experiential cluster associated with it, any more than it can be reduced to the semantic relations it maintains on the symbolic landscape - because the process of linguistic categorization inevitably involves two dynamics rather than one: The identification of distinctions (and connections) *between* sets, and the identification of similarities *within* sets. The attempt to squeeze these two dynamics into a relationship between a signifier and a single level of meaning (semantic or pragmatic) has failed. The two dynamics take place at two different levels of meaning: The symbolic landscape distinguishes and connects between sets; experience identifies similarities within them.

2.7 The web of emergent connections between signifieds on the symbolic landscape creates a generalized view of the world that speakers can talk about. It is different from the world of experience. The strongest connections create an implicit map of the essential, abstract properties of this world: Event structure, argument structure, quantification and so on.

The semantic associations between the signifieds on the symbolic landscape do not amount to definitions. They do, however, specify how all the different things which inhabit the symbolic landscape are related to each other *on* the landscape: The property *fragile*, for example, is semantically associated with a host of fragile objects, with the event *break*, with other properties such as *frail* and *delicate*, with the antonyms *sturdy* and *strong*, and so on. Together, they construct a web of relations, a semantic field, which reflects a certain logic, neither the logic of things as they are in the real world, nor the logic of things as they are in private experience, but the logic of things as they are imagined through the social prism of the symbolic landscape. In experience, things are

never *either* fragile *or* sturdy, either frail or strong. In the world of experience, eventualities in which things break or do not do not arrange themselves neatly on the two sides of a clear line of demarcation. On the symbolic landscape, they do.

On the symbolic landscape, moreover, things are arranged in semantic relations which are not just digital variations on the analogue relations identified in experience, but are actually radically different. In experience, for example, not all things have gender. Chairs and stools, as we experience them, are not male or female. In the linguistic landscapes of many languages, however, they are. It may well be the case that the demarcation of inanimate objects into male and female reflected very deep experiences of those people, who in the distant past made the social decision to assign gender to inanimate objects in their language – but for us this is no longer the case. Nevertheless, if we speak any of these languages, and if we wish to talk about inanimate objects, we are obliged to adopt the logic of the symbolic landscape, and refer to them as either male or female, according to the digitized convention stabilized in the language.

It is important to understand that the web of semantic relations on the symbolic landscape emerges as a necessary consequence of the foundational act of *demarcation*, which gives rise to the landscape in the first place. If different animal species are to be distinguished on the symbolic landscape, the signs referring to them *as species* must abstract away from everything else that distinguishes between them in experience, including their sex. If male and female animals need to be distinguished on the symbolic landscape, the signs referring to them *as male or female* must abstract away from everything else, including species membership. Once these two demarcations are established, however, they are forced into a standardized relationship which stereotypically reconnects what was originally disconnected and isolated.

All this goes much deeper when we consider signs referring to *eventualities*, such as **break**, **dance**, **eat**, **know**, **tall**, and so on (what we usually think of as verbal and adjectival meanings.) To receive their own signs, eventualities do not just have to be digitized away from their experiential complexities – they have to be abstracted away from the very entities which *embody* them. Think about that famous line from Yeats' *Among School Children*: 'How can we know the dancer from the dance?' There is nothing in the event of dancing that is not, at the very same time, a fact about the dancer. Dancing

only happens as long as the dancer dances, and the dancer is only a dancer as long as the dancing event takes place. Watching a dance (visually experiencing it) is looking at the dancer. To assign the eventuality itself a sign, the symbolic landscape has to disembody the dance, and leave the identity of the dancer *unspecified*. The sign referring to the dancer, however, must be semantically connected to the sign referring to the eventuality, because a dancer is to be identified as different from other people on the sole basis of the fact that he or she participates in a dancing event. What this means is that eventualities come to be associated on the symbolic landscape with their stereotypical participants. This is what *thematic roles* are all about: they are social entities.¹⁰ The eventuality of **breaking** comes to be associated with a stereotypical *patient*, the thing broken, which is stereotypically breakable, and with a stereotypical *agent*, the one who breaks. The eventuality of **dancing** comes to be associated with a single participant, the *agent*, the one who dances, and so on.

Let the most robust semantic relationships on the symbolic landscape, the ones that constantly recur, make up a generalized meta-logic - a relational meta-abstraction from the digital abstractions of the symbolic landscape – which I will call the *generalized map*. Think of the generalized map, metaphorically, as a closed web of interrelated relational statements, which, taken together, describe the main essential properties of the symbolic landscape as a world. Let the generalized map include statements such as the following (which vary in different languages):

- (i) Eventualities have stereotypical participants: They have *agents*, *patients*, *instruments*, *authors*, *recipients*, and so on. (Each specific eventuality-type has a certain subset of these.)
- (ii) Eventualities are related to other eventualities as *causes*, *reasons*, *goals* etc.
- (iii) Eventualities take place within a time-frame: They either happened in the *past*, happen in the *present*, will happen in the *future*, and so on.
- (iv) Eventualities take place in a specific *location*, at a specific *time*.
- (v) Eventualities have different *manners* of happening.

¹⁰ For a more detailed discussion of thematic roles as part of the social level of meaning, see Appendix A.

- (vi) Entities have gender: They are either *male* or *female*.
- (vii) Entities are either *singular* or *plural*.
- (viii) Singular entities are either *definite* (the only tokens of their type in the context shared by speaker and listener), or *indefinite*.
- (ix) Entities have *properties*.
- (x) Entities always participate in *eventualities*.

Think of these statements as sets of emergent agreements between speakers, about the general properties of their mutually-identified world: "In the world within which we communicate, things are either singular or plural, male or female, animate or inanimate; they participate in events of certain well-defined types; these events take place in time, and they have causes and goals ...".

3. The Communication Protocol

3.1 The communication protocol is a set of mutually-identified, prescriptive conventions which guide the speaker and the listener in the processes of linguistic production and comprehension. The protocol is not a grammar. It governs linguistic communication in the way law governs social life: Not all laws are equally internalized; not all laws are equally enforced; not everybody obeys the law all the time.

What the symbolic landscape does for the speakers' general experiences of the world, the communication protocol does for their *experiences of communication*. Different speakers experience events of communication – instances of communicative success and communicative failure, confidence and confusion, understanding and misinterpretation, trust and suspicion, sincerity and deception - in an infinite variety of ways. They occupy different positions in these communication events; they look at them from their own particular perspectives (as speakers, as listeners). They come into them from different histories and different social positions, with different habits and sensitivities, different capacities (for expression, understanding, trust); different attention spans, sets of beliefs, expectations and interests. They share a moment of communication, but this moment is nevertheless an attempt to construct a very narrow, very shaky bridge over the constant

experiential gap. This is a very precarious endeavor: It can only work to the extent that the participants in the collective effort manage to coordinate their efforts; and in order to be able to do that, they have to be able to mutually-identify, at every given moment, what the other person is trying to *do*.

The communication protocol, then, emerges through the entire history of a language's development, from the constant social struggle over the conventionalization of communicative behaviors which are to be thought of as mutually-identified by the community. Think of the rules and regulations comprising the communication protocol as specifications of certain ways of speaking, those ways which the members of the community have already mutually-identified, highlighted and marked (not necessarily in an egalitarian fashion) as *co-operative* behavior.

When the members of a community agree on a behavioral convention, what they say to each other, across the experiential gap, is: 'Here we *do* the same thing.' (And, again, they never really do.)

At its core, the communication protocol includes a practical guideline for the actual production of utterances - a set of ordered procedures which lead the speaker, step by step, through the process of experience-to-speech translation. This set of procedures is related to the actual process of production in the same way as law is related to the dynamics of social life: It specifies the socially-established ways in which the process should take place. To the extent that it is internalized, and to the extent that it is enforced, it actually regulates and constrains the process in real time. The set of procedures does *not*, however, do more than regulate and constrain. The activity of communication is not brought to life by the conventions which regulate it. It originates from, and is motivated by, the very will to communicate. This is the reason for another crucial similarity between linguistic communication and its conventions and social life and the law: Communication is always regulated by conventions, but it is also, at the very same time, a constant attempt to break away from them.¹¹

¹¹ This conception should help us come to terms with a host of phenomena that reside at the boundaries of linguistic communication – the will *not* to communicate, the sense of *betrayal of*

The communication protocol is not a *grammar* in the cognitive sense of the word. It does not supply the speakers with a set of formal representations that demarcate grammatical from ungrammatical utterances. Rather, it tells the speakers *what to do* in the process of speech production. Together with the symbolic landscape, however, the procedures of the communication protocol do maintain a principled relationship with the notion of *grammaticality*: When a speaker follows the conventionalized rules to the letter, the end-product (the actual fragment of speech) is judged by the other members of the community as grammatical. When making this judgment, the other members of the community are not interested in the fragment of speech itself. They ask themselves whether the *speaker* has followed the conventionalized procedures established by their community. They do not ask: 'Is this sentence grammatical?' They ask: 'Does the speaker sound like one of *us*?'. Grammaticality judgments are always identity judgments (as the sociolinguistic literature shows). Eliciting grammaticality judgments from native speakers is thus always a *political* act, and grammaticality judgments are always methodologically complex, variable and vague. Speakers are defensive about grammaticality judgments: As far as they are concerned, it is their very membership in their speech communities, not the sentence they are presented with, that is being tested.

As linguists, we sometimes say to native speakers: 'Forget what they taught you at school about the way the language *should* be spoken; just tell us how this sentence sounds to *you*'. We think we are going beyond conventional norms; touching something much deeper. Actually, we are telling our native speakers: 'Forget about the norms you have only *heard* about; pay attention only to those you have already *internalized*.'

We, as linguists, have internalized the distinction between *descriptivism* and *prescriptivism* to such an extent, that we no longer realize that the object of Linguistics as

self that accompanies the linguistic socialization of a private intent, the notion of the *secret*, and *Freudian slips*.

a descriptive science is itself a system of prescription. Grammaticality judgments are *judgments* for no other reason than the fact that they are based on prescription.¹²

3.2 Speech acts are socially-negotiated, stereotypical communicative behaviors: They are part of the protocol. (The two sides to the Searle-Derrida debate thus highlight the two faces of this reality: They are both right.)

Our communicative intents very rarely, if ever, emerge in our minds as digitally-demarcated intents to either *say* or *ask* something, to *order*, or *promise*, or *predict*, or *deny*, and so on. They are multi-layered, variable, vague, dynamic, analogue. We wish to *express* something, and what we wish to express is as complex as the experience within which the communicative intent emerged. Coupled with the foundational fact of the experiential gap, this analogue complexity constitutes a major obstacle to communication: We very often find it difficult to understand what the person speaking to us is trying to do (*'Is this a promise or a threat?'*). And our experiential histories often lead us to the wrong conclusion.

Think of speech-acts, then, as socially-negotiated, stereotypical communicative behaviors, highlighted and isolated from the experiential continuum of communication, which, when practiced according to a set of mutually-identified conventions, allow for the successful mediation of stereotyped forms of communicative intents across the experiential gap. When conventionalizing a speech-act, what the members of the community agree on is this: 'From now on, when we behave *this way* – when, in these particular contexts, we use this intonation, this word order, this gesture - we mean to ask a question (or make a promise, or tell a story)'.

We have learned from J.L. Austin that speech acts are ways of *doing things with words*. But where do they come from? What is the source of the distinctions between different speech acts – between assertions and questions, promises and requests? Austin himself, and later John Searle, maintained that speech-acts are *natural* phenomena, distinguishable

¹² All this, of course, follows a tradition of thought that conceptualizes about language as a normative-conventional system, and usually begins with Lewis' (1969) definition of convention. See, for example, Itkonen (2008) and Zlatev (2008).

from each other (and dependent upon each other) on the basis of their logical properties as forms of communication. The formal properties of speech acts (as analyzed, for example, by Searle and Vanderveken), are facts about the very essence of language as a logical system of communication: When a speech-act is performed under certain, well-formedness conditions (*felicity* conditions), the speech act is successful. When the felicity conditions are not met, the speech act fails.

This conception has been famously attacked by Jacques Derrida, who has claimed that meanings are always polysemic, contexts are never totally determined, intentions are never transparent – and speech acts are hence always different, opaque, indeterminate. The option of failure is not something that lies outside the domain of language, something that can be systematically avoided by following a certain formal logic. It is always there, within language, as "the very force and law of its emergence", because the reality of language is always an "endless alternation of essence and accident", never the stable, idealized system envisaged by Austin and Searle. The essence of speech acts lies precisely in the fact that they cannot be exhaustively defined.

Derrida is right, but so are Austin and Searle. Speech acts *do* have specific formal properties, but they have them not because they reflect an inherent logic *of* communication, but because they are socially-conventionalized *norms for* communication – conventionalized norms which are necessary because (not in spite) of the fact that the experiential world within which linguistic communication has to take place looks just like Derrida describes it. What Derrida is talking about is not language itself, but the experiential gap that is "the very force and law of (the) emergence" of language, and because of that, Derrida is perfectly right: The experiential gap is never *really* bridged. Austin and Searle, on the other hand, are talking about language and its conventional norms, and because of that – they are not idealizing. They are *describing* the idealized conception that lies at the very basis of language as a conventional system - the idea that perfect communication across the experiential gap would be possible if only all members of the community followed the linguistic conventions to the letter.¹³

¹³ The relevant references here are Austin (1962), Searle (1977) and Derrida (1977) - and see Kenaan (2002) for a discussion of the Searle-Derrida debate.

4. Linguistic Production and Comprehension

4.1 Linguistic production is a process of iterative translation. In the first stage, the speaker's private-experiential intent is translated into a social-semantic message. This is done in prescribed stages: The message is built top-down, from generalities to details. The complete message is then translated (again, in a staged manner) into an utterance, which is then actually articulated.

The following is a radically simplified first step towards a technical description of the process of linguistic production. Throughout the analysis, I will use a simple sentence in English as an example. Nothing is implied by this choice, certainly not that the conventions involved are universal. The analysis of the example serves demonstration purposes only – it should not be read as a full technical analysis. The discussion intends to describe the foundational properties of the process, its general architecture – and thus, eventually, to develop a general framework within which the analysis of the more complex issues should take place.

Note, moreover, that this is not a description of linguistic processing in the sense of cognitive science – but of the process as it is *prescribed* by the conventions. In actual speech, speakers do not necessarily do things in exactly the same ways: The habits of communication make up their own shortcuts – both in the form of expectations and constructions.¹⁴

¹⁴ This description of the process bears some obvious similarities to the models currently discussed in the psycholinguistic literature – most of which follows the path of Levelt (1989): Thus, for example, the description proceeds from the production of lexical meaning (the lemma) to the production of lexical form (the lexeme). There are, however, crucial differences: All the models are models of on-line processing; models of performance based on grammar-like theories of competence with a rich syntactic component; they all assume a relatively transparent translation between a pre-linguistic meaning that is already conceptual and the linguistic 'lemma'. None of them conceptualizes about message production as an iterative process that produces a semantic hierarchy.

Let us begin with the stages of message construction. In the first stage, which I will call DECIDE, the speaker is required to make the most radical act of translation – to abstract away from the analogue complexities of his or her communicative intent, and translate it, *reduce it*, into a stereotyped kernel of meaning, the *message-kernel*, that the language - and thus the listener - will be able to identify. The message-kernel is a convention: Different languages present speakers with different varieties of message-kernels. At the moment, let us concentrate on one variety, which seems to be very widespread. In this variety, the message-kernel is a meaning unit of the type –

SPEECH-ACT_{(topic-entity)(eventuality)}.

Think of this message-kernel as a conventionalized demand - formulated by the community, and addressed to the speaker - to specify, from the very beginning, what (or whom) he or she is going to talk about, what happened (or happens, or will happen) to it, and what he or she intends to say about it – tell a story, ask a question, make a prediction, and so on. It is a social requirement, which came to be conventionalized through a long and difficult process in which listeners had to interrupt speakers and insist, time and again: 'Before you go on, can you please tell me what we are talking about?'.

The message-kernel obligates the speaker to decide, from the very beginning, whether he or she wants to *Assert*(John)(dance); *Ask-whether*(Bill)(tall); *Promise*(child)(eat), and so on. These meaning triplets are as of yet untranslatable into English sentences: They lack an indication of tense, for instance. Each of them is a first step towards the construction of a message, towards the socialization of the speaker's intent. The speaker is only allowed to select one speech-act out of the pre-established set of speech-acts offered by the conventions of the language, and is only allowed to select one entity and one eventuality, out of the pre-established set of entities and eventualities on the symbolic landscape, as the objects of the speech-act. The entire set of possible message-kernels, everything that the language allows for, is already-specified prior to the act of communication.

It is important to appreciate the challenge posed by the message-kernel prescription: Imagine that you come home from work and find that a very precious object (an expensive vase, for instance) has been broken. You have every reason to suspect that it was a certain person who did it, but you are not quite sure. He is still there, but he is not

saying anything. You are upset, angry, confused. You decide to approach him. What are you going to say? Are you going to *ASSERT*_{(vase)(break)}? *BLAME*_{(you)(break)}? *ASK-whether*_{(you)(break)}? *ASK-why*_{(you)(break)}? *ASSERT*_{(I)(angry)}? *THREATEN*_{(I)(hit)}? *DEMAND*_{(you)(compensate)}? All of these are probably going through your mind simultaneously: The will to communicate emerging in your mind is as complex as your experience. And yet, each of these options, once highlighted and isolated, has an entire array of discursive implications: Assume you *ASK-why*_{(you)(break)}, and then discover that it wasn't him; or assume that you *ASSERT*_{(vase)(break)}, and the person reacts with what seems to you like a fake expression of surprise. What are you going to do then? To be sure, language allows you to formulate a *series* of consecutive utterances. However, it also obligates you to formulate them *one after the other*. The question remains: Which are you are going to formulate *now*?

In the second stage of message construction, which I will call *SPECIFY*, the speaker is required to flesh out the message-kernel and develop it into what we shall call a *basic message*. The speaker is presented with a closed set of extensions to the message-kernel, each of which is specified by those statements of the *generalized map* which associated with the specific *eventuality* chosen for the message-kernel. The speaker is required to *fill in* contents for some of these extensions, which have by convention become *obligatory*, and is allowed to specify contents for the remaining extensions. This can be thought of as second step of the collective attempt to enforce a certain order on speakers: 'Now that we know what you are talking about, can you please tell us, in a way that we will understand, what *happened*?'

A speaker of English, who selects a message-kernel such as *ASSERT*_{(boy)(break)}, for example, is presented at the second stage with the following set of obligatory and voluntary extensions, each of which is based on one of the statements of the generalized map which is associated with the eventuality *break* (starred extensions are obligatory¹⁵):

¹⁵ Obligatory extensions are normative and prescriptive. They are there because communities decided to obligate their speakers to include certain meaning components in their messages, *whether or not* these components featured prominently in their communicative intents. By turning a categorical distinction from the symbolic landscape into an obligatory component of every possible message, the members of the community, in their role as listeners, impose a *minimal*

	* <i>patient</i>	_____
	<i>agent</i>	_____
	<i>instrument</i>	_____
	<i>goal</i>	_____
Extensions of event-type (<i>break</i>):	<i>reason</i>	_____
	<i>manner</i>	_____
	* <i>tense</i>	_____
	<i>time</i>	_____
	<i>place</i>	_____


The resemblance of this scheme to an official form is no coincidence. Language is *formal* in this sense exactly: It provides the speaker with a closed set of pre-determined forms of expression, and forces him or her to translate his or her intent into these forms. Through the scheme, the speaker is presented by language with a closed set of things that can be said (or should be said) about eventualities of *breaking* – precisely those which were isolated from experience and mutually-identified by the community.

Assume, then, that our speaker fills in the form in the following way:

specificity requirement on themselves as speakers. I will later claim that this plays a crucial role in the complexity surrounding the Sapir-Whorf hypothesis.

	* <i>patient</i>	<u>vase</u>
	<i>agent</i>	<u>boy (topic)</u>
	<i>instrument</i>	_____
	<i>goal</i>	_____
Extensions of event-type (<i>break</i>):	<i>reason</i>	_____
	<i>manner</i>	_____
	* <i>tense</i>	<u>past</u>
	<i>time</i>	_____
	<i>place</i>	_____


In the third stage, COMPLETE, the speaker is required to further develop the basic message into what we shall call a *complete message*. Each of the selections made by the speaker in the former stage now leads to another round of extensions, some of which are obligatory. These extensions are filled in, and, if necessary, prompt another round, and so on, until all obligatory spaces, and all the spaces chosen by the speaker, are specified according to convention. The process may look like this:



Having specified the vase as the *patient*, for example, the speaker is presented with a set of extensions to the patient-vase, a set of conventionalized questions, some of which *must*, some of which *may* be answered: Is the patient a single entity or a plurality of

entities (a vase or vases)? Does it have any property that needs to be mentioned? Did it (or does it, or will it) participate in another eventuality which has to be reported? The choice of singularity leads to another obligatory extension: Is this entity already mutually known by the speaker and the listener (is it definite)? Think of this as the third step in the collective effort: 'Now that we know what we are talking about, and what happened, you may add some *detail* of the type that we can understand'.


The end-product of these three stages, the *complete message*, is thus a formal representation of a selected set of interconnected primitives from the symbolic landscape, *hierarchically arranged by order of selection*:



The next three stages in the process of production are dedicated to the construction of the utterance on the basis of the complete message.

In the first stage, REPLACE, the signifieds of the complete message, hierarchically-ordered as they are on the basis of their order of selection, are replaced by their respective signifiers. Think of this process, metaphorically, as similar to the copying of a DNA

sequence into RNA. In most cases, replacement takes place one-to-one: Most signifieds on the symbolic landscape are, by their very conventional nature, uniquely attached to single signifiers. Sometimes, replacement would take place many-to-one. Let me call the product of this stage – *the basic utterance*:




In the next stage, *LINEARIZE*, the signifiers comprising the basic utterance are arranged in linear order – following the *linearization conventions* of the language. Let me call the output of this state – *the linear utterance*. This is the second step towards the construction of the *full utterance*, which must be ready to be transmitted as a stream of speech. (This is why the process of linearization takes place only *after* *REPLACE*, and applies to signifiers, not to signifieds: It can only apply to structural entities with a temporal dimension.)

As a first approximation, let the process of linearization *in English* be governed by the following set of conventions (there are additional complications, of course, which are not of interest at the moment):

- (i) linearization takes place iteratively, bottom-up on the selection hierarchy;
- (ii) quantificational extensions (such as *singular-plural*), and *property* extensions, should appear to the left of their bases; the remaining extensions should appear to the right of their bases;
- (iii) when obligatory and non-obligatory extensions appear on the same side of the base, obligatory extensions should always be on the left;
- (iv) the *topic-entity* of the basic message should always be on the left of *eventuality*.

The first convention of the above set implies a process in which every base is always located in the utterance with its extensions *adjacent to it*. This is what strict word order (which is a fact about *English*) is all about: It copies the semantic relations between the signifieds into the linear order relationships between the signifiers. And again: Strict word order is a prescriptive convention - 'always mention the thing and its details together, so that we will know that the details are *about* that thing'. The other three conventions organize linear order on the basis of the properties of the signified message, as they accumulated throughout the process of production. The prescriptive convention in this case is this: 'Situate the different participants in your message in mutually-identified positions along the sentence, so that we will be able to recognize which is which'.

In our rudimentary example, then, linearization would proceed in the following manner:



In the next stage, CONNECT, the by-now ordered signifiers of the linear utterance are structurally connected to each other, and a set of *phonological* procedures are applied to them. Two sets of conventions are involved here: *Morphological* conventions determine the extent to which two signifiers or more, which are adjacent to each other, are to be phonologically *fused* into a single word or phrase; phonological conventions determine what such fusion should sound like. This is exactly why the pre-theoretical notion of a *word* is in the final account an intuition about *phonological boundaries*. (And this is also why optimality theory seems to be in the right direction with respect to phonology: It is best interpreted as a social theory of the competition between different sets of norms.)

In the last stage, SPEAK, the full utterance goes through a set of *phonetic* processes (deriving from the phonetic conventions of the language), and is actually realized as a stream of sound.

4.2 The process of production reveals itself most clearly in broken speech - in the real-time struggle, within the speaker, between the private will to communicate and the social conventions of the language.

In regular, spontaneous speech, the entire process of message construction – from the selection of a message-kernel to the specification of the complete message - is almost always subconscious. The first time we get acquainted with our own message, then, is when we *hear ourselves speak*. This is why we so often feel (and sometimes say): 'That was not what I wanted to say', 'I didn't mean it that way', 'I know exactly what I want to say, but for some reason I can't *find the right words*'. The process of message construction (subconscious or not) is always a struggle to produce the best possible approximation to the intent. It does not always work. We construct a message, translate it into an utterance, start speaking it, listen to the first words, realize that this is not what we *meant to say*, stop, hesitate, try again, and so on. We enter a feedback loop with ourselves, and sometimes with our listeners, through which we struggle to approximate our intent.

General theories of language have always concentrated on the fully grammatical utterance as a reflection of linguistic competence (specific or general-cognitive). Broken speech – as a fact about performance - has been relegated to the outskirts of the discipline. It has attracted attention in Psycholinguistic research, especially in the context of *speech errors*, and in Conversational Analysis, especially in the context of *repair strategies*.¹⁶ The understanding of language as a communication technology, however, positions the phenomenon of broken speech at the center. Broken speech is a foundational fact about the technology: In most cases, when the technology is used for actual communication, this is what it produces.

This does not mean that speakers' judgments are not important. They are crucial. But our object of inquiry should be the discrepancies between the judgments that the speakers make *as members of their linguistic community*, and the utterances that they actually produce *as individuals in actual communication*.

¹⁶ On speech errors, see Levelt (1993) and Poulisse (1999). On repair, see Sacks (1995), Schegloff, Sacks and Jefferson (1977) and Schegloff (1992). Broken speech also plays an important role in clinical psychology.

4.3 Listeners follow the mutually-identified conventions of the protocol as they go through the process of comprehension, all the way from the analysis of the stream of sound to the production of the experiential interpretation. The challenges they face, however, are different: Ambiguity, for example, is always a listeners' challenge. Because of that, the conventions of the protocol reflect a history of struggle between speakers and listeners over the regulation of linguistic communication.

As they move along the process of comprehension, listeners go through four consecutive stages. Let me call them IDENTIFY, REPLACE, RECONSTRUCT and RECOVER.

In the first stage, IDENTIFY, listeners apply a set of sound analysis procedures to the incoming stream of speech. Using the conventions of *phonetics*, and then *phonology* and *morphology*, the listeners attempt to achieve the goal of utterance-identification. The first goal is achieved once the listener has managed to break down the stream of speech into a set of identified signifiers - still arranged in linear order but already disconnected from each other. The output, then, is the *linear utterance*.¹⁷

It is already here that a major difference is revealed between linguistic production and linguistic comprehension: The level of *proficiency* required for the task of comprehension is always lower than that required for the parallel task of production. Identifying a speaker as belonging to a certain community (performing the task of judgment) is always easier than speaking in an identifiable way, and breaking down the stream of speech into identifiable signifiers is always easier than fusing a set of signifiers into an identifiable stream. This is so, among other things, because there is an inherent element of *opportunism* in the process of comprehension which cannot be a part of production. When the speaker is missing a word, communication is halted (speech is *broken*); when the listener fails to identify a signifier or two at the first stage of comprehension, the entire process that lies ahead provides him or her with ample opportunity for compensation (an opportunity that may or may not be used; compensation is a serious challenge). There is a constant element of *guessing* in the process of comprehension which eventually allows

¹⁷ At this stage, listeners also already attempt to identify the extent to which the speaker belongs to their community. This is why accent plays such a crucial role in the linguistic construction of social identity (and, eventually, of social power).

listeners to produce reasonable interpretations, at different levels of complexity, of utterances which they themselves might not have been able to produce. It is a pivotal fact that is always there, from the first stages of language acquisition all the way to the ability of adult readers to understand (and *identify with*) works of literature that they could not have written themselves.

In the second stage of comprehension, REPLACE, the by-now disconnected and identified signifiers are copied onto their respective signifieds. This is the stage in which *ambiguity* asserts itself, for the first time, as a constant danger: Every *homonym* that the language allows for may send the listener to the wrong signified. This is what the semantic literature calls *lexical ambiguity*. The fact of ambiguity (lexical and other) reflects the fact that speakers and listeners are not simply engaged in parallel-but-inverse tasks of coding and decoding. Ambiguity is always a problem of comprehension: Speakers are never confronted with homonyms, for example, because they always proceed from a selected signified to its associated signifier. The problem of *synonyms*, on the other hand, is always a problem of production: Listeners are never confronted with the need to choose between two signifieds of very close meanings, associated with two different signifiers.

The inherent difficulties and necessary capacities involved in production and comprehension are essentially different. This is why the conventions of language emerge from a constant struggle between speakers and listeners. Speakers have to make sure that their language allows them to express themselves. Listeners have to make sure that they understand.

In the third stage of comprehension, RECONSTRUCT, the listener faces the challenge that lies at the very heart of what we usually think of as semantic analysis: The challenge of the re-organization of the by-now copied signifieds *on the generalized map*, and the reconstruction of the speaker's complete message in its original, hierarchical form. This is where the entire set of message construction conventions comes into play, together with the conventions of linearization (in languages where they are relevant). This is the stage in which *structural ambiguity* becomes a constant concern: Listeners face the danger of re-associating the signifieds in ways that are licensed by the generalized map, but do not correspond to the speaker's original message. Many of the conventions developed in

different languages, around what is often called the syntax-semantics interface, are there to minimize this risk.

Crucially, the end result of RECONSTRUCT – a mental representation, in the mind of the listener, of the speaker's complete message – is only a temporary stage on the way towards understanding the speaker's intent. Having reconstructed the complete message, the listener now has an ordered list of instructions, formulated in terms of the internal logic of the symbolic landscape, that he or she now has to use in order to imagine an experiential interpretation. Linguistic comprehension is made possible by the conventions of language, but it is also, at the very same time, totally dependent on the ability of the listener to go beyond these conventions, to figure out what it was that the speaker had to digitize, stereotype, socialize, in order to be able to communicate in the first place. All comprehension is eventually pragmatic.

4.4 Semantic analysis only provides the listener with a scaffold for experiential imagination. Interpretation never stops there. Implicature is the rule, not the exception. In literal interpretation, the process of implicature simply requires less cognitive effort.

In the last stage, RECOVER, the listener has to imagine an experiential interpretation of the speaker's utterance - on the basis of the reconstructed message and everything else in his or her experience. The listener must re-connect the signifieds of the reconstructed message and their semantic relationships to their experiential clusters, re-analogize what was originally digitalized by the speaker, re-contextualize what was de-contextualized, and bring back to the private domain of experience what was originally translated into the language of social agreement.

This implies a major re-thinking of the essence of pragmatic interpretation, and its role in general linguistic theory. Pragmatic theories of implicature emerged in the second half of the 20th century in an environment in which semantics was still thought of as the only necessary tool for the analysis of meaning. The theories tried to show, quite modestly, that in some particular cases semantic analysis was not enough, and that in other, more extreme cases, it actually produced the wrong results. Pragmatic analysis was needed only in those instances where it seemed that the speaker's intent was not fully reflected in

semantic interpretation of the utterance. The notion of implicature, as it was formulated by Grice, was thus constructed not as a general concept of interpretation, but as a supplementary concept, something that needed to be introduced into the analysis of meaning only in cases where semantics was not enough – after the listener decided that the speaker could not have meant what the utterance literally said. Sperber and Wilson's Relevance Theory marked a dramatic departure from this conception: The definition of Relevance as the ratio between contextual effect and processing effort implies that listeners have to calculate the relevance of each and every utterance they are provided with. Relevance Theory, then, goes beyond the illusion created by language itself – and thus reflected in the theory of semantic interpretation – that the symbolic landscape of a language, together with its communication conventions, maps onto the experiential worlds of the speakers in a perfect manner, and thus, in effect, obliterates the experiential gaps between them: As long as the speaker follows the conventions of the language to the letter, the semantic content of the speaker's message is all the listener needs in order to fully understand what the speaker attempted to say.

This is an illusion not because Semantics cannot do *all* the work, but because it does not do *any* of the work. The complete message is nothing but a scaffold for experiential interpretation. (In this sense, Relevance Theory, like all other Pragmatic theories, including Levinson's, still accepts some of the illusion.) What this means is that the distinctions between the different modes of interpretation described in the literature – from literal interpretation, through conventional implicatures and explicatures (or implicitures), all the way to genuine conversational implicatures (involved, for example, in the understanding of irony and novel metaphors) – are determined by the extent to which the conventions of language save cognitive effort for the listener. This conception solves a host of complex problems in Pragmatics, that I will not discuss here.¹⁸

¹⁸ For some of the complexities involved see Grice (1967), Sperber and Wilson (1986), Levinson (2000), Carston (2004) and Bach (2006).

5. Syntax

5.1 In the above description of the process of linguistic communication, no reference has been made to a level of hierarchical syntax. Hierarchical relationships emerge in the process of message construction: They are semantic. Syntax as such is linear. The syntactic tree is an epiphenomenon.

It seems important, in retrospect, to understand that Generative Grammar was never an antithesis to American Structuralism – but rather a direct continuation. In his earlier writings, most importantly in *Syntactic Structures*, Chomsky attempted to take structural analysis a step beyond where it was in the analyses of Bloomfield, Zelig Harris and their colleagues. Chomsky parted ways with his teachers in the sense that he characterized language as a system that allowed for the generation of an infinite number of sentences from a finite set of primitives. This meant that an explanatory theory of syntax could not be deduced from any finite set of sentences: Linguistic theory should go beyond the analysis of corpora, and search for *a-priori principles* which could explain the infinite generative capacity of language (and, eventually, the fact that children are capable of acquiring it). As far as the claim concerning the autonomy of structure from meaning, however, Chomsky remained in complete agreement with his predecessors. The idea that distributional patterns should be explained on a purely formal basis was already there, at the center of the Structuralist enterprise. What Chomsky was looking for was a new formal approach to the old problem of autonomous structure. As he shows in *Syntactic Structures*, the first natural candidate for a formal explanation of this type would be a linear theory, a Generative Grammar that works, so to speak, from left to right. If we manage to account, for every word as it appears in every sentence, which types of words can or cannot precede or follow it, we would indeed capture the desired phenomena. This was the strategy adopted and developed by the American Structuralists, and as Chomsky showed, it could not produce the results. There had to be something that determined distribution in a *hierarchical* fashion, transcending simple considerations of linear order – and what Chomsky offered was a *top-down* generative mechanism, the branching tree structure. The tree structure was designed as a technical solution to a problem that was stated in purely formal terms, and in this sense, it was actually a major step forward. The American Structuralists attempted to force a linear logic on a system whose logic was

both linear *and* hierarchical. Realizing that, Chomsky developed a mechanism that was capable of generating representations, whose building blocks were related to each other both linearly and hierarchically.

The problem, however, is this: The hierarchical relationships that Chomsky acknowledged and represented on the tree are semantic relationships – not structural ones. They are the semantic relationships between *the signifieds of the complete message*, as they gradually accumulated in the process of message construction. The linear relations, on the other hand, are indeed structural, just as the American Structuralists thought: They are the linear orderings of the *signifiers* of the *utterance*, and they are determined by the linearization conventions of the language. These are two different types of relationships, between two types of entities. One is a semantic relationship between signifieds, the other is a linear relationship between signifiers. The syntactic tree, which attempts to conflate them into a single, constitutive level of representation, is thus a category mistake – an *epiphenomenon*.

Consider, if you will, the sentence you are reading right now. Ask yourself: Where is the tree structure of this sentence? The graphic representations of the signifiers you see on the page in front of you are ordered linearly – not hierarchically. You read the sentence, launch a process of comprehension, identify the signifiers, replace them with their respective signifieds in your mind, and arrange them hierarchically in terms of their semantic relationships. The signifieds are now arranged hierarchically, but the linear relationship is no longer there. Look back at the page: The signifiers are again arranged linearly with respect to each other – but now the hierarchy is gone. Again: The issue is not just that the two types of relationships appear, and then disappear, at different stages of production and comprehension. The point is that they are relationships between two different types of entities.

5.2 The two sides to the debate over the autonomy of syntax, the Generativists and the Functionalists-Cognitivists, have been thinking in terms of a relationship between syntax and a single level of meaning. The realization that two levels of meaning are involved changes the very formulation of the question. In terms of the original debate, Chomsky has been right all along: The syntactic structures of language are indeed divorced from the complexities of non-linguistic, private-

experiential meaning. But the structures are correlated, and very intimately so, with the complexities of normative-linguistic meaning. All this is dictated by the functional logic of the system. The facts of syntactic autonomy should not be explained away: They reflect the autonomy of prescriptive, social meaning from private, experiential meaning.

Noam Chomsky's characterization of language as a non-functional entity was not just another component of his overall, philosophically-inspired conception of language as an innately-given, universal cognitive capacity. At its core, it was founded on a set of basic empirical observations, which seemed to indicate that the structural patterns manifested by languages could not be correlated with patterns of meaning. The argument based on these observations was simple: The products of a system whose essence lies in its function have to ostensibly reflect this function in their structures. If the essence of language lies in the function of the communication of meaning, then the observable patterns of structure it manifests should correlate with patterns of meaning. But they do not. They seem to be *autonomous*.

Most of the debate that emerged around the claim of autonomy thus concentrated on the empirical question of the precise degree of autonomy of the relevant structural patterns from considerations of meaning. Functionalists and Cognitivists kept demonstrating that there were striking resemblances (mostly defined in terms of Prototype Theory) between the relevant structural patterns and general patterns of meaning and communication. Generative grammarians, for their part, kept insisting that all this was not enough: Correlation had to be *exact*. The question came to be: How *striking* should a correspondence be in order for it to be considered *exact*?¹⁹

Everything we have said so far, however, indicates that the entire struggle was misguided from the very beginning. Chomsky's demonstrations of structural autonomy do not imply non-functionality, and they do not need to be explained away in order to salvage the functional essence of language. It is precisely in Chomsky's demonstrations that the

¹⁹ For three perspectives on the debate, see Newmeyer (1998, 2007), Croft (1995, 1991, 2001), and J.M. Anderson (1997, 2006).

essence of language, as a system of instructive communication, is most clearly revealed. This is so, because two different levels of meaning, rather than one, are involved in the relationship with syntax – the level of private, experiential meaning, and the level of socially-prescribed, semantic meaning. The very functional essence of language lies in the fact that it allows for the (partial) translation of meanings of the first type into meanings of the second. All of Chomsky's demonstrations are about the relationship between the patterns of regularity manifested by utterances, and the observable patterns of non-linguistic meaning - in our terms, the patterns of private, experiential meaning. His observations are thus correct: The formal properties of utterances are autonomous from experiential meanings (as indeed they have to be.) Utterances, however, are built *from* messages: Their formal properties are there because they signify the internal semantic relationships of their messages. Chomsky's structural patterns, then, are intimately correlated with the complexities of the level of socially-prescribed meaning. Far from indicating non-functionality, the formal complexities of language are actually the key to the function of language as a *social* system of communication.

Claims of autonomy claims have been made with respect to many different types of structural observations. Appendix A presents a very informal discussion of four representative types: The paradigmatic questions of the definitions of *parts-of-speech* and *grammatical functions*, and the syntagmatic questions of *syntactic selection* and *long-distance dependencies*. The arguments, in a nutshell, are these: (i) Parts-of-speech cannot be defined on a notional basis, as functions from experiential meaning to linguistic structure. They should be defined on a prescriptive basis, as structural reflections of the social-normative categorizations of language. Not: 'If something in the world of experience is a person, a place or an entity, then the word referring to it is a noun', but: 'If a word referring to something in the world of experience is a noun, then the thing in experience that the word refers to is to be thought of, by prescription, as a person, place or entity'. (ii) Grammatical functions are regulative conventions. The subject in English, for example, does not stand for the speaker's intended topic. It forces the speaker to highlight a single topic whether or not such a topic is significant for the speaker's purposes. This is why speakers' judgments concerning the obligatory subject are so different from their actual performance. (iii) Syntagmatic relations within the utterance are determined by the patterns of convergence between the social conventions that speakers follow in the course of production. Habitual patterns of convergence are eventually remembered by speakers

as *constructions*. (iv) Island constraints reflect differences between the sets of conventions regulating the production of sentences of different types. The conventions regulating the production of interrogative sentences, for example, are more restrictive (due to independent, functional reasons) than the ones regulating the production of simple assertive sentences. Because of that, the set of producible interrogatives is always smaller than the corresponding set of assertive sentences.

6. The Sapir-Whorf Hypothesis

6.1 The theory of language as a communication technology allows for a major re-formulation of the Sapir-Whorf Hypothesis – as a question about the influence of the normative meanings of language on the process of private experiencing. As with any other technology, we must learn to experience in specific ways in order to make language work. The extent to which language-specific modes of experiencing influence the process of experiencing in general depends on a long list of parameters: Every individual stands at the center of a vector space of influences, experiential and linguistic, and the actual contribution of language depends on the susceptibility of the individual to these influences, and on the direction, temporal order and strength of the different vectors. The two opposite (and variable) processes - the influence of language on private experiencing, and the experience-based invention of new linguistic forms - connect language and experience in a never-ending cycle of bi-directional influence, always spiraling over the foundational constant of the experiential gap.

A certain consideration has been hitherto left out of our discussion: We have seen how pieces of private experience find their way, through mutual-identification, into the fabric of linguistic meaning, and we have also seen how the constant friction between the normative meanings of language and the private world of experiential meanings determines the functional logic of language and its architecture. But doesn't language also influence the way we experience? And if it does, doesn't that mean that the experiential gaps between the speakers of a particular language are effectively reduced?

Two separate questions are involved. The first is very general: Does language, as a system of instructive communication, allow other people to influence the way we experience? The answer must be positive. Our experiences are influenced by what other people tell us - about their experiences and what ours. We are directed towards certain types of experiences, and we avoid experiences that others have warned us against. We seek advice, exchange opinions, argue about the best plan to adopt. We very often compare our experiences and judge them collectively. Our private experiences are almost always entangled with our perceptions of what others say (or might say) about them (and about us). All this is beyond doubt.

The second question is more specific, and the answer is less obvious: To what extent do we let the semantic categories of our particular language direct us in our private processes of experiencing? According to *Linguistic Determinism*, the view most explicitly formulated in the *Sapir-Whorf Hypothesis*, language fully determines the way we experience. We look at the world *through* its socially-constructed categories. Importantly, the hypothesis was formulated in terms of the relationship between the *structures* of language and the patterns of our *habitual thought*. It was almost immediately accepted by the social sciences, and it played a very important role in the advance of the linguistic turn. It gave social scientists a reason to believe that they could decipher the logic of social thinking, through language. In the Linguistic sciences of the Chomskian era, on the other hand, the idea was rejected with contempt. Steven Pinker called it a conventional absurdity. The rejection was rooted in the universalistic, Fodorian conception of the *language of thought*, but it was also based on the simple fact that the hypothesis was founded on conviction rather than empirical work: Whorf's most famous claim was that Hopi had no tense, but it actually does. For a long time there was no communication between the opposing camps, but in the 1980's, a revolutionary, interdisciplinary research program began to investigate the hypothesis as a scientific question. In the pioneering works of John Lucy, Steven Levinson and Dan Slobin, and then Lera Boroditsky and others, Whorf's deterministic conception was replaced with a weaker notion of *influence*, and ingenious experimental techniques were developed to determine *correlations* between structural patterns in different languages and the habitual patterns of thought manifested by their speakers. The main challenge was to capture the patterns of thought independently, *not* through language, so as to avoid the circularity of Whorf's argument

(the very circularity that still haunts some of the social sciences when they rely on speaking for social understanding.)²⁰

The accumulated results now show that speakers of different languages are indeed sometimes inclined to think in ways that are correlated with their languages. Many of the experiments were designed to show correlations between language and cognitive patterns manifest in non-linguistic tasks. John Lucy and Suzanne Gaskins, for example, presented speakers of English and Yucatec Maya with different sets of three objects, and asked them to choose which two were more similar. Two of the objects were made of the same material substance (wood, plastic and so on), and two had the same shape. Most Maya speakers made their judgment on the basis of substance; most English speakers made their judgment on the basis of shape. As Lucy and Gaskins show, this is correlated with the properties of nouns in both languages. In another set of experiments, Steven Levinson and his colleagues asked speakers of Dutch and Tzeltal (another Mayan language) to pick up physical objects and re-position them, in a certain way, somewhere else. Dutch (like most other languages) regularly refers to the positions of objects with respect to each other in *relative* terms: *To the right of, to the left of, in front of, behind* and so on. Tzeltal, quite remarkably, works only with absolute terms (like our north and south): *Uphill, downhill* and so on. Levinson and his colleagues, then, found that most of the speakers of the two languages moved the objects in space in a way that was correlated with the system of their language.

Characteristically, research projects of this type attempt to make the case for *causality* – for the claim that the non-linguistic judgments and behaviors are shaped by language – by demonstrating that the correlations appear, in much younger children, only when they acquire the relevant properties of their language. Lucy and Gaskin, for example, claim that Yucatec Maya children below the age of seven actually prefer shape to substance, and only change their preference after they master their language's complex nominal system. This is indeed a viable interpretation of the results, but temporal order as such does not necessarily imply causality. It may just as well be the case that the Maya

²⁰ Important references here are Whorf (1956), Pinker (1994), Lucy (1992a,b), Lucy and Gaskins (2001), Levinson (1996), Slobin (1996), Boroditsky (2003), Shen and Gil (to appear), Li and Gleitman (2002), Lakoff and Johnson (1980) and Choi and Bowerman (1991).

children take seven long years to become experientially acquainted with the importance of materials and substances in their physical culture, and only manage to acquire the nominal system when they can already understand it. And there are other possible explanations. The fact that linguistic structures and patterns of habitual thought are sometimes correlated may be enough to refute the universalistic perception that all humans think in the same way, but it does not imply that speakers of different languages think in different ways *because* of their languages. The overall argument for causality, as it is presented in these research projects, is not strong enough.

Additional theoretical weaknesses haunt this new and creative domain of research. Both sides of the equation, the structural patterns of language and the habitual patterns of thought, are too narrowly-defined: Language may affect us not only through its structures, and it may affect much more than just our thoughts. Linguistic structures and patterns of thought, moreover, are entities of very different types: It is not at all clear why (and how) they should be causally related to each other. Most importantly, the question is usually framed as an attempt to separate cognitive patterns which are shaped by language, and are thus variable (and necessarily acquired from external input) from cognitive patterns that are universal (and thus potentially innate). This leaves out everything in the individual's experiential life that serves as non-linguistic input for learning. Children and adults do not either know a-priori or learn from language. They learn from by doing, touching, playing, watching and listening. Most importantly, they learn about social life not only through language, but also through active participation (always *privately* perceived) in social activities. All these weaknesses leave the empirical findings under-theorized. We need a major re-formulation of the hypothesis itself.

To understand the influence of language on the way we experience, we must first move away from the philosophical conception of language as an interpretative scheme. Language is a communication technology, not an epistemic organ. We do not experience through language. We experience *and* use language. The question should be conceptualized in terms of the relationship between the two levels of meaning that make language what it is: Linguistic meaning and experiential meaning, socially-constructed and privately-accumulated, digital and analogue. The question is: How and to what extent does linguistic meaning influence experiential meaning? As opposed to a causal relationship between structure and a single level of meaning, which is hard to

conceptualize, a relationship between two levels of meaning, both mentally represented by the speakers of the language, makes simple causal sense.

We already know that handling the technology of language requires a very specific experiential skill: Speakers must be able to experience-for-speaking, to pay attention to components and properties of their experiences that are required by the language's norms. The categories of language, then, determine the way we experience-for-speaking. The idea was first formulated by Franz Boaz, and it was revived in Dan Slobin's classic article on *thinking-for-speaking*. It acknowledges the functional essence of the patterns of experience that emerge in our minds: They serve activities of different types, and there is a mode of experiencing that is specifically oriented towards the activity of speaking. Slobin maintained that language only forces speakers to think-for-speaking where it presents them with obligatory categories, but this restriction seems to be too strong: Whenever speakers wish to use any of the categories of language in the attempt to describe an experience, they have to experience-for-speaking in terms of that category.²¹

It must be understood, however, that this is a *conditional* statement. Speakers may only use a linguistic category to the extent that they manage to experience-for-speaking it, but this does not mean that they always do. Speakers avoid linguistic constructions they do not understand. Moreover, the very essence of the process of mutual-identification implies that speakers do not necessarily experience the same things when they experience-for-speaking a certain linguistic category. The experiential gap is always there. `

²¹ This is already good enough as an affirmative answer to the original philosophical question: It goes without saying that the mental activity which Western philosophy calls 'thinking' – that mode of experiencing that serves such specifically linguistic activities as argumentation and persuasion, negotiation and explanation, affirmation and refutation – is very heavily influenced by the requirements of language. Humans are only capable of such a mode of experiencing because language evolved into a tool that serves these communicative activities, and children – and adults – only manage to develop the capacity for this type of experiencing to the extent that they participate in them. (In this sense, the theoretical evolution of the original philosophical question was probably heavily biased by the fact that the problem had mainly been investigated by individuals who lived their lives with words: This is a particular case of Bourdieu's *scholastic bias*.)

The fact that experiencing-for-speaking is just one mode of experiencing, functionally-dedicated to the usage of language, has recently been demonstrated by Psycholinguist Yeshayahu Shen and Linguist David Gil in their *hybrid perception* project. Shen and Gil presented speakers of different languages with drawings of imaginary hybrids they had never seen before. The hybrids were made, half and half, from images of humans, animals, plants and inanimate objects (like a hammer and a gun) – the four stages of the *animacy hierarchy*. In one experimental setting, the subjects were asked to *describe* the hybrids. In another setting, the subjects were asked to make non-linguistic judgments about the hybrids - allocate the drawings, for example, to different sets. A statistically-significant majority of the subjects took the animacy hierarchy into account in their descriptions. They described one of the hybrids, for example, as 'a person with an upper body of a fish', not as 'a fish with human legs'. This, however, was not how they behaved when asked to perform non-linguistic tasks with the drawings. There, the hierarchy did not seem to play a significant role. The remarkable fact about these experiments is that the subjects were *inexperienced* with respect to these objects. What would happen, then, if the subjects kept on experiencing hybrids and talking about them? Would they eventually zoom in on the animacy hierarchy in their non-linguistic categorization as well? This is how the general question of the influence of language on experience should be formulated: Under which conditions, and to what extent, does experiencing-for-speaking influence experiencing-as-such? This is no longer a question that awaits a simple yes-or-no answer. The extent of influence depends on a very long list of variables.

Let us, then, think about the entire experiential world of any individual, at any moment, as a *vector space* of experiential and linguistic influences, pushing and pulling in all directions. As a first approximation, we may say that the accumulated impact of these influences depends on two types of variables – the *experiential profile* of the individual standing at the center of the vector space, and the *strength, direction and temporal order* of the vectors themselves.

The first type of variables has to do with the general susceptibility of the individual to different types of influences. Individuals have different experiential styles. They are different in their ability (and will) to allow language to interfere with their private experiences. Some individuals are more directable than others. The individual's experiential style is related to the extent to which the individual actually spends his or her

time experiencing or speaking. The more he or she looks at the world in order to talk about it, and the more he or she learns about the world from language, the more influential the worldview of language is in the arrangement of his or her experiences. People may also differ in terms of their trust in the carriers of the linguistic worldview, and in the general truthfulness of language. There is an entire causal continuum here that is generally ignored: When speakers are convinced that they speak a sacred tongue, reflecting godly truth, they are deeply affected by their language. When speakers believe that the language they have to use is an artificial construction, enforced for purposes of thought-policing, they very often do learn to experience-for-speaking the language (they have no choice), but in their own world of experience, they nevertheless keep themselves away from the language's worldview. (This was the attitude adopted by many speakers towards the linguistic dictates of the great totalitarian regimes of the 20th century.) Between these two poles, we find all the more complex variations: Children, for example, may sometimes adopt the categorizations of language simply because they trust their parents and teachers, but they may also learn to speak about things in terms of these categories in order to sound normative when talking with adults - without letting the categories interfere with the actual way they experience. And there are probably additional variables: Even the extent to which we reflexively understand that language affects the way we experience, changes something in the way that it does. We may live our lives by metaphors, but the actual effect of new metaphors on our experiences is reduced after we read *Metaphors We Live By*.

Two preliminary conclusions emerge: First, the influence of language on experience must be *variable* – even within the same linguistic community and with respect to the same linguistic category. In this sense, the most important thing to note about the entire set of experimental results accumulated in the field is that they are all statistically-significant results. There are always exceptions. Some speakers are not influenced as much as the others. Second, many of the considerations mentioned above suggest that children are the most natural candidates for deep influence. They are inexperienced and trusting, at their most flexible moment, eager to know, eager to speak – and eager to conform.

The second type of variables determining the overall effect of language on experience has to do with the vectors of influence themselves. The linguistic vector and the experiential ones may be different in *strength*: In extreme cases, the individual may be exposed only

to language, or only to experience, with respect to a particular realm of meaning. In between, some vectors may be stronger or weaker than the others. The vectors may also appear in different *temporal orders*, together or one after the other. And they may arrange themselves with respect to each other in different *configurations*: The linguistic vector may direct the individual towards a way of experiencing he or she is attracted to anyway by experience. Different experiential influences may direct the individual in different directions, one of which is correlated with the categorization of language. Or the linguistic categorization may fully or partially conflict with the experience of the individual.

Let us look at these configurations one at a time. The first is the most straight-forward: The individual has to learn to experience-for-speaking in a way that he or she finds experientially natural, and the two vectors *converge* to produce a strong effect. Language serves two functions in this configuration: First, as just another vector, it strengthens the inclination towards the relevant way of experiencing. Second, and much more importantly, it provides the individual with the sense of *social affirmation* that comes with the experience of mutual identification. This is the best-of-all-worlds scenario. Language and experience emerge in the individual's mind genuinely correlated, and thus relatively easily translatable into each other. Note, however, that in this configuration, the individual does not end up experiencing in a certain way *because* of language. The correlation is there, but the causal effect is secondary: In the absence of language, the individual would zoom in on the same worldview.

Think, for example, about the system of absolute directions in Tzeltal. The notions of *uphill* and *downhill* have real correlates in the physical experience of the speakers. This is what the terrain they live in feels like. It seems reasonable to assume that for many of them, when they were children, the linguistic terms strengthened an experiential inclination that also emerged on its own. Note, however, that this was not necessarily true of all of them. There were probably those who had a harder time reading the terrain by themselves, and these probably benefited more from language: It actually directed them towards a way of experiencing they would not zoom in on their own. And there were probably those who found it difficult to find their way around, with language or without it. So, we see a variable process, with variable functions played by language vis-à-vis different speakers, that eventually results in a statistically-significant pattern of language-

experience correlation in the entire group (a pattern which, as Lila Gleitmann and her colleagues have shown in their re-examination of Levinson's results, is also sensitive to real-time changes in spatial experience).

In the second configuration, different experiential influences direct the child in different directions, and language helps (or forces) the child to choose between them. Consider, for example, Melissa Bowerman and Soonja Choi's comparison of English and Korean speaking children. Whereas English distinguishes between being *on*-something and being *in*-something, Korean distinguishes between maintaining a *tight*-fit and a *loose*-fit with something: In English, a disk is inside its pocket like an apple is inside a bowl, and a cover is on a pot is like a cup is on the table. In Korean, the disk and the cover are in tight-fit, whereas the apple and the cup are loosely-placed with respect to the bowl and the table. Bowerman and Choi show that Korean and English speaking children do indeed pay more attention to the categorization of their own language. What seems to be happening is this: Both ways of categorization are independently supported, to this extent or the other, by experience. Children, as we all know, are fascinated by both. Quite obviously, the fact that language highlights one of them at the expense of the other sends them a very clear message: One way of looking at the world is more important, more relevant than the other: The adults mention it all the time. The children, of course, must learn to experience the first categorical distinction anyway - for speaking - which means that most of them would naturally choose to focus on the worldview suggested by their language. It is there, it is valued, it is more easily communicated, and it mutually-identified. And yet, even here, it is important to understand that the effects would be variable. Some English speaking children would still be most impressed by the aesthetics of tight and loose fit, and the other way around.

In these two configurations, then, language plays an auxiliary role in the experiential processes of some of the individuals involved. It may appear on stage before or after the relevant experience, and the temporal order may affect its overall contribution. If the experience is already there, language may strengthen it and give it social legitimacy. If language arrives first, it may effectively direct the individual towards the experience, but this is far from trivial. A linguistic construction - a word, an expression or a metaphor - cannot direct the individual anywhere if he or she cannot assign an experiential interpretation of some sort to the construction. Pieces of language that the individual

cannot understand have no effect. Because of this, individuals must be already close enough to the relevant way of experiencing for language to direct them there.

The third configuration is more dramatic: The linguistic vector and the relevant experiential vectors (for a certain individual) point in opposite directions. Lera Boroditsky's work on *grammatical gender* is the only research project I am aware of which may be interpreted as an attempt to tackle such opposition. The main fact here is that many languages do not just assign grammatical gender to animate entities, but also to inanimate ones. In these languages, entities such as chairs, knives and apples have gender just like humans, dogs and cats. If the children acquiring these languages have any experiential understanding of the difference between males and females, they simply *must* know that chairs and knives do not really have gender. When they wish to talk about these objects, however, their language forces them to relate to them as if they did. So, does this have an effect on the way they actually perceive the objects? Boroditsky shows that it sometimes does: Speakers of Spanish and German were asked to mention three adjectives that came to their mind to describe different objects, and the sets of adjectives, taken together, amounted to descriptions of characteristically masculine or feminine entities, in correlation with the grammatical gender assigned to these objects in both languages. The word for *key*, for example, is masculine in German and feminine in Spanish - and most of the descriptions of keys provided by speakers of the languages conformed with this distinction. In this case, language seems to have at least a partial effect on experience when experience itself contradicts it. This is a very important finding: Children are forced by language to memorize the gender affiliation of entities which are not experienced as gendered, and they probably find this easier to do, less arbitrary to grasp, by *imagining* that the entities do have gender.

It is crucial to remember, however, that the effect of all this is much more limited than it seems: Speakers of German and Spanish do not end up believing that keys have gender. The tension between both worldviews remains. In the case of inanimate gender, the tension does not seem to have serious communicative, experiential or practical implications. It is just there. In many other cases, the tension positions itself at the very center of our mental and emotional lives. The discrepancy between the way things are experienced, externally and internally, and the way they are described (especially the way

they are *habitually* described) is a constant factor in human life. Clinical psychologists know about this more than anyone else.

Conflicts between linguistic and experiential influences may thus be partially resolved by an experiential shift, and they may simply remain unresolved – a constant obstacle for communication (and for understanding the world). They may, however, be resolved in another way: Language and experience may send opposing messages to the individual, and he or she may actually let his or her experience shape language. Modern Hebrew, for example, pluralizes the word *woman* with the masculine marker *im*, and the word *father* with the feminine marker *ot*. Children very often refuse to use these plural forms. They re-arrange the morphology, and say *aba'im* and *ishot*. Children, and adults, do this all the time.

What does all this amount to? To begin with, the influence of language on its speakers looks like the influence of any other technology on its users: It depends on the extent to which individuals use the technology, depend on it, feel comfortable and with it and trust it, and it depends on the specific properties of the technology and their functionality *for* the individual.) Other things being equal, those who feel more comfortable with the worldview of language will find it easier to communicate. Second, in order for a linguistic category to participate in the shaping of the individual's worldview, the individual has to be able to identify it with *something* at the relevant experiential level. So, language may strengthen an experiential tendency, and it may help the individual choose between different experiential tendencies, but when the relevant category does not connect to anything in the individual's experiential world. No speaker ever uses everything that language has to offer. Third, even where a linguistic category does influence the experiential worldviews of some individuals, there is no real reason to assume that it influences all of them in the same way. Different speakers assign it with different experiential interpretations.

Last, and most important, language and experience do not maintain a unidirectional, linear relationship between them. The relationship is bi-directional, dialectic, highly dynamic and cyclical. At any given moment in the life of a linguistic community, different individuals and groups of individuals find themselves occupying different positions in the cycle: There are always those who find that language, as it is, lacks

something they need in order to express their intents – or in order to understand their interlocutors. They invent new ways of speaking: A word, a phrase, a metaphor, a politeness marker, a pattern of argument. The inventors are not necessarily aware of their special status: They are driven by the need to communicate. Other speakers may adopt the inventions, because they echo their own experiences, or because they wish to sound as if they do, or because they reveal something they did not see before, or simply because they have to. Social power plays a crucial role in all this. Yet others may reject the inventions or ignore them, because they go against their own experiences, or because they wish to distance themselves from such experiences, or from the inventions, or simply because they do not understand them. In the process of negotiation that ensues, some of the inventions may be gradually accepted into some version of the language, and begin to influence more and more individuals – always to variable degrees and in variable ways. As inventions come to be stabilized as part of the obligatory arsenal of linguistic norms (very often in a form that the original inventor would not recognize), children have no choice but to learn to experience-for-speaking them. Some of the children, depending on their susceptibility and positioning on the vector of influence, would find themselves being directed towards the relevant ways of experiencing – with the necessary level of variability between their experiential interpretations. By that time, new inventions will already be on their way into the language, and new processes of negotiation will be launched.

7. Universality and Variability

7.1 The theory of language as a communication technology offers a new perspective on the universality of language as a human phenomenon: Language is a universal fact not in the sense that all human minds have language, but in the sense that all human *societies* have it. Universality is a social fact, not a cognitive one.

This is not the replacement of one meta-theoretical presupposition with another. It is, quite simply, the replacement of an invalid empirical observation with a valid one. Not all human minds have language, but all societies do. There are many human minds without language: The minds of all human children before they acquire language; autistic minds (not all of them); aphasic minds; the minds of people with a wide range of brain injuries

and affective pathologies; and most importantly – the minds of human individuals who have not been exposed to the social activity of language at the right age. All these tell us something important about the limits of the universality of language as a cognitive capacity. There is, however, not a single human society without language. This is an *absolute universal*. We should start from there.

7.2 The properties that are shared by all human languages are the constitutive determinants of its functional specificity as a technology of instructive communication.

The universal characterization of language as a social system of instructive communication is definitional in the most concrete sense of the word: We identify a certain system of communication as a language only if it manifests the constitutive properties that allow for instructive communication. Everything that is absolutely universal about all human languages is thus a constitutive, functional property of the system as such: All languages are socially-negotiated systems of communication, comprising a large symbolic landscape with triple-natured signs, and a protocol of prescribed procedures for the production and interpretation of utterances. All utterances, in all languages, reflect in some way the compositional and recursive nature of the process of message construction, and the formal nature of the process of utterance construction. All languages reduce experiential complexities into sets of semantic stereotypes, and necessitate a long process of cognitive internalization. All this is very general, but also, at the same time, very particular: All the languages of the world are technologies of the very same type.

7.3 The social-functional universality of language implies that empirical research in Linguistics should always be founded on the default assumption of absolute variability. Languages cannot be understood unless they are analyzed on their own, in their own terms, one at a time – and only then compared to each other.

The point here is not just the fact, recently widely acknowledged in the field, that languages are extremely different from each other, and not even the fact that each and

every language that we know is a closed book for the great majority of humans.²² The point is that things could not be otherwise. Languages are technologies developed by communities, within communities, for *themselves*. This is the essence of language as a communication technology, the heart of its functional specificity.

Every particular language (every dialect, every variation) is a unique system of instructive communication. Every language embodies, in each and every one of its conventions, a totally unique social history of negotiation and struggle, in a unique community, with unique individuals, at different times and places, under uniquely different social conditions. Every component of every language - every word, every semantic category, every behavioral prescription - emerged and came to be established as part of the effort to allow for the communication of different types of experiences, across uniquely different types of experiential gaps. Every language reflects a different set of communicative interests and compromises, a different set of realities and perceptions, a different history of innovation. Every language is an entity *sui generis*.

In terms of the distinction between *core* and *periphery*: Language is a complex system, and because of that, the observable facts of language are determined by a multiplicity of causal factors. Some of these factors, however, are more crucial, more essential to the system as a whole. The others are secondary, contingent. Linguistic analysis is thus bound to lead the researcher, working on a particular language, towards a tentative distinction between *core* and *periphery* phenomena. Certain facts about the language are bound to assert themselves as reflecting something that is essential to the entire system; others are bound to stay mute. There is no way to avoid the distinction, and there is no need to: It is exactly what we need in order to understand the internal order of that multiplicity of causal factors which eventually determines the way a particular language is.

²² In a very influential BBS article, Evans and Levinson (2008) thoroughly deconstruct the myth of language universals, and show, on the basis of typological research, that languages "vary radically in sound, meaning, and syntactic organization". Importantly, Evans and Levinson still conceptualize about the radical variability of languages in terms of the cognitive essence of language. For different perspectives on the question, see the reactions to the BBS article, and also Comrie (1981), Croft (2002), and Christiansen, Collins and Edelman (2009).

The assumption of universality, however, imposes on the distinction between core and periphery a notion totally foreign to the essence of language: If the universal human capacity for language is the essential causal factor determining language - every language - then the question of core and periphery becomes an *inter-linguistic* question. Core and periphery in a particular language are determined by the extent to which they share architectural properties with other languages. Those properties that the language shares with the others are thought of as core phenomena: They are determined, as in any other language, by the universal capacity. Unique architectural properties, developed by the language 'on its own', thus require additional explanation: In the generative literature they are usually demoted to the status of contingent, peripheral phenomena, or explained away as surface phenomena that hide a deeper universal reality.

When variability comes to be equated with contingency, linguistic theory loses its very foundation. Variability is not a contingent fact about language: It is its very essence. The distinction between core and periphery cannot be inter-linguistic. It must be language-internal. There is a certain logic to every language: An emergent complexity of a very unique type, a certain implicit order in the way signs are arranged on the symbolic landscape; a certain understanding of the essence of convention; a different etiquette of communication; a different regime of prescription; a particular level of tolerance and capacity for change; a very particular division of labor between semantics and pragmatics; a unique pattern of social stratification; a unique set of relationships with other languages; and, of course, different signifiers connected to different signifieds and different experiential clusters; different conventions of message construction; different word orders; different morphologies, and so on. What turns out to be a core phenomenon in one language may be a peripheral phenomenon in another, and totally absent from a third. A phenomenon that appears in many languages may be a core phenomenon in some and peripheral in others. Whether or not a certain phenomenon in a certain language also appears in other languages is perfectly immaterial to the explanation of that phenomenon in that language. Languages cannot be variable only at their peripheries. They must also be variable at their cores.

Everything that is essential in every particular language (along with everything that is contingent about it) emerged and came to be essential (or contingent) through a very particular social history, located in space and time, in which different communicative

problems forced members of different societies to find new and different ways to communicate new and different experiences. The social histories of the languages are different, and the languages themselves must be at least as different.

7.4 As the typological literature shows, there are patterns of similarity between specific properties of different languages: These *restricted universals*, however, are not foundational to language. They are observable patterns reflecting emergent commonalities between languages, points where different linguistic communities found similar solutions to similar problems.

There is a fundamental difference between universalistic theories and variability-based theories: Universalistic theories, by their very nature, have no real way of handling variability: Variable observations are either treated as peripheral phenomena; explained as surface manifestations of deeply-hidden universality; captured as variable parameters of a universal scheme; or simply ignored. A theory of variability, on the other hand, has the very same respect for universal (or near-universal) observations that it has for highly-variable ones. If all the members of a set are different from each other, then they cannot all be different to the *same degree*: If that were the case, the set would be governed by a universal generalization. Because of that, the members of the set must also manifest patterns of partial similarity between them, and among those, there may very well be some properties that are shared by most or all of the members of the set: 'All languages have property (p)' is a private case of 'some languages have property (p)'.

7.5 The attested variability between languages does not hide a deeper level (external to language) of functional universals. Notions such as *iconicity*, *information flow*, *topic* and *focus* and so on are themselves the results of the social negotiation of meaning. They must be much more variable than we have usually assumed.

The functionalist perspective assumes that a large class of fundamental linguistic phenomena results from the adaptation of grammatical structure to the function of language (as the general-purpose communication device). Where Chomsky imagines a formal universal grammar underlying all the languages of the world, functionalism imagines an underlying universal set of meanings and communicative functions. There

are culture-dependent meanings and functions, of course, but they reside at the periphery. The core is universal.

The emerging consensus among Functionalists is this: There is a set of functional determinants of linguistic structures. They include, among others, such notions as *iconicity* and *dominance*, *semantic hierarchies*, *information flow*, *topic* and *focus*, and so on. They are all universal and external to language. They are properties of the universal human capacity of communication. These functional universals (also termed *external motivations*) determine the way linguistic structures develop. For most Functionalists today, this is mainly a diachronic process: The relevant universals influence the historical process of *grammaticalization*. The universal functional determinants, however, do not all influence the structures of languages in the same way. They produce potentially conflicting results, and because of that, they compete. In different languages, certain determinants have a greater influence than others, and the end result is a web of restricted universals, partial similarities between large subsets of the world's languages.²³

Here is a specific example: Functionalism uses the notion of *iconicity* to refer to the idea that structures are the way they are because they *resemble* their meanings. Talmy Givon explains the logic of the notion in the following way: 'All other things being equal, a coded experience is easier to store, retrieve and communicate if the code is maximally isomorphic to the experience'. Think, then, about the following iconicity principles: (i) old information appears first in the utterance; new information appears last; and (ii) ideas that are conceptually close to each other are placed together, or close to each other, in the utterance. It is easy to see that the two principles would arrange the elements of the utterance (as a first approximation, the words) in two different ways. The first principle would linearize the words by novelty, the second by proximity. The two principles, then, are in competition (and there are additional principles in competition with both of them). In different subsets of the world's languages, different principles are more powerful than others, and the result is the variety of word orders that there are. All this makes much more sense than Chomsky's Universal Grammar. It is a theory that respects variation, a theory that attempts to explain it rather than explain it away. There is also no doubt that

²³ See Croft (2002), Givon (2002) and Bybee (2006). On grammaticalization, see Traugott and Heine (1991), Hopper and Traugott (2003) and also Deutscher (2005).

the relevant functional categories – including the iconicity of novelty and the iconicity of proximity – have been important in the lives of languages. There is, however, no reason to assume that they have ever been universal.

The attested variability between languages is not a reflection of a deeper level of experiential universality. Behind the structural variability that is observed on the surface there is nothing but a deeper level of experiential variability. The ways languages eventually carry meanings - the meanings themselves, the regulation of information flow, the semantic hierarchies in the symbolic landscapes - all these are social compromises, emergent patterns of prescriptive order that are there in the first place only because the very foundation of language lies in the variability of experience. The two types of iconicity we have discussed above have influenced the structures of languages not because all human beings share a natural tendency to begin by talking about what is already known and then proceed to the news, and not because we all share a natural tendency to arrange our sentences by conceptual proximity – but precisely because we do *not*. Look at children, especially when they are excited: They begin with the news; they jump back and forth between the different facets of the experience they are trying to express. They are not interested in iconicity. They are trying to communicate. The two types of iconicity are prescriptive principles. They are demands made by the community and addressed to the speakers: 'Could you please tell me about the first thing first, and then the other?'; 'Can you please refer to the properties of the referent you talk about together with the referent itself, so we can tell they belong together?' The principles are there (in those languages that have them) because some innovative individuals realized, at certain points in the evolution of languages, that communication would become more effective if speakers could be made to arrange the elements of their utterances in ways that somehow reflect the order of their experiences. These innovative individuals were probably more sensitive to iconicity than others, more attuned to the problems of unordered communication: They may have experienced more difficulty than others in understanding what was being said, and they may have found themselves looking for ways to express something about the internal order of their own experiences, because this mattered to *them*.

Either way, their innovations launched different social processes in different languages, in which they had to struggle, among other things, with other prescriptive innovations that

were trying to pull the languages in other directions. There is always competition, and in this sense the Functionalist perspective is on the right track. But the competition is always language-internal. The complex web of partial similarities between the iconicity-related patterns in different languages does not result from a competition between universals; it results from a competition between *variables*.

There is not even a good reason to assume that the very notion of iconicity was developed in the same way everywhere. The distinction between new and old information, for example, has probably taken different shapes and forms in different communities throughout history: It is not obvious. We are not born with it. It encapsulates an entire socially-developed philosophy. Meaning and communication are just as variable as the structures of languages. There would be no reason for the degree of attested structural variability had it not reflected a parallel level of variability at the foundational level of the social negotiation and struggle over meanings and their communication.

7.6 The general object of linguistic inquiry cannot be the human cognitive potential for language. It must be the human collective capacity for linguistic creativity. Instead of looking for the common denominators of all possible languages (including those that have never existed), we have to discover the actual range of variability between the languages that there are. (Sign languages are extremely important here.)

The question is sometimes formulated this way: Could we all of a sudden find a new human language that is simply different from all the languages we know, in all imaginable ways? The answer must be a resounding Yes. We would only call the new entity a language to the extent that it showed the fundamental architectural properties of a system of instructive communication – but beyond that, the only limit is that of collective creativity. There is a good chance that the speakers of the new language would have found some similar solutions to the ones found by most other language communities. But there is no necessity. There is a functional task that all languages perform: They all allow for the communication of experiences, translated into messages, across the experiential gap. How this task is performed in the next language to be discovered, how it is regulated, what criteria of communicative success the community there sets for itself, what types of tools it has developed throughout the years – all these are totally open questions.

Once we interpret the question of the human potential for language as a question about the human potential for linguistic creativity, we immediately turn it into a question about variability – not about the universal. The potential does not manifest itself in those solutions found and implemented by the majority of linguistic communities around the world. These were probably the easiest solutions (and it would be interesting to ask why). The human potential for linguistic creativity manifests itself first and foremost in the range of differences between all those languages that reside, from the point of view of the universalist paradigm, at the periphery - the range of differences between exceptional languages.

In all this, we should pay special attention to sign languages. The presupposition of universality came to be established as a dominant notion in linguistics years before sign languages were recognized as a topic of linguistic interest. Less than fifty years ago, they were still thought of as forms of pantomime. The idea that full-fledged languages may take shape in a modality other than sound was theoretically unimaginable. Early typological reports thus included such totally obvious absolute universals as: All languages have vowels. The devastating understanding that there are actually many languages without vowels should raise the suspicion that languages are much more variable than we thought, that there is, in other words, something fundamentally wrong with the presupposition. There is nothing to be gained, apart from the salvation of the presupposition itself, from the attempt to show that sign languages do have vowels after all, with the only difference being that they are pronounced by the hands.²⁴ The superimposition of the categories of spoken languages on sign languages misses out on the very essence of sign languages. It does not let them speak out in their own terms.

²⁴ See Sandler and Lillo-Martin (2006). The authors provide systematic analyses of the properties of sign languages, and they go about it almost in the order of resemblance to spoken languages. Among other things, they attempt to analyze the building blocks of the utterance – hand shapes, positions, movements and directions – in terms of the phonology and phonetics of spoken languages.

8. Knowledge of Language

8.1 The conception of language as a communication technology releases linguistic knowledge from the shackles of the universalist presupposition: Human individuals are extremely different from each other in their linguistic skills. Like any other technology, language has its experts, its role models, its innovators, its guards, its rebels and its users – active and passive, efficient and inefficient to different degrees. There is no such thing as an idealized human speaker (just as there is no such thing as *Homo Economicus*).

For language to function, it has to be installed in the minds of individual speakers. To be able to communicate their experiences across the experiential gap, individual speakers have to perform the function of translation *inside their minds*. Speakers do not have to know much *about* their language. They do not have to be able to explain it. They have to be able to *do* things with it: To speak, express, ask, interpret, convince, describe and explain – and to do all these in ways that have been mutually-identified with the other speakers of the language. Knowledge of language is practical knowledge. It is a know-how, a skill.²⁵

The realization that knowledge of language is a skill highlights a crucial fact: It is also highly variable. Speakers are extremely different in their linguistic skills. They are different in the ways they manage to express themselves, to translate their experiences into the semantic structures of language; they are different in their ability to understand,

²⁵ Three generations of linguists have already been brought up on this foundational metaphor: Knowing language is like knowing chess, it is not like knowing how to ride a bicycle. Well, this is simply wrong: Knowing language is almost exactly the same as knowing how to ride a bicycle – and the fact of the matter is that the same is true of chess. Language and chess certainly share an element that is absent from cycling – there are *rules* to be learned (constitutive for chess, regulative for language). But learning the rules of chess or language is a very small part of the process of learning to *play* chess, or to *speak* a language (the rules of chess may be learned in half an hour; learning to play the game takes years.) All three activities – playing chess, speaking language and riding a bicycle – require *skill*, and skill can only be acquired through a long process of active *experiencing*.

to analyze the structures of complex sentences, to translate semantics back into experience, to follow discourse; they are different in their ability to persuade, to explain, to spot a contradiction, to tell (and understand) a joke, solve a crossword puzzle, get away with a lie. Some people write novels, litigate in court, or move crowds to action by speaking at a demonstration. Many can read (some of) the novels and enjoy them (*identify* with them), and follow the discussions in court. Others find these tasks very demanding, quite often impossible. Knowledge of language is not equally shared by all its speakers.

This is a fact that we all know from our experiences as speakers, and it has attracted much attention in the psycholinguistic research on language processing, acquisition and learning.²⁶ General theories of language, however, have ignored it for a very long time: It has been assumed that individual differences (observationally indisputable as they are) are peripheral to the cognitive essence of language. They may affect linguistic *performance* and *second* language learning, but nothing beyond that. The essence of linguistic knowledge lies in linguistic *competence*, which comes into play in language *acquisition* – and it is perfectly universal. Everything said here so far indicates that this is untenable. The linguistic skills that different humans possess (the skills, not the capacities) are perfectly identical only *before* acquisition begins. All new-born babies are non-speakers to the very same degree. From that moment on, differences begin to emerge. As time moves on, and as the linguistic experiences of the different children gradually diverge, additional layers of variability gradually emerge. The end result, the overall distribution of linguistic skills across a community of adult speakers, is a complex pattern of variability. A general theory of language must accept – not ignore - this variability as a foundational fact in need of explanation.

8.2 The question of nature and nurture should be re-formulated in terms of the dynamic interaction between two complex levels of *variability*: Different individuals come into interaction with language with variable innate capacities (both general and language-specific), and they interact with it in variable ways.

²⁶ Elizabeth Bates has probably contributed more than anyone else to the development of psycholinguistic research on individual differences. See Bates, Bretherton and Snyder (1988) and Bates, Dale and Thal (1995). See also Gathercoal and Baddeley (1993), Ehrman, Leaver and Oxford (2003), Kovas and colleagues (2005), and Prat, Keller and Just (2007).

Consider a human skill, like chess playing, which has never been claimed to be universal. Here, the question of nature and nurture would be naturally formulated in terms of the dynamic interaction between two complex levels of variability: We would probably start with a detailed description of the experiencing processes of different individuals – all the way from non-players to the great masters. We would be interested in such questions as: At what age did the different players learn the game? From whom? How? Which strategies were they taught? How many hours a day did they dedicate to the game? Then, we would turn our gaze to something which in the case of chess, we would not hesitate to simply call *talent*: A set of variable innate capacities, unevenly distributed among the different players: Capacities involved in learning, remembering, analyzing, planning, inventing and implementing chess moves. We would find that different players, with relatively similar experiencing processes, had nevertheless developed different levels (and types) of skill. Some players would be able to play by the book, as they were taught. Others would gradually show that deeper understanding of the game which eventually produces innovative ways of playing – and winning. As a first approximation, we would find that skill is correlated with some complex cumulative function of experience *and* talent, nurture *and* nature - as variable concepts. We would obviously reject the idea that knowledge of chess is innately-given: The fact that it is not universally and evenly distributed would warn us against such a move. We would, however, say the following: Playing chess requires a certain complex set of skills, and because of that, it exposes a corresponding set of differences in the innately-given capacities that individuals bring with them to the world (and to the game). Our innate dispositions do play a crucial role in the way we learn and play chess – but they do that as emergent variables, not as foundational constants.

This is exactly how we should think about knowledge of language, with three additional, *quantitative* observations – which change nothing in principle (but nevertheless raise new questions, to be discussed further on): First, unlike chess, the technology of language does appear in all human societies. Second, unlike chess, virtually all humans are exposed to the technology at a very early age. Third, the variable set of innate dispositions which are exposed when individuals approach the task of language acquisition (and usage) seems to have a robust *lower limit*. With language, acquiring a basic skill is within the capacity of almost all children (definitely not all of them). We will get back to these observations and

explain them later on. At the moment, the point is only this: Nowhere, in the entire story of linguistic knowledge, is there anything *foundationally* universal. The observable patterns of variability at the level of skill (including those that are near-universal, or even perfectly universal) do not hide a deeper level of cognitive universality, masked somehow by the contingencies of learning. Variability goes all the way down.

8.3 The question of language acquisition is not how individual children acquire their languages. The question is: How do different children, with their different experiential tendencies, different innate capacities, different types and levels of exposure to language, nevertheless manage to acquire a language that are similar-enough for them to be able to communicate with each other (and with the adults around them)? The answer should be: Language acquisition is a collective project, not an individual one. Children do not learn their language on their own. They develop it together with everybody else in their immediate linguistic community, their parents and other adults, but most importantly their peers.

The entire process of the evolution of language is recapitulated wherever there are children around the world - with two major differences: First, children today are already fully prepared (in variable ways) for the challenges involved in the acquisition of full-fledged language. Those early humans who were involved in the first stages of the process were not (see section 9.5). Cognitive evolution, then, is recapitulated as *maturation*. Second, the children develop their language in a context that already offers a language, spoken by the adults around them. The children may thus *pick up* elements from that language and *introduce* them into the one they are developing. The development of language by children does *not* require full exposure to adult language – definitely not of any single child, but also not of a whole group. Just like our ancestors, when children have nothing to import, they invent what they need.²⁷

²⁷ The remarkable fact of the invention of new sign languages around the world (Nicaraguan Sign Language, Al-Sayid Sign Language and others) allows for a rare view into this very same process – in the (almost) total absence of external input. See Senghas and Coppola (2001), Senghas et al. (2004), Senghas et al. (2005), Aronoff et al. (2008), Meir et. al (2010) and Kisch (2008).

None of this implies that the interactions between children and their parents do not play a crucial role in acquisition. Of course they do: Tomasello's seminal *Constructing a Language* is the best exposition of the contribution of these interactions. It is also true that children (at least in the West) sometimes go through the first stages of acquisition before they begin to spend more time with their peers; that a single child, who only ever communicates with adults, would be capable of acquisition; and that children need to communicate not just with their peers, but also with the adult members of their linguistic community. There is, however, a single fact that overrides of this, a fact characteristically acknowledged only in the literature on *second* language learning: The language that children eventually speak is almost always more similar to that of their peers than to that of their parents. Many of the languages that communities of children develop are *very* similar to those of their parents, and this is indeed an indication of the relative stability and homogeneity of their linguistic communities. But in many other cases, probably most of them, children experience a constant tension between the language they develop with their peers and the language they use with their parents, and when they have to choose, they almost always choose the former. This is most dramatically demonstrated in immigrant children. A single child who only ever talks to adults would naturally (if sadly) see the adults as his or her identity group, and would develop his or her language together with them, which would probably result in a very similar variation indeed. But this is the exception, not the rule. Children grow up with other children, and identify themselves, first and foremost, in terms of their relationships with their peers.²⁸ Research on language acquisition should turn its gaze away from the home (especially homes where the adults are themselves linguists) to the kindergartens and parks. The real story takes place there.

Children are *not* engaged in an attempt to acquire the language of the adults around them (the *target language*). The adults' language runs way over their heads: It is uninteresting for them not just because it is too complicated, but mainly because the adults use it to communicate about experiences the children do not understand. The children are interested in the communication of their own experiences – their *childish* experiences –

²⁸ Judith Rich Harris makes the very same claim, with respect to children's emotional and intellectual development, in her book *The Nurture Assumption*.

and because of that, they have to begin to mutually-identify some of *their* experiences with other children. This is a children's project: It takes place in their world. In Western cultures, parents often take the time to actively participate in the project, and when they do, they speak in ways that indicate that they are now mutually-identifying with the child. This is what *motherese* (or *child-directed speech*) is all about. In many cultures, however, parents very rarely do that. They speak around the child and to him or her, but do not position language as such as an object between themselves and their children. The process of mutual-identification takes place among children, in their own world.

Thus, in the ordinary situation where the adult language is already there, the process may be characterized by the following spiraling dynamic: Throughout the entire process, children look for *correlations* between conventional elements of the adult language that they can *isolate* and experiences (or later on, relations between experiences) that they can *identify*. Different children, of course, identify different correlations. Where the parents are willing to participate, the children mutually-identify the correlations with them; where they are not, they do not. Then, the children introduce their findings into their discourse with their friends: They use them for communication with them, and the findings function exactly like the inventions of our ancestors. A multiple process of mutual-identification (and consequently *reciprocal imitation*) is launched in which some children learn correlations they have not yet noticed, from their peers who already have. What is then accepted into the children's language, and what is rejected, is determined by the children's communicative needs (needs that emerge from *their* experiences), their social relationships, the power dynamics involved, and the general level of capacity manifested by the majority of the children. (Suggestions by precocious children, for example, are very often rejected, and precocious children do indeed sometimes find it easier to talk to adults.) As their language gradually develops, and as the children develop the skills that allow them to communicate at a certain level, they manage to understand a bit more of the adult language spoken above their heads. They may now identify and isolate some additional correlations between conventions and experiences that they could not have identified before, and introduce them into their language, and so on – as long as the process goes on. The language of the children is thus gradually constructed from the bits and pieces brought by the different children into their discourse; the levels of communication that the children reach between themselves constantly rise; and the children mature and develop their skills with their increasingly complex language.

For a very long time, the discourse on language acquisition concentrated on the question of nature and nurture: Is it innate capacity or parental input? In *Rethinking Innateness*, Jeffrey Elman, Annete Karmiloff-Smith, Elizabeth Bates and their colleagues tore the question to pieces: It must be *both*. Their notion of innateness was revolutionarily different from Chomsky's, but their conception of nurture remained at the level of parent-child interaction. What this implied was that the task of acquisition remained an individual one. Every child had to figure out the entirety of language on his or her own. The conception of language acquisition as a collective effort, on the other hand, implies a dramatic reduction in the magnitude of the challenge that every particular child has to face. Children must have to learn to use a system of conventions for the communication of their experiences. The conventions, however, never present themselves *as such* in the speech of the adults. Every adult utterance is a product of a long process of production governed by a complex cluster of conventions, one on top of the other. The real challenge is the isolation of conventions, but then the question comes up: Where do you even start? Well, different children start at different points - those which are more natural to *them*. The first words that children pick up are highly variable; some of them are more fascinated by the sounds of the words, others by their meanings; some children are attracted to logic from a very early age; others relish metaphors. Having isolated a piece of language (by themselves, or with the help of their parents), the children have done most of the work for all the others. The convention now appears in the very simple speech of the children, either on its own or surrounded by those items which have already been mutually-identified by the children before: So much easier to learn.

At the bottom line, the nature-nurture debate has always been about a single question: Is it reasonable to assume that a child can actually acquire the entirety of a language on the sole basis of external input? The conception of acquisition as a collective effort, propelled by the variable capacities of different children, ends up siding with Chomsky: A child does not acquire language on the basis of external input. Groups of children do. Language was invented collectively, and it is still acquired in the same way.

9. Evolution

9.1 In the last two decades, an entire field of inquiry has developed around the question of the evolution of language. This question is the most important bottleneck that any general theory of language should be able to squeeze through: A theory should be thought of as explanatory to the extent that it allows for the construction of an explicit, detailed and multi-dimensional speculative narrative of the evolution of language *as it is analyzed by the theory* – a narrative that meets the requirements of current evolutionary theory.

In the final account, this is the deepest paradox of the universalist program: If language is founded on a universal human capacity – internal or external to language - then we should have somehow evolved to get there; but if the capacity is universal, equally spread across the entire human population, there is nothing that a theory of evolution could say about it. For this, only Evolutionary Theory should be blamed. The idea of perfect universality does make sense in other worldviews – such as the one which has the gift of language bestowed upon all human souls by the powers above. But the replacement of creationism with the evolutionary perspective carries a certain implication that cannot be ignored: The complex patterns of variability that we see in this world came to be what they are as a result of a long and dynamic process, in which patterns of variation changed across time. This is what evolution is all about.²⁹

²⁹ Chomsky himself, in the first four decades after *Syntactic Structures*, consistently refused to discuss the question of the evolution of language. It is a mystery, he used to say, not a scientific problem. Some of his followers (see Pinker and Bloom (1990), Jackendoff (1999), and Piatelli-Palmarini (1989)) have claimed over the years that there is common ground to be found between the Generative program and evolutionary theory, but Chomsky was actually right all along: The evolution of a perfectly universal trait in a biological species is indeed a mystery. In the last decade, however, Chomsky seems to have relented. In an article he wrote with two evolutionary biologists (Hauser, Chomsky and Fitch 2002), he offers a tentative conceptual solution to the problem of the evolution of language. The authors distinguish between two types of human faculties for language. The first, FLB, the Faculty of Language in the Broad Sense, includes components that can be found in other species, and are thus not problematic from the evolutionary point of view. The secret of language, however, lies in the FLN, the Faculty of Language in the Narrow Sense, that includes just one component – the capacity for recursion: The potential

There is, of course, something deeply ironic in all this: Without Chomsky's original insistence on the universal innateness of language, the very question of the evolution of language (in its modern form) might not have emerged – and would have definitely not assumed the central position it has in the linguistic discourse of the last two decades. In this period, a vast amount of high-quality work has been done on the question, and virtually all of it – almost regardless of theoretical inclination – has been devoted to the attempt to salvage language from the Chomskian paradox of unevolvable innateness: To find a way to make it (innate or not) *evolvable* again. Most of the more sophisticated new theories share a serious commitment to the logic of evolutionary theory, an interdisciplinary approach to the question of evidence, a deep understanding of multiple causality, a basic suspicion towards the idea that actual pieces of linguistic knowledge are encoded in our genes, a much more sophisticated understanding of learning, and a firm belief in the very idea that Chomsky rejected: That language evolved because it was functional. Some of the theories are also committed to the idea that social evolution played an important role in the evolution of language. At their core, however, almost all these theories are still formulated as attempts to answer Chomsky's own question – which he himself has refused to answer: What happened to the human *mind* (or brain) that eventually allowed it to carry language? This is still a cognitive question, and the mind is still one. Our first goal, then, should be re-thinking the question itself: Re-arrange the theoretical relationship between the question of the evolution of language and the question of the evolution of its speakers.³⁰

capacity for discrete infinity (which is never observed in reality, because certain properties of FLB, such as memory and processing limitations, prevent it from ever materializing). How did this capacity evolve? It "may have evolved for reasons other than language," such as navigation, number quantification, or social relationships. Well, this is not really a solution. Not even a tentative one. There is nothing here but a desperate attempt to keep the universal essence of language (whatever is left of it) in the realm of mystery – away from the domain of evolutionary explanation.

³⁰ Here, again, I will not attempt to review the literature: It is massive. Some important references: Arbib (2003), Bickerton (1981, 1990), Botha (2003), Botha and Knight (2009), Christiansen and Kirby (2003a,b), Corbalis (2003), Deacon (1997), Donald (1991), Dor and Jablonka (2000, 2001, 2010), Dunbar (1998), Fitch (2010), Gardenfors (2004), Hurford (2007), Hauser, Chomsky and Fitch (2002), Jackendoff (1999), Knight (1998, 2007, 2008), Knight, Studdert-Kennedy and Hurford (2000), Larson, Deprez and Yamakido (2010), Liberman (1991, 2007), Pinker and

9.2 The question of the evolution of language is a question about the dynamic interaction between two evolutionary processes: The technological evolution of language – its original invention, development and sophistication, propagation and diversification, and involvement with other social dynamics - and the evolution of human individuals as speakers. Crucially, the entire process was driven from the very beginning by the technological evolution of language: Human individuals evolved into speakers because language forced them to.

The literature on the evolution of language widely acknowledges the co-evolutionary nature of the process, but still considers the process of cognitive evolution to be the primary one: Language could only emerge when human cognitions were ready for it. This may sound trivially true, but actually is not - for two complementary reasons: The first is the fact that the emergence of complex technologies has very little to do with individual human cognition. No-one would try to explain the evolution of the car, or the Internet, in terms of cognitive evolution. Complex technologies are always collective inventions: Many different individuals – very often exceptional individuals, in exceptional circumstances - are involved in their development. And complex technologies always take a long time to develop: They are founded on the collective capacities of many successive generations.

The second reason has to do with the more general relationship between innovation and individual behavior, cognition and genetics. In the last decade or so, many of the most sophisticated new theories in Evolutionary Biology, especially in the domain of evolutionary-developmental theory (*evo-devo*), have been informed by the understanding that behavioral innovation plays a much more important role in the evolution of biological species than has previously been assumed.³¹ The idea, most clearly formulated by evolutionary biologists Eva Jablonka and Marion Lamb in their *Evolution in Four Dimensions*, is this: Consider the individual members of a biological species, fairly reasonably adapted to the conditions of their environment, who all of a sudden find themselves losing ground because the environment has begun to change. This is the

Bloom (1990), Steels (2001), Tomasello (1999, 2008), Tomasello, Carpenter, Call, Behne and Moll (2005), Zlatev (2007), Zlatev, Racine, Sinha and Itkonen (2008).

³¹ See Jablonka and Lamb (2005) and West-Eberhard (2003)

simplest scenario of evolution. Now, in the traditional, gene-centered view of evolution, the individuals must inevitably go on behaving as if the environment has not changed (their arsenal of possible behaviors being genetically fixed), while waiting so to speak for a helpful genetic mutation, and letting natural selection determine their fate. This, however, is not what actually happens. Biological organisms *react* to environmental changes and launch a process of *exploration* in which they try all kinds of behaviors they have not been genetically adapted to before. Terrestrial mammals are not adapted to swimming, but when they find themselves surrounded by water they nevertheless do what they can to keep afloat. When the environment gets colder, animals look for shelter in places they have never entered or even noticed before. The new environmental conditions raise the level of stress, and stress brings about behaviors that are unordered, exploratory, accidental, sometimes even frantic. Most behaviors do not help much, but from time to time an animal stumbles upon a behavior that actually helps it survive, at least for a while. In these cases, the animals survive not because they were genetically prepared for the new circumstances, but because they were capable of behaving outside the confines of their behavioral envelope.

Biological organisms, then, are adapted in different ways to their environments, but way and above these adaptations they are also adapted – in different degrees and ways – to the fact that their environments keep changing. They are capable of innovation. In the biological literature, this capacity is referred to as the capacity for *plasticity*. There is, of course, a genetic foundation for the capacity of plasticity, and different species are capable of different types and levels of innovation. But it is also crucial to understand that this genetic foundation has very little to do with the actual behavioral *products* of the exploration process. Innovative organisms are genetically prepared for the *search* for behaviors that break the specific, genetically selected-for mold of the regular patterns of their lives. The innovations themselves emerge from this search process.

Luckily stumbling upon a successful behavior, however, is only the beginning of a much longer process, in which the new behavior has to be *stabilized* as part of the organism's behavioral arsenal. The organisms have to identify the successful behavior, isolate it from other accidental behaviors which were not helpful, understand something (at the relevant cognitive level) about the causal connection between the behavior and its functional output, learn how and when to initiate it *systematically* – and, eventually, *get used* to it.

Again, different species are different in terms of their ability to stabilize a new behavior. Assuming, then, for the sake of simplicity, that in our scenario the environment changes and then stabilizes again with a new set of conditions, those individual members of the species that would survive the ordeal, regardless of their genetic makeup, would survive on the basis of what they actually managed to achieve with their innovations – everything that they managed to learn and apply in their relationships with the new environment.

The next step in the argument takes things to an entirely different level of complexity. Assume that some of our innovative individuals manage to survive in the new environment - and eventually multiply. Their offspring will now be born into a world in which the new stabilized behaviour is simply there: They would not have to re-invent. They would have to *learn*. None of the offspring would be genetically-adapted to the task: The capacity for the innovative behaviour of their parents was not passed on in their genes. The fact that the behaviour is now already in their world, however, would radically change the way their genes would *express themselves* in the process of their ontogenetic development. The young organisms would have no choice, in the course of their development, but to launch an exploratory process of their own, recruit as many of their genetically-given capacities as possible (capacities that evolved for other purposes), combine them in innovative ways – and attempt to do whatever they can, with the tools that they have, in order to master the behaviour. In the course of this effort, then, a totally new pattern of cognitive and genetic variability between the learners would be exposed - variability in the types and the qualities of the genetically-given capacities they can recruit and combine for the new learning task.

This variability would then assert itself in two complementary ways: First, different learners would eventually adopt (and stabilize) different *strategies* for the learning task. This is so, because they would rely on different capacities, of different qualities, and would thus have no choice but to attack the learning problem from different angles. Second, different learners would eventually master the behaviour to different degrees. Some would find learning relatively easy, many would find it challenging but possible, others would find it very difficult, even impossible. To the extent that the behaviour remains obligatory for survival, these differences would reflect themselves in the ability of the learners to multiply – which means that they would also be reflected in the patterns of genetic distribution in the next generation. Specific combinations of genes, which only

came to be functionally related to each other because they were recruited for the new task, would now be selected for – and the next generation would actually be *more genetically-adapted* to the behaviour (only *more* adapted, never totally). This process is called in the biological literature *genetic accommodation*: Genes accommodating themselves to innovations. Capacities accommodating themselves to behaviours. Not the other way around.

In the simplest scenario, then, the process that leads from change in the external conditions of life to change in the distribution of genes across the species involves necessity, renewed exploration, luck, innovation, stabilization, and then learning, exploration again, recruitment and re-combination, exposure of cognitive and genetic variability, strategy stabilization, and only then, eventually, genetic accommodation. It is a very different process indeed from the one envisioned in the traditional, gene-centred conception of evolution.³²

It is important to understand, then, that in the course of this complex process *new* capacities emerge which are not just re-combinations of capacities that were already there. Behavioural innovation produces cognitive novelty. New behavioural patterns are forced into existence by necessity; they are gradually curved by experience to approximate, and hopefully meet, their specific functional goals; they become objects of learning, and eventually mould capacities in their shape. It is thus not the case that behaviour is based on existing capacity: Capacity actually emerges *from* behaviour. We are never capable of doing something before we begin doing it. We become capable *as we try*. Skill emerges *from* practice, not the other way around. Quite obviously, new capacities are never totally unrelated to their past: Pre-adaptations play an important role in the story. New capacities, however, emerge from the interaction – made possible by

³² Note that none of this implies that the more traditional mechanisms of evolutionary change are no longer there. Mutations, and other molecular changes, such as genetic re-shuffling, still occur. The point is that the innovations have a direct influence on the way the products of these genetic changes end up expressing themselves. They change the general pool of genetic variability, and are selected for, or against, on the basis of their contribution to the effort of invention and stabilization.

plasticity - between old capacities and new necessities, and because of that, they are never mere reflections of their past. They really are new.

As Jablonka and Lamb show, once a process of this type is launched, and as long as certain conditions are met, the emergence of new capacities may lead to further development and refinement of the innovation itself. As noted, the capacities for plasticity manifested by different species (and different members of these species) vary – but they all share a common property: They are all *finite*. The capacity for invention, and for the ensuing learning process, is never completely open-ended. The further an innovation is from the envelope of the adapted behaviours of the individual, the more difficult it would be for the individual to invent and stabilize it. Hence, those individuals of later generations, who have by now adapted themselves, at least partially, to the behavior invented by their ancestors, would now be able – if required by necessity – to invent and stabilize behaviours that were outside the capacities of their ancestors. To the extent that they prove useful, such additional innovations would launch another process of learning, exposure of genetic variability and eventual genetic accommodation, followed by more innovation - as long as necessity is there. Jablonka and Lamb refer to all this as the *assimilate-stretch* dynamic: Innovative behaviors become easy to accomplish because of genetic assimilation, individuals can then stretch their behavioral envelopes by further innovation, then assimilation occurs again, and so on and so forth. The evolutionary paths of the innovation and its users find themselves entangled in co-evolutionary spirals.

9.3 Throughout the evolution of language, the cognitive challenges that speakers had to face evolved together with language itself. As the challenges increased, and as they became more language specific, patterns of variability between different individuals began to be exposed: Some turned out to be better at handling the technology than the others. Eventually, individuals began to be selected for their linguistic capacities. Because of that, modern humans do have innate capacities for language, but these capacities are *derivative, emergent, variable* and *partial* – not constitutive, foundational, universal and complete. Innateness is *a-posteriori*, not *a-priori*. (And it is not only cognitive: The deepest indication that we have evolved for language is not the fact that we can do it – but the fact that we need it: Our minds are language-craving.)

As Eva Jablonka and I show in our model of the co-evolution of language and speakers, pre-linguistic humans must have been cognitively and socially ready for the beginning of the process of exploration that eventually gave birth to the first prototype, or prototypes, of language - but they could not have been language-ready before language came into being. Language was invented *before* its speakers were fully prepared for it. The first prototype, and its subsequent modifications, exposed new patterns of variation between its speakers: In the course of their attempts to use the technology, different individuals recruited different sets of capacities, and managed to meet the challenges of language to varying degrees. The first generations of speakers gradually learned how to use the new technology efficiently, what the problems were and how they were to be solved, and they gradually shaped their linguistic behaviours into increasingly specific patterns. In the active social exploration of the behaviours, specificities came to be conventionalized: Linguistic habits, mutually-identified by members of the community, began to be formed as objects for learning. As the technology advanced and became more difficult to learn and use, and as it gradually established itself as an indispensable element of human life, the variations between speakers began to assert themselves more dramatically: Individuals who could not find a way to join the social activity of language gradually began to fall off track.³³ Later generations were thus more adapted - socially, cognitively and genetically - to the realities of language, and were consequently better suited not just for the tasks of learning and speaking, but also for the task of further innovation. New innovations, and especially the revolutionary ones, brought language to new levels of

³³ The fact that the adaptive value of language registered most clearly, from the very start, at the level of *group selection*, probably served to delay the process of the individual selection: The first groups of speakers (and then the better groups) could do things *together* that the others could not even imagine: Language allowed them to share information, plan and decide collectively, fight differently, construct a totally new level of social cohesion, develop a collective understanding of the world that lay beyond the here-and-now of private and social experiencing – on progressively higher levels of efficiency. All this was beneficial for the entire group, *including* those who found it very hard to cope. There were probably many individuals on the way who failed to meet the cognitive challenge of language, who nevertheless enjoyed some of the benefits of the technology simply because they were born into a group of speakers, and were thus (other things being equal) better protected, better fed and so on. With time, these individuals eventually found themselves losing ground within the group, both as co-operators and as potential mates.

complexity and efficiency, extended its expressive envelope, raised its functional significance in the general fabric of social life, consequently contributing to the further selection of better speakers over others – and so on and so forth, stage by stage, all the way to the language-biased minds of modern humans.

Moreover, the speakers of more advanced languages probably acquired their linguistic skills at a younger age than their ancestors: Not just because language was already there, as a 'natural' element of their social lives, but also because there were specific advantages to earlier acquisition: The young are the least experienced, the most exposed and the most dependent. They actually need language more than anybody else. Individuals were thus selected not just for their general ability to reach a certain level of proficiency, but also for their ability to do so *quickly*.

And it was not just a cognitive matter. The deepest indication that we, modern humans, are already adapted to participate in the social activity of language is not the fact that we can do it – but the fact that we *need* it, that we crave it. Throughout the evolution of language, individuals were not only selected for the capacity to participate in the activity of language – but also for their *will* to do it. Those who were more deeply attracted to the evolving technology, who were more desperate to understand and to talk, who longed for mutual-identification, spent more time exploring language. They spent more time mutually-identifying from an earlier age, and they invested more energy in the acquisition of linguistic skills. Their fascination with language increased their chances of survival, and was thus partially genetically-assimilated. We are born with different levels of this fascination. Our minds are language-*craving*. This is why, as children, we actively seek language. Not because we already know it, but because we *want* it.

We also crave experiential communication, and in this sense we are not very different from many other species. Other animals, like us, can feel lonely without social contact. Only humans, however, hunger for the type of social contact that can only be achieved by mutual-identification. It is important to look at Kanzi the bonobo in this light. Kanzi saw the humans around him trying to teach his mother to use a system of graphic signs, and he managed to recruit enough cognitive capacity to master the language at the level of a three year old. The question of the cognitive differences between Kanzi and human children is thus far from trivial (see section 9.5). The real difference between Kanzi and

human children is the fact that Kanzi very rarely (if ever) *initiated* events of mutual-identification with the human researchers who invented the language. What would happen, then, if Kanzi had not been exposed to the teaching process? Assume that he would simply be surrounded by humans using the same sign system for communication between themselves: There is a very good chance that he would remain uninterested. A two-year old human, however, would be strongly attracted to the sign board.

9.4 Like all other technologies, language evolved from a first prototype, in a long line of gradual modifications and technological revolutions. Throughout the process, its functional envelope gradually expanded, its efficiency improved, new problems were created and solved, and language diversified and spread.

Two general types of questions should be asked about the evolution of language as a technology. The first has to do with the spread of the technology across human space and time. We will never know how many times, and in how many communities, language was originally invented, how many times it was invented, forgotten and re-invented before it stabilized, in some communities, as a regular element of social life. But we may safely assume that just like any other technology, language was invented *somewhere*, in a certain community or in a few communities of early humans – definitely not everywhere at the same time. From a very local beginning, language must have very gradually spread across more and more communities, splitting into ever-changing dialects and varieties, until it reached the universal spread it now has across human societies. In some cases, language may have been imported (or stolen) from another community, as technologies usually are. In other cases, language may have spread together with its speakers, who used their revolutionary technology to defeat their non-speaking rivals in actual wars or in the more mundane struggle over resources. Most of what we would like to know about this process is probably lost forever: Where exactly was language first invented? When? Did it happen just once or a few times (and if it happened more than once, how did the first prototypes of language differ)? How did language spread in different places and times? How long did it take before language reached the moment of universal spread? How long ago did communities of non-speaking humans exist?

The second and much more important question has to do with the actual technological development of language. We may safely assume that the first prototypes of language, the

first versions of the technology that actually worked, were very different languages as we know them. They were much less complex, much less sophisticated, much less efficient: All first prototypes are. A first prototype of any invention makes a difference because it makes certain things possible which could not be done before – very often at a very rudimentary level of success. As the first users begin to work with the prototype, as they start to accumulate their experiences *as users*, they begin to learn things about the interactions between the invention and the environment. They gradually understand more about the capacities of the technology and the ways it should be used, and they discover totally new problems which arise as the technology interacts with its environment. For a long time, and very often unintentionally, different users thus introduce many slight modifications to the prototype (and to the way it is used), which in their turn have a cumulative and gradual effect on the general efficiency of the system. The effect is quantitative: The system remains the same, it just gets better. In some cases, this long line of accumulated improvements eventually leads to the stabilization of the invention in its final form. In other cases, the quantitative process eventually translates into a qualitative effect – a technological *revolution*. The revolution may occur as a direct result of the gradual accumulation of modifications: The system reaches a critical point at which totally new patterns of usage are suddenly possible. It may also occur due to the accumulation of problems: As the system improves, new problems arise, and a new set of modifications, which emerges as an attempt to solve these problems, turns out to open up entirely new functional capacities. Sometimes, the revolution is made possible only when the environment (physical, social and technological) changes in a specific way. This way or the other, the system that eventually emerges should be properly thought of as the first prototype of the *next generation* of the technology. It is qualitatively different in its architecture, and it can do things outside the functional envelope of the first generation – again, at a very rudimentary level of success. Consequently, the system now enters a new phase of gradual evolution, in which new capacities and new problems arise and many slight modifications are introduced. Then another revolution occurs, the next generation of the technology appears, and so on: The process goes on as long as it remains both necessary *and* possible. All the complex technologies that we use today are the products of such evolutionary histories.

Finally, all technological systems, as they evolve, impose more and more *system constraints* on the possible venues of their own future evolution: The system acquires a

certain specificity that gradually makes certain types of changes more difficult to incorporate, and thus in effect participates in directing its own evolution. The characteristics of the early generations determine many of the characteristics of the later ones.

9.5 The key to the origin of language is the fact that it was not the first major human invention: It was preceded by a revolution in material technology, and then by a consequent revolution in experiential communication - the invention of mimesis. Throughout this very long process, groups of humans gradually developed the collective capacity of mutual-identification – totally absent in the apes – which provided the foundation for the invention of language. The constitutive components of language emerged from mimetic communication, and language itself began to be used for genuine instructive communication as a result of a rise in the levels of social trust in human populations.

Most of the theories about the origin of language begin with a negatively-formulated question: Why *don't* the apes have language? The theories thus think in terms of a single-step dynamic leading directly from ape cognition (or culture) to the human capacity for language (or the human capacity for proto-language). This is the common denominator of the hypotheses put forward by Chomsky and his co-authors, Bickerton, Jackendoff, Deacon, Dunbar, Hurford and others. Donald and Tomasello, on the other hand, go a step further and hypothesize a two-step dynamic: Donald proposes a mimetic stage in between; Tomasello postulates that humans had learned to co-operate at a higher level before the emergence of language. In this sense, both authors move away from the negative formulation of the question: It is no longer simply what the apes lack, but what happened to humans *after* they became a separate species. The two hypotheses thus take a significant step away from the deeply-rooted equation of the essence of humanity with the existence of language: We are not the talking ape. We are first and foremost the co-operative ape, or the mimetic ape.

The realization that language is a collectively-invented technology, allows for the further development of this conception:³⁴ Early humans began by further developing the material

³⁴ For the full version of this claim, see Dor (2011).

cultures they shared with the other apes, especially with chimpanzees. They became better tool inventors and makers. The more advanced material inventions implied a more difficult learning task, and humans gradually developed into better individual learners, more efficient *imitators*. From a certain moment on, however, imitation was no longer enough, and the skilled individuals who had to transfer their knowledge to the next generation had to go beyond the facilitation of learning, and develop communication technologies that allowed them to *teach*, to *point out* specific experiences and to direct the attention of the others to them. The new mimetic tools, both physical and vocal, revolutionized experiential communication, solved the problem of the transfer of material skill and allowed for the stabilization of more collective inventions. They changed human cultures in fundamental ways, brought about the invention of *re-presentational* communication (art, ritual and theater), and allowed for revolutionary levels of co-operation between individuals. As mimetic communication evolved, systematic associations between certain experiences pointed at, and certain vocal experiences, came to be established. The associations allowed communicators to direct others to the specific experiences they were interested in, and as such, they were only useful within the here and now of collective experiencing. Once the levels of epistemic trust among some individuals reached a certain level, the associations came to be used for the first time for instructive communication.³⁵

From the vantage point offered by this hypothetical narrative, the negatively-formulated question – why don't the apes have language - seems almost irrelevant. But it is not. This formulation has forced us to take a very close look at the apes and their capacities, and what we now know is that they have what is required at the *cognitive* level for the acquisition and use of a first viable prototype of language. The answer to the question is not a matter of individual cognition. For a long while, it has been assumed that the apes do not have a *theory of mind* – and therefore they have no language. But as Tomasello and his colleagues have shown, it turns out that they do. For a long while, it has been

³⁵ Such levels of trust were required because in instructive communication, as opposed to experiential communication, the receiver has to *believe* the sender: The message cannot be experientially verified. The technology of language thus created a totally new problem – the problem of *truth*. Paradoxically, the rise in the levels of trust allowed speakers to use the new technology for a revolutionary form of deception – the *lie*.

assumed that the apes cannot *learn* language. What Kanzi and his sister Panbanisha managed to learn in Sue Savage-Rambough's laboratory proves that they can. (Note that science very rarely goes beyond demonstration of viability into the actual realm of proof. What Savage-Rambough and her colleagues achieved is tremendous.) It is often claimed - by proponents of Chomskian innateness - that Kanzi's linguistic behavior actually shows that the apes *cannot* acquire human language the way we know it. His capacity does not exceed that of a three-year old. This is true, of course, but it is also totally beside the point. The crucial fact is that the apes in the wild not only lack fully-developed languages – they lack what Kanzi managed to learn in the lab as well. The negatively-formulated question should thus read: Why don't the apes in the wild have language at the level that Kanzi managed to learn? What Kanzi and the other apes, observed in the wild and experimented with in different labs, actually demonstrate is only this: Ape cognition is enough for learning and using a first prototype of language. When an intelligent ape is provided with a language already invented and stabilized by others, the ape is perfectly capable of learning it to a level that allows for efficient instructive communication. The apes in the wild never had a language not because they could not have learned it, but because they could not have *invented it by themselves*. And they could not do this not because there was something lacking in their individual cognitions, but because the invention of a system such as language requires a whole set of social and technological pre-conditions, an entire pre-history of social and technological development that the apes never went through.

The long and complex evolutionary history of pre-linguistic humans brought human *collectivities* to the point where the invention of language became possible. As opposed to individual learning, and to those individual activities associated with a theory of mind (such as the following of another individual's gaze), the human activities that eventually paved the way to language – teaching and directing, being directed and learning from teaching - were all collective activities. They take place between individuals. The level of success that a community of individuals reaches in these activities does not depend on any of the individuals involved as such, and it does not depend on the simple sum of the capacities of all the individuals put together: It depends on the quality of the interaction, the communication technologies involved, and the level of *reciprocity* that is achieved in the process of the mutual identification of experiences and their significations. The capacities that determine all these are collective capacities, and these are precisely the

capacities that the apes are missing – not as individuals, but as collectivities. Looking for the origin of language in individual cognition thus misses out on the very essence of the human species: We are collective animals (sometimes co-operative, very often not). We could only invent language because (and when) our cognitions came to be *de-individuated* to a sufficient degree. The apes never invented language because they have always been far too individualistic for the task.

9.6 From the moment of origin, language gradually demarcated itself from mimetic communication – and asserted itself as a formal tool. The process was constant, gradual and *difficult*, and it probably went through at least two technological revolutions: Following the original invention, language began to assert its *autonomy* on the two poles of the process of communication, with the beginnings of the symbolic landscape on the one hand, and phonetics on the other. The second stage was launched by the revolutionary invention of compositionality, and saw the beginnings of language-internal *structure*, with phonology and morphology. The third revolution was the invention of the *protocol* as a regulative mechanism.

The first prototypes of language probably allowed their speakers to engage in what the acquisition literature calls *holophrastic communication*: Signs referred to whole experiences, and speakers whose communicative intent had an experience of a type already mutually-identified as its object could use the appropriate sign and communicate their intent to their listeners. The impetus for further invention and mutual-identification of additional signs was obvious: Every new sign increased the expressive power of the language. It stands to reason, then, that the first prototypes of language, following their initial stabilization, underwent a long process of gradual increase in the number of their signs. Slowly at the beginning, and probably faster later on, speakers came to spend more (collective) energy on the process of mutual-identification. More and more types of experiences were isolated from the continuum of private experience, and highlighted by social agreement. The holophrastic signs that were multiplying on the symbolic landscape were not yet connected to each other: Each of them marked the mutual-identification of a single isolated experience. Neither were the signs categorically demarcated: They did not yet refer in different ways to what we now think of as experiences of different types – concrete entities, activities, properties and so on. The mutual-identification of isolated categories of experiences had to wait for the future. At the moment, however, this is much

less important than the fact that the symbolic landscape was gradually taking shape. A symbolic worldview was beginning to emerge.

Changes, moreover, were taking place on two more levels: It is not unreasonable to believe that holophrastic *conversations* began to appear and develop. Members of speaking groups started to engage in exchanges of signs, one at a time, in which the interlocutors experimented (not necessarily consciously) with new ways to achieve radically new communicative goals. A system that allows its users, for example, to signal that there is a certain entity, out of sight, in a certain direction, is bound to produce proto-conversations in which the interlocutors express disagreement about the right direction to look at, signal their understanding or confusion, or actually *ask* (for the first time in the evolution of life on the planet): Is it here or there? A system that allows its speakers to direct collective action by defining its goal in advance, as in naming the thing to be looked for, opens the door for attempts to negotiate the goal, and for the beginnings of collective planning. Most importantly, a system whose very function depends on the foundational act of mutual-identification must have dragged its speakers into endless meta-conversations in which signs were uttered again and again, imitated and compared, with and without the pointing, with and without the meeting of the eyes, in the constant attempt to verify their collective value. Mutual identification is still, today, the most difficult challenge involved in the acquisition and learning of language, and its failure (when it fails) is still the foundational source of miscommunication. It would be naïve to assume that the first human speakers found it easier than we do.

And there was another challenge that gradually asserted itself as the new technology developed. Communicative success depended not just on the mutual identification of experiences, but also on the mutual-identification of the *sounds* associated with them. The signs had to be mutually-identified against the background of individual *variability* – vocal and emotional. On the analogue continuum of presentational vocalization (both at the mimetic and the pre-mimetic level), individual variability was a primary tool of expression. It still is: We say much more with our voices, with our tones, than we usually imagine. Language, however, has required speakers and listeners, from the very beginning, to abstract away from all this and learn to produce and identify the *same* sounds across the continuum. A speaker may be old or young, a man or a woman, angry or sleepy, thirsty or not – whatever he or she utters must be recognized by the listeners as

one of the sounds already mutually-identified. We, modern humans, do this with perfect ease, but this is only because we evolved to be able to do it. The first speakers must have found it unbearably difficult. (To get a sense of the difficulty, think of *tone identification* in music. When the same tone - or pitch - is played by different instruments, with different levels of loudness and different colors, the actual sounds that are heard are very different indeed. It is not at all easy to hear that the clarinet's whisper, the violin's cry and the piano's pounding are actually the same tones. Most of us require practice to do it.) All this means is that the early speakers of language must have begun to assemble sets of sounds that could be used, with reasonable confidence, in the process of linguistic communication – and isolate them from the continuum of presentational vocalization. Language began to *sound* differently. *Phonetics* was born.

The same thing, then, was happening on all three levels: Speakers were gradually isolating language from everything else that was already part of their experiential world. Their experiences of linguistic sound were gradually demarcated from their experiences of presentational vocalizations; their experiences of linguistic communication were gradually demarcated from their experiences of presentational communication (mimetic and pre-mimetic); and the socially-constructed worldview of the symbolic landscape was gradually demarcating itself, in their minds, from their worlds of private and collective experiencing. Language was making its first steps towards *autonomy*.

As language slowly developed in all these different directions, the lives of those individuals who were close enough to the new technology began to change – not just because but they were teaching themselves, as Austin has put it, how to do things with words, but also because their entire world was beginning to expand. They were still living in the here-and-now of private experiencing (we still do), but certain elements were gradually penetrating their worlds from the outside: Things they did not experience, but were told about. If there ever was a Copernican revolution in human history, this was it: Almost all the animals on the planet experience; the apes know how to follow the experiences of others in order to learn about the world; pre-linguistic humans learned how to direct the experiences of the others, and let the others direct theirs – but with language, humans could finally begin to experience *for others*, and let others experience *for them*. Because of that, the first speakers gradually saw space and time opening up in completely new ways. Things that happened to other individuals, in their absence, in places they had

not been to, slowly became relevant for their lives: They began to learn how to *imagine* them, and how to take them into account in their own decision-making for the *future*. In their conversations, they started to share judgments, and argue about them, making their first steps towards collective descriptions of the world. In the process, moreover, memories from the *past* gradually turned into new objects of communication. Active remembering-for-speaking – keeping an experience fresh in your mind until you tell others about it - became a useful capacity. As memories began to accumulate, as collective consciousness began to emerge, individuals found themselves increasingly relying on language in their own private process of experiencing: They found themselves going to places they had never been to, looking for things there that the others had told them about. Their experiential worlds expanded. The limits of the here-and-now within which they were living gradually stretched. With new experiences, came more signs, which allowed for more types of experiencing, and so on. Language and experience began to spiral together. All this implied an entire set of enormous and unique challenges: The foundational challenge of belief; the experiential and linguistic challenges of mutual-identification, signification and digitization; the challenges of articulation and identification of phonetic sounds; the social challenges of face and insult, co-operation and turn-taking, trust and deception; and most importantly, the challenges involved in understanding a new world, gradually revealed through language - a world that spread beyond anything they had ever experienced privately.

Of all the gradual processes which the first prototypes of language, their speakers and their communities went through, one development – an almost technical one – would pave the way towards the next generation of the technology: As signs accumulated on the symbolic landscape, some of them began to *overlap*. Certain clusters of experiences came to be associated with more than one sign. Speakers thus found themselves in a new situation: They could use more than a single sign to refer their listeners to the experience they wished to communicate. Assume, for example, that they already had a sign for the experience of *chasing*, which they used together with pointing to order to get the others to start running in a certain direction. And assume that they began to invent signs for clusters of experiences in which they saw different kinds of animals. Speakers who saw a rabbit, and wanted to launch a group chase, all of a sudden found themselves associating two signs with their experiential intent. They had a new choice to make: They could either use one the two signs - the more regular strategy - or they could try to use both, one

after the other. It was a moment that changed language forever: *Compositionality* was born.

The speakers who began to utter more than one sign for a single experiential intent might not have noticed it, but they were actually presenting their listeners with a radically new challenge: The listeners were no longer just required to bring up from their memories clusters of experiences associated with mutually-identified sounds. They were asked to imagine the experiences associated with the sounds, and then *calculate the intersection* between them: To concentrate on chasing-experiences, and on rabbit-experiences, and then realize - the speaker is talking about the experience of rabbit-chasing. The speakers were not thinking in terms of compositionality: They were just using the signs which they associated with their experiential intent. The listeners, for their part, were not thinking in terms of intersections: They invoked the relevant clusters of experiences in their minds, and some of them, some of the time, found themselves imagining them together. All this was far from easy: It must have taken a lot of time and a lot of energy to stabilize. It was revolutionary for many different reasons, but most importantly because it allowed for communication about the intersected cluster of experiences (the cluster of rabbit-chasing) *without* the prior mutual-identification of the cluster itself. The prior mutual-identification of chasing-experiences and rabbit-experiences was enough. Members of the language community could now communicate not just about experiences they had mutually-identified, but also about different combinations of these experiences. (Which also meant that they could communicate about entities they had constructed in their own imaginations, from pieces of experiences. The consequences were tremendous: As I show in Dor (2011a), the full scale of the revolution asserted itself in the *lie*.)

What all this implied was a great leap forward in the expressive power of language: The number of messages made available by language was no longer equal to the number of signs. In principle, speakers could now produce as many utterances as they could compose from the signs at their disposal. Every sign, then, was now worth much more than before. This, in turn, implied much higher dividends on the mutual-identification of new signs. The function from the number of signs to the number of messages, which was up to now a linear one, became exponential. The symbolic landscape was being populated at a much higher rate. Language entered the period known in the acquisition literature as the period of *naming explosion*. (It is important to see the dialectics here: New signs

became much more valuable at the very moment that certain possible candidates for new signs – like a sign specifically dedicated to rabbit-hunting – were no longer needed.)

Not every two signs, of course, could be used together for the communication of the same experience: Not all of the mutually-identified experiences intersected with all the others. Because of that, those signs which were associated with intersecting experiences, and were thus repeatedly used together, came to be associated with each other on the symbolic landscape. *Semantics* was born. Very gradually, the socially-constructed worldview of the symbolic landscape, which up to now included sets of isolated experiences, each mutually-identified on its own, turned into a categorized system: This experience always appears with this one, sometimes with that one, never with a third. The same general experiences could now be *differentiated* into subtypes, looked at from different angles, *specified* in ways that were not possible before. Human communities launched a huge practical science project: The collective project of *classification*. As a result, and very slowly, the signifieds themselves were grouping up into subsets of their own. They were no longer all of the same status. When the sign for *chase* came to be systematically associated with signs for different types of prey, for example, it was already on its way towards becoming a verb. The signs of the different types of prey were on their way towards becoming nouns.

And there was another implication to the emergence of compositionality: Speakers who were trying to use more than one sign per utterance found themselves obliged to pronounce them one after the other. *Linearization* was born. It required some new skills, cognitive and articulatory. As speakers gradually improved their linearization capacities (to different degrees and in different ways), they managed to combine more signs and produce longer utterances, and they managed to do so with greater ease. Speakers began to speak *faster*. This was a pivotal moment: We shall get back to it in the next section. With the rise in speed, the mutually-identified sounds of language began to be pronounced more closely. This was an articulatory challenge of the first degree. It also presented the listeners with a new *auditory* challenge: The real-time challenge of breaking the stream of speech back into its separate signs. As the speed of speech gradually increased, the sounds of language began to arrange themselves in systematic relationships. Combinations of phonemes that were repeatedly used came to assume a certain character: Speakers and listeners gradually became familiar with their sound

patterns. Mutually-identifying these patterns was in the best interest of both speakers and listeners. *Phonology* was born. Phonetics and phonology began to evolve together: Phonemes began to be selected and modified to allow for the swift move from one sound to the next. Newly invented signifiers began to comply with the emerging sound patterns of the languages.

As the first prototypes of the second generation of language took shape, the same thing was happening on both ends of the technology: From the first generation's symbolic landscape and phonetic system, emerged semantic and phonological structures. Language was making its first steps towards *internal complexity*. It was by now a very powerful tool of communication, and it was probably already being used by many more individuals in more human communities, for many more functions. It was gradually asserting itself as a major entity in human life.

Along with success, however, a set of new problems emerged. These could only be solved by the regulation of linguistic communication. Norms of speaking began to be isolated and mutually-identified, prescribed and internalized. The *protocol* was born. Major parts of the protocol were dedicated to the resolution of a variety of social problems accentuated with the advent of language - problems of conduct and status, insult and face, suspicion and deceit. Communities developed implicit collective understandings regarding the appropriate structure of conversation, politeness and respect, communicative co-operation and social power.

At the core of the protocol, emerged sets of rules for the regulation of the actual process of speaking. They were only required because of the increasing internal complexity of language itself. In the life of every invention, a moment comes in which its own internal complexity begins to reduce its overall efficiency; at such a moment, the complexity of the second generation of language, made possible by the invention of compositionality, began to produce problems of *interpretation* of a totally different order. The speakers, skilled as they came to be in the art of composition, gradually produced longer and longer utterances – more and more complex translations of their experiential intents into semantic messages. The listeners, however, who had to re-arrange the signs into a semantic structure that could be re-translated into an experiential representation, found this more and more difficult. As long as utterances were composed of no more than two

or three familiar signs, interpretation was probably still relatively simple. The longer utterances, however, were much more difficult to decipher. They were *ambiguous* - the signs could be re-arranged in different ways to produce different messages (and interpretations) - and they were increasingly more *opaque*: In the holophrastic conversations of the first generation of languages, speakers referred listeners to isolated experiences, each of which had been mutually-identified by both sides. In languages of the second generation, speakers referred listeners to combinations of experiences that had not been mutually-identified as such. The listeners had to imagine a reasonable intersection of the different experiences that would produce a coherent interpretation. Each of the individual signs was still associated with a sufficiently demarcated experiential cluster, but the combined experience had no sharp outlines. It was fuzzy. Listeners must have found themselves misinterpreting speakers at increasingly higher rates. Speakers were getting better, but the general level of communicative success was declining. (It should be remembered that not all individuals played the roles of speakers and listeners to the same extent. Many were mainly listeners, and many of these were precisely those who found it hard to cope with the technology. Compositionality made their task so much harder.)

The new problems of interpretation were aggravated by another factor: As the different communities of speakers grew larger, the levels of experiential variation between individual speakers, as well as subgroups within the community, grew deeper. The task of mutual-identification of meaning across the entire community was gradually becoming more and more difficult. Patterns of linguistic variability within the same community began to emerge, and they contributed to the reduction in the communities' overall levels of interpretative success. All this was an inevitable consequence of the success of language, but it was a serious problem nevertheless.

Innovative speakers, then, began to identify different communicative problems and explore different ways of speaking that could resolve them. This was not a theoretical project: The inventors did not look at their languages as objects of description or explanation. When they spoke, they tried different ways to make the others understand; when they listened to speakers talking to them, or to a third party, they tried to discover what the speakers wanted to say – what it was that had not been understood. The best innovators, however, must have been those who were most sensitive to the multi-faceted

dynamics of linguistic communication, those who really managed to turn their exploratory gaze from their experiences of the world to their experiences of the discourse, and in this sense, they were the first linguists – indeed, the greatest linguists of all times. Everything that we now have a name for in Linguistics – the verb, the noun, the sentence, the subject, assertion and negation, the speech act, the metaphor, strict word order, the case system - was invented at some point or other in by these linguistic giants and their successors. The safest general characterization of the entire tradition of Linguistics is that it consists of a series of footnotes to their work.

The architecture of the second generation of language allowed for a whole array of innovative strategies that could, to the extent that they were mutually-identified and accepted, raise the levels of communicative success between the speakers. All strategies were probably taken up simultaneously, by different innovative speakers, and they began to interact with each soon afterwards. Some speakers, for example, began to experiment with *definitions*. The technology of compositionality allowed them to narrow down gaps between different interpretations of signs by talking about them. Other speakers began to recruit components of presentational communication, so much older than language, for the challenge explanation. They waved their hands, nodded their heads and put expressions on their faces in new, voluntary ways, which were then conventionalized into language-related *gestures*. Certain speakers began to regulate discourse, not just in terms of trust and politeness, but also in terms of *relevance*. The most radical strategy, however, was to regulate the actual activity of experience-to-utterance translation. This posed two challenges: Speakers had to agree on the general properties of the proper way to speak about complex experiences; and they had to agree on ways in which these general properties should be signified at the level of the utterance. In other words, they had to develop correlated procedures for experience-to-message translation and for message-to-utterance translation.

The first challenge implied a full-scale investigation into the nature of complex experiences and complex messages, and into the nature of their relationships in the minds of the speaker *and* the listener. Speakers began to mutually-identify general *types* of complex experiences, and gradually came to understand some of their properties *in terms of their translation into the symbolic landscape*: 'This experience involves two participants; that one only one.' They began to develop prescriptive notions about the

relative importance of these properties: 'To understand experiences of this type or the other, it is essential to know this, not necessarily that'. They established norms of *aboutness*: 'This experience is told from this perspective, not that.' With time, they mutually-identified certain understandings of the *speech-act*; isolated and categorized *ways* of experiencing: 'This I saw with my own eyes, that I only heard about'. They made their first steps towards the development of *logic*. Certain signifieds were added to the symbolic landscape whose sole purpose was to point to mutually-identified relationships *between* experiences. They were the first *function words*. Eventually, on the basis of everything they had already achieved, they began to regulate stages in the production of the message: First, what the listeners had to receive - only then, what the speakers wanted to add. An etiquette of instructive communication.

The second challenge implied a full-scale investigation into the nature of the relationships between complex messages and complex utterances – again, in the minds of both speaker and listener. Speakers began to experiment with *linear order* and with *adjacency*: They uttered certain signifiers close to others, preceding or following them, to signify their intended relationship. Norms of *iconicity* emerged. Different types of combinations of signifiers came to be habitually used to signify different types of complex experiences. This was the beginning of *constructions*. When phonological material began to be reduced *within* constructions, *morphology* emerged. Certain heavily-used signs came to acquire the status of function words: The process of *grammaticalization* was launched. And speakers also experimented with prosody, with stress patterns, with tones and intonations: Everything that could be mutually-identified, at the level of pronunciation, and used to mark different aspects of the message.

As always, the entire process was highly variable. The technical solutions proposed by different innovators were beneficial and relatively easy for them and for some of their interlocutors – but not necessarily for everyone. In the process of social negotiation that ensued, many innovations were rejected or simply ignored); others were readily adopted by the entire group; some were accepted in a compromise form; yet others were enforced. The protocols that emerged in different languages echoed the individual interests of the inventors, their capacities and their insights about language and discourse – as they came to be reflected through the ever-changing social prism of negotiation and prescription. As different conventions began to spiral together, new opportunities and new problems

emerged, which required new solutions, and so on. The levels of complexity that emerged on both sides of the technology came to weave a web of complexity *between* them. Languages were on their way to maturity.

Again: Nothing in this process was *natural*. The conventions did not reflect a universal set of foundational notions about the logic of communication. The logic itself was gradually invented and developed, and was never universal. When speakers agreed that a description of a single experience required at least an indication of an entity and an eventuality (*what* happened to *whom*), and that any such minimal description should be covered, from beginning to end, by a mutually-identified intonational pattern, they actually invented the *proposition* – and the *sentence*. When they conventionalized norms of aboutness, and marked them by linear order or morphological markers, they invented grammatical functions. When they developed norms of conduct for co-operative conversations, they invented the maxims of pragmatics. When they decided that transitive experiences should be looked at and described from the perspectives of both participants, they invented *voice*. When they began to systematically relate the communicated experience to the moment of speaking, they invented *tense*, the *pronoun* system, *locative* markers – and so on and so forth. These are now so familiar to us that we have come to believe they have been there all along. But they have not. They are human inventions. They are always, still, *emergent* phenomena, variable, dynamic, imperfect.³⁶

³⁶ Consider *formal semantics*, for example. Among all the inventors in the history of language, there were also those who were implicitly interested in logic. Their inventions reflected the attempt to streamline the process of interpretation, to clean linguistic communication from ambiguity. Quantifiers, modalities, negation, scope – all these were introduced into language by logically-oriented inventors just because speakers do not necessarily take them into account. They emerged from attempts to understand: 'Let me make it clear for myself: *Did* they or *didn't* they? *All* of them or only *some*?'. All this evolved into the categories of language that we know through invention, persuasion, propagation and acceptance, which means that the products of the process cannot – and should not – look like reflections of the explicit categories of mathematical logic. A Formal Semantics that ignores all this participates in the process of prescription. It does not describe language. It describes a logician's utopia. Let formal semantics, then, investigate *whether* the relevant structures of language are founded on logic; how and to what extent; how pieces of logic are conventionalized; how they are actually used. Let it take a look at communicative

All of this does not mean that the process was entirely unconstrained. To begin with, the range of possible solutions to any given problem of communication was already heavily determined by the state of language itself at that moment. The process was also constrained, at least to a certain extent, by the difficulties of acquisition: Inventions which were too complicated or cumbersome were probably rejected or gradually simplified. In a very general way, the process was constrained by the contingencies of human bodies: Sounds that nobody could produce, for example, did not make their way into language. Most importantly, every new invention had to be functionally efficient: Speakers had to be able to use it for speaking about *their* experiences. Certain inventions probably proved their efficiency almost immediately. They just happened to capture something that was already shared by many of the speakers. Other inventions, however, took more time to be established: Most speakers did not experience the world in a way the invention could signify. The invention was prescriptive: The speakers had to learn to look at their worlds through the invention's lens. They had to learn to *experience-for-speaking*. We still do.

Because of all this, the languages of the world as we know them today manifest such complex typological patterns of variability and similarity. The problems that required technical solutions in the different languages, and the technical solutions themselves, were just as variable as their communities, their histories and their experiential worlds. But they were also, all of them, problems of - and solutions for - instructive communication. Many of the fundamental problems we have discussed simply came up repeatedly in language after language: It is impossible, for example, to continuously develop a language of the second generation and never to confront the problem of ambiguity. Many other problems, on the other hand, were more variably distributed. By the same token, many of the fundamental solutions emerged in language after language – not necessarily for the same problems - whereas others were only developed here and there.

interactions between logically-oriented speakers and others. Let it try to understand the *urge* to avoid opaqueness.

9.5 When language was deep in the second stage of its evolution, fast speech became the most important individual challenge. From the speech revolution, which took place within the social-linguistic environment of a human species not yet specifically adapted to language (most probably *Homo Erectus*) – emerged *Homo Sapiens*, as a new human species physiologically adapted to fast speech. It is *not* the case that we have language because we are species of speakers. We are a species of speakers because we invented language.

As we have seen, languages of the second generation brought with them a completely new set of challenges – the challenges involved in the translation of complex experiences into sets of signifiers, the calculation of the intersections between signified experiences, and the construction and segmentation of the linear utterance. Most importantly, a new parameter of linguistic skill emerged: Better speakers could speak faster, and better listeners could understand faster speech. Speaking faster was useful for three interrelated reasons: First, and quite obviously, it allowed for faster co-operation in collective activities; second, it made it more difficult for unwanted listeners to understand the conversation; third, and most importantly, just like Zahavi's peacocks and their tails, fast speakers indicated – actually *proved* – that they could speak fast *and* still make sense, express themselves clearly, produce clean segments of sound, and so on. Linguistic capacity gradually turned into a yardstick for intelligence and cooperativeness. As language took over human communities, as humans came to think of themselves as speaking animals, fast speaking came to play a major role in *mate selection*. Humans began to be violently selected for their ability to speak fast. This was the *speech revolution*: It resulted in a new human species.

These are the facts: We, modern humans, are physiologically adapted to fast speech in a variety of ways. Our lungs and the muscles surrounding them, our vocal tracts, the muscles around our mouths, and most importantly our tongues – all these are different from those of our ape relatives in ways that can only be explained as adaptations *for* fast speech. Our hearing is also specifically adapted for the interpretation of speech: As Philip Lieberman notes, when we listen to non-linguistic sounds, such as clicks or drum beats, the individual sounds merge into an indistinguishable buzz at rates exceeding 15 sounds per second (this is a statistical average: Some of us are better at this than others.) When we listen to speech, however, sounds are transmitted at about 20 to 30 phonemes per

second, and we hear each and every one of them. Over the years, researchers have tried to use fossil records in order to discover when it was that humans began to adapt to speech. Much of the work has concentrated on the size of the holes in the skull and vertebrae, through which run the nerves that innervate the lungs and the tongue: Fast speech requires much higher levels of voluntary control, which implies larger nerves, and thus larger holes. The consensus in paleontology is that in *Homo Erectus* the holes are still within the Chimpanzee range, but in early *Homo Sapiens*, and probably (even if less consensually) in the *Neanderthal*, they are similar to ours.

Most paleontologists agree that the speech revolution must have been driven by the social pressure to speak faster, which means that language must have been there first. Language, then, must have been invented by the ancestor of the early Sapiens, an Erectus of some kind. We already know that Erectus made the greatest innovative steps in the evolution of humankind, with the control of fire, the building of shelters, and the manufacture of complex tools. A genius species. As Donald argues, Erectus also invented mimesis, and with it art, music and ritual. It stands to reason, then, that in the million years or so of their evolution, Erectus communities turned into mimetic societies, then invented language – and then invented compositionality. When languages of the second generation began to take over, when Erectus communities began to turn into linguistic societies, human bodies - muscles, nerves, skulls and ears - began to adapt to their social realities. The speaking species emerged.

It is widely recognized in the literature that the physiological adaptations to language would remain totally nonfunctional without parallel adaptations in the *brain*. Deeper innervation requires a brain more specifically dedicated to speech. What the literature usually neglects to mention, however, is that the brain could not have adapted to the speech part of language alone: The competition was not just about the fast vocalization of utterances. It was about fast linguistic *communication*. Speakers and listeners must have been selected for their quickness in dissecting complex experiences into compositional messages, translating messages into experiential interpretations, and so on - everything that enhanced their capacity to join the discourse at the pace determined by the quicker speakers. Because of that, the Sapiens brains that emerged must have been ready – in a great variety of specific ways - for linguistic communication in a much broader sense.

As languages of the third generation emerged and stabilized in Sapiens communities, their speakers came to be selected for their general ability to internalize and follow the increasingly more complex regulations of the protocol: To understand (and *accept*) the conventional logic of language production and comprehension, and to apply entire clusters of rules to everything they said and heard - in real time, and in the right order. No conventions were assimilated as such.

Then, very gradually, the better speakers' selective advantage (in the biological sense) began to shrink. Speaking better (and acquiring quicker) than others remained a major *social* advantage – we still select people on the basis of their linguistic capacities – but the advantage gradually lost its *direct* effect on the number of offspring. This was part of a broader cultural process, but it was also, in a crucial way, specific to language: At a certain point, virtually all speakers - whose ancestors had survived the entire ordeal of the evolution of language – managed to acquire and use languages of the third generation to a certain rudimentary level, a socially-determined level that allowed them to find a place for themselves in the daily activities of language in their societies. This, however, does *not* imply that we are now, by our innate nature, less different from each other than our ancestors were. We have all climbed the ladder of language together – generation after generation of variable capacities. Reaching the rudimentary level of skill is still difficult for some: Their patterns of linguistic behavior are now thought of as linguistic *disorders*. The rudimentary level of linguistic success is thought of as the *normal* capacity for language.

Appendix A

Syntactic Analysis and the Question of Autonomy

The two sides to the debate over the autonomy of syntax, the generativists and the functionalists-cognitivists, have been thinking in terms of a relationship between syntax and a single level of meaning. The realization that two levels of meaning are involved changes the very formulation of the question. In terms of the original debate, Chomsky has been right all along: The syntactic structures of language are indeed divorced from the complexities of non-linguistic, private-experiential meaning. But the structures are correlated, and very intimately so, with the complexities of normative-linguistic meaning. All this is dictated by the functional logic of the system. The facts of syntactic autonomy should not be explained away: They reflect the autonomy of prescriptive, social meaning from private, experiential meaning.

Claims of autonomy claims have been made with respect to many different types of structural observations. This appendix presents a very informal discussion of four representative types: The paradigmatic questions of the definitions of *parts-of-speech* and *grammatical functions*, and the syntagmatic questions of *syntactic selection* and *long-distance dependencies*.

1. Parts-of-Speech and prescriptive categorization

Functionalism traditionally claims that *parts of speech* should be defined *notionally*, in terms of their meanings: A noun is a word that names a person, a place or an entity; a verb is a word that names an event or a state; an adjective is a word that names a quality, and so on. Parts of speech, then, are definable in terms of what they *do*: They carry meanings of different types. The generative claim (which in fact has a much longer history) is that notional definitions of this type do not account for the observable structural facts. The most ready examples are words that by the notional definition should be verbs or adjectives, but are in fact nouns: Words such as *fire*, *crime*, *party* and *war* name events and states, but they are not verbs; words such as *beauty* and *health* name properties, but they are not adjectives. The verb *dance* and the noun *dance* clearly name the same type of eventuality, but they play very different roles in the syntax. Generativists thus insist that

terms such as *verb* and *noun* should be defined on the purely formal basis of *distribution*: The paradigmatic question can only be answered syntagmatically.

As different as these positions may seem to be, they nevertheless share a similar understanding of the concept of *naming* – as a systematic function *from things* (physical entities, places, events, states, qualities, either in the real world or in people's experiences of it) *to signs*. Under this shared conception, the notional perspective would be right only if the properties of signs could be exhaustively determined on the basis of the experiential clusters associated with them. We already know that this is not the case. There is nothing in the experience-clusters associated with such words as *war* or *beauty* which implies that they should be nouns. Under this shared reading of naming, then, the Generativist perspective is quite right. Parts of speech may be *prototypically related* to experiential meanings – but as such they are autonomous.

According to LCT, however, naming is in fact a much more complicated function. It is not a descriptive function, from the world of experience to the world of signs, but a prescriptive one: From signs to the world of experience - *from the realm of social convention to the realm of private experiencing*. The symbolic landscape does not mirror the complexities of experience. It constructs its own simplified, stereotyped world view, *superimposing* on the analogue world of experience a set of digital distinctions and generalizations which the speakers of the language found necessary for the purposes of communication. What this means is that the traditional notional definition of parts of speech, which is wrong as is, is perfectly right in reverse order. Not: 'If something in the world of experience is a person, a place or an entity, then the word referring to it is a noun', but: 'If a word referring to something in the world of experience is a noun, then the thing in experience that the word refers to is to be thought of, *by prescription*, as a person, place or entity'. Not: 'If a thing in the world of experience is an eventuality, then the word referring to in language is a verb', but: 'If a word referring to a thing in the world of experience is a verb, then that thing in experience which the word refers to is to be thought of, for the sake of communication, as an eventuality'.

When naming is understood this way, the fact that the correspondences between parts of speech and experiential meanings are, at the very same time, both striking *and* imprecise, is exactly what we should expect to find. The correspondences should be *striking* – and

thus describable in terms of prototypicality - because the semantic categories of the symbolic landscape should capture something that is real enough in the experiences of the language's speakers to be mutually-identified as a reference point for communication. The digital categorization of things into entities and eventualities, which is signified by the distinction between nouns and verbs, bears a striking resemblance to something that is already present, in analogue form, in experience. The very fact, however, that it is a digital categorization which is superimposed on an analogue set of experiences necessarily implies that the resemblance cannot be *exact*. There always have to be mismatches between a linguistic categorization and the experiential complexities it is associated with. And indeed, this is what makes linguistic communication possible in the first place. To be able to mediate between the different experiential worlds of its speakers, the socially-constructed symbolic landscape has to be, at least to a certain extent, mismatched with each and every one of them.

The fact that the digital categorizations of the symbolic landscape are superimposed, by social convention, on the world of experience, thus means that these categorizations may find themselves superimposed on experiences which might not, at first sight, suggest themselves to this speaker or the other as candidates for the relevant category. This is the essence of prescription. Such words as *war*, *dance* and *beauty* are nouns in English not because there is something inherently nominal about them. They are nouns because the speakers of the language decided, at some point in the cultural evolution of the language, that it would be useful to be able to refer to the entire sets of war-experiences, dancing-experiences and beauty-experiences, *as if they were entities*. The capacity of referring to wars, dancing and beauty as entities is functional in the most foundational sense of the word: It enhances the expressive capacities offered by the language. Being able to refer to wars, dancing and beauty as entities, speakers may now qualify them ('unusual beauty'), count them ('three wars'), describe them as being involved in eventualities ('the dance went on for hours'), and so on.

Parts of speech, then, are indeed autonomous, but not because they are autonomous from meaning. They are autonomous because the semantic categories of the symbolic landscape, with which they are correlated by prescription, are autonomous from experiential meanings.

2. Grammatical functions as regulative conventions

Consider, for example, the *subject*. Functionalists claim that the syntactic notion of subject can and should be defined in terms of its communicative function – as the *topic* of the sentence: The subject is what the sentence is *about*. Generativists contend that this is sometimes true, but not always, and it is definitely not true when the subject seems to have no semantic significance (when it is *expletive*). Thus, in the English sentence *It is raining*', the subject *it* is not what the sentence is about, simply because it is very unclear what the *it* stand for in the first place. For Generativists, this has been taken as evidence for the contention that subjecthood is a formal notion, not a functional one: A subject is whatever appears in the sentence in subject position. As no reference is made in this analysis to topichood, or any other meaning-related concept, the notion of subject position, and its properties, turn out to be, quite clearly, non-functional.

As the debate stands, again, the structuralist perspective is correct. The subject cannot be defined as a function from a certain experience of communication (the experience of aboutness) to a certain structural configuration, for the simple reason that it appears in situations where the experience is not there. How can we move away from this unfortunate dilemma between an inexact functional definition and a non-functional structural stipulation? Well, we may go back to the notion of topichood, but this time the other way around: Not as a natural function from experience to language, but as a prescriptive function – from language to experience.

Think of the functionalist contention, that the subject is what the sentence is about, not as a description of something that is simply there, but as a conventionalized, regulatory principle, which a community of speakers decides to impose on itself precisely *because* the speakers of the community do not always naturally specify for their listeners what it is that they are talking about. Think of it as a prescriptive principle, through which the members of the community (acting as listeners) tell themselves (as speakers): 'From now on, whenever you tell me something, make sure that you remember to tell me what (or who) you are talking about. It will help me understand'.

For this prescriptive principle to be efficiently implemented, two separate issues have to be seen to: First, a special convention needs to be added to the set of conventions of

message construction of the language, that will oblige speakers, whenever they set out to construct a new message, to select a single entity from the symbolic landscape and designate it as the topic. This is exactly what the specification of the *topic-entity* in the *message-kernel* is there for. It forces the speaker to decide, at the very outset of the process of experience-to-speech translation, what he or she wants to talk about. Second, a special convention (or set of conventions) needs to be added to the utterance construction conventions, that will oblige speakers to mark the topic-entity in a specific way, mutually-identified with the listeners, so as to allow them to distinguish between the topic-entity and the rest of the entities in the sentence. In English, the convention chosen for this task is a *linearization* convention: The topic-entity has a designated position in the linear order of the sentence.

The fact that all this is a prescriptive effort, the fact that the specification of a topic is not a natural tendency of all speakers, crucially implies that its application cannot be a function of the experiential intents of the speakers. The conventions oblige speakers to choose and mark a topic-entity, *whether or not* there is indeed a certain entity which they themselves experience as the thing they intend to talk about. It is exactly because of this that the subject is required (in English) even in those cases where it is expletive. This very fact, that the notion of subject is indeed autonomous from experiential meaning, is thus the best indication of the specific type of prescriptive function that it fulfills.

Much of the debate about subjecthood, between generativists and functionalists, was also, at the same time, a debate about competence and performance. It is well-known that speakers of English often produce sentences without subjects, as in the answer 'Went out to get something to eat' to the question 'where is John?' In situations of this type, it is clear that the subject is not functionally necessary: The question already determines what the answer is going to be about. For Functionalists, this meant that the subject was not really obligatory in English, and that its distribution was indeed a function from the experience of communication to the actual usage of language. For Generativists, the important point was not what was happening in the course of performance, but what types of grammaticality judgments speakers produced when asked to consult their linguistic intuitions, and here speakers of English very clearly indicated that they knew that the subject was obligatory. They did not always behave according to their knowledge, but they knew nevertheless. The two competing camps thus highlighted (and discredited)

different types of empirical observations in the course of their argumentation, with the debate quite often deteriorating into an argument about presuppositions.

The realization that the subject is neither a natural function of communication, nor a purely-formal entity, but a regulative mechanism, a set of prescriptive conventions designed to organize the experience of communication, immediately reveals that the empirical observations used by *both* sides of the debate are not just valid – but also significant. In most cases, speakers clearly indicate in their grammaticality judgments that they have indeed internalized the relevant conventions of their language: *They know the law*. When explicitly asked about their knowledge of the law, they demonstrate their knowledge. They thus indicate that the subject is normatively obligatory in English. They do not, however, always behave according to the law: Social norms are not always followed. Specifically, speakers may feel, in certain communicative situations, that the normative behavior is superfluous, because whatever the norm was supposed to regulate in the experience of communication is already there. This is exactly the situation in the exchange above: The speaker who asked the question has already determined what the exchange was about, in a way that was already mutually-identified by the two sides before the speaker providing the answer started to speak. The responder may thus feel free to break the law, and produce a subjectless sentence, without taking the risk of communicative failure. This is exactly what we do whenever we encounter regulative laws: We usually refrain from crossing the street on a red light, but when we do – we (hopefully) do it in situations which we judge as relatively safe.

3. Syntactic selection and the essence of constructions

The basic syntagmatic observation is the fact that there are systematic restrictions on the co-occurrence of word-types in utterances: Different types of words are allowed and disallowed to appear with other types of words, in different types of constructions. The question is: How are these restrictions to be explained? It was around this question – the question of *selection* - that the struggle between the generative and functionalist approaches eventually evolved into an attempt, shared by scholars from both schools of thought, to reach theoretical common-ground. A new perspective emerged, now called *Constructionism*, - which is, at the same time (and for the first time) both functionally-

oriented and empirically-robust.³⁷ At the core of the new perspective lies the understanding that the syntagmatic observations of selection, for a long time used as evidence for autonomy, are actually correlated with much finer-grained notions of meaning than the functionalists and cognitivists traditionally invoked. These finer-grained notions of meaning, moreover, are not general functional principles of communication, but the specific meanings of the words and other syntactic constructions involved in the observable distributional patterns. The following examples, from Charles Fillmore's early article *On the Grammar of Hitting and Breaking*, are by now classic material. The verbs *hit* and *break* are sometimes allowed to appear in the same syntactic configurations – but are also often disallowed to do that:

- (1) a. John *broke* the stick
- b. John *hit* the stick

- (2) a. The stick *broke*
- b. *The stick *hit*

- (3) a. *I *broke* John on the leg (meaning: I broke *John's leg*)
- b. I *hit* John on the leg (meaning: I hit *John's leg*)

As Fillmore showed, these behavioral patterns (and many others) are not just manifested by these two particular verbs, *hit* and *break*. Verbs such as *bend*, *fold*, *shatter* and *crack* behave exactly like *break*; verbs like *slap*, *strike*, *bump* and *stroke* behave exactly like *hit*. The patterns are thus associated not with verbs, but with verb classes, and the classes are distinguished on the basis of a set of fine-grained semantic categorizations. *Break*, *bend* and *fold* are *change-of-state* verbs: They indicate, as part of their meanings, that their object came to change its state as a result of the relevant eventuality. *Hit*, *slap* and *stroke* are *surface-contact* verbs: They indicate that their object came to be touched on its surface as a result of the eventuality denoted by the verb. In English, all change-of-state

³⁷ In the following discussion, I will use the term 'construction grammar' to refer to the entire family of theories in this new field of research. Specifically, I will not discuss the important differences between Goldberg's 'construction grammar' and Croft's 'radical construction grammar'. Some references here: Fillmore (1967) Goldberg (1995, 2006), Croft (2001), Tomasello (2003), Jackendoff (1983, 1990), Langacker (2005), Levin (1993), Levin and Rappaport-Hovav (1995) and Dor (2000, 20005).

verbs appear in syntactic configurations (1) and (2), but not in (3). All surface-contact verbs appear in (1) and (3), but not in (2).

The syntactic restrictions in the above examples, then, are correlated with the semantics of the classes of verbs involved. How is this to be explained? In construction grammars, all parts of syntax, simple and complex, are *constructions* – systematic pairings of form and meaning. Sentences are systematic mappings of thematic roles onto grammatical functions: They construct *perspectives* on eventualities. Different Constructionist theories, then, suggest different technical variations on the explanatory principle which Adele Goldberg calls the *Semantic Coherence Principle*: Words and other types of constructions may appear together in the same sentences if the semantics of the different words and constructions are *compatible* with the meanings of the sentences. The construction in (1) maps an *agent* onto the subject, and either a *patient* or a *theme* onto the direct object. It thus provides a way of describing transitive eventualities from the point of view of the entity that instigated the eventuality. The construction in (2) maps a *patient* onto the subject: It provides a way of describing eventualities in which an entity undergoes a certain change – from the point of view of that entity. The construction in (3) involves the possessed body-part construction, whose intrinsic semantics indicates that it can only be used to refer to *themes* (the details of this semantics are not important at the moment). The different constructions may thus appear together in the relevant sentences only if their semantics are compatible: This is why *break*, with its patient, may appear in constructions (1) and (2), and *hit*, with its theme, may appear in (2) and (3).

All this makes good functional sense. It is still, however, formulated in a way that severely reduces its overall theoretical significance. Two issues are involved: The first has to do with the semantic categorizations which turn out to determine the relevant syntactic patterns. The question is this: What exactly are those semantic categorizations? Where do they belong in the general picture? The Constructionist view, formulated most clearly by Ray Jackendoff, is that the relevant semantic categorizations emerge in general cognition, from *conceptual structure*. Language, within this conception, provides structural tools for the communication of meanings – which are themselves independent of language.

Everything we have said so far, however, indicates that the relevant semantic categorizations cannot be general-cognitive – that they are the specifically-linguistic categorizations of the symbolic landscape: The categorizations highlight specific facets of experiences, abstracting away from the totality of their complexities; they isolate eventualities from the very entities which embody them; they assign stereotyped roles to these entities; they define classes of verbs, not the verbs themselves, and the meaning differences within classes are grammatically invisible; they are demarcated digitally, not analogically. Two additional arguments are especially important here: First, as many Constructionists have explicitly noted, the sentences which turn out to ungrammatical *because of the semantic incompatibility between their parts* – sentences like ‘the stick hit’, or ‘I broke him on the leg’ – are nevertheless perfectly comprehensible in terms of general meaning: We can tell what they are supposed to mean, but we also feel that this is not the way to say it *in English*. If the problem with these sentences is, at the very same time, linguistically-specific *and* semantic, then the semantics should be linguistically-specific. There seems to be no way around this conclusion. Second, the most important fact about the categorizations seems to be that there are so few of them.³⁸ The number of categorizations that determine syntactic distributions in languages is surprisingly small. There is a huge number of meaning categorizations one could think of, which participate in the way we conceive of the world, that are not syntactically relevant in any language: The distinctions between *friend* and *foe*, *important* and *trivial*, *expensive* and *cheap*, *endangered* and *safe* – and so on. And note: There does not seem to be anything in experience that would explain, on the basis of the way we understand the world, why certain categories turn out to determine syntactic patterns and others turn out not to. The categorizations which turn out to be grammatically relevant are not necessarily the most important, or the most robust, or the most visible. They are the categories most suitable for the social process of mutual-identification and signification.

Constructionist analyses, then, rather than merely describing sets of systematic relationships between linguistic form and general meaning, are actually extended investigations into the very essence of language as a *social* system of *meaning construction*. The observations of grammaticality, which form the basis of the analyses, are thus, in the final account, the observable traces of the social process of pointing and

³⁸ This has been discussed, among others, by Levinson (1997) and Frawley (1992).

naming, and then re-pointing and re-naming, through which the symbolic landscape is constructed as the mutually-identified worldview of the community. New signifieds are born into the world from the constant social struggle over the construction of common ground for communication, and because signifieds are born *together* with their signifiers – every moment in the process always involves pointing *and* naming – the behavioral patterns of the signifiers are always correlated, from the very beginning, with the semantics of their respective signifieds.

All of which leads us to the next issue: Constructionist theories are formulated in terms of *cognitive representation*. The semantic coherence principle, for example, in its different versions, is thought of as a fact (or a hypothesis) about the logic of the arrangement of linguistic representations *in the mind*. The principle, in other words, is taken to be a component of a *grammar*. A functionally-oriented, semantically-based grammar – but a grammar nonetheless.

Once we think of the principle, however, as a *prescriptive* convention - as the *order* to connect elements to each other only if they are semantically compatible - we can easily see that it actually captures the very essence of the process of linguistic production: It is precisely what *happens* when we produce grammatical sentences, and what fails to happen when we produce ungrammatical sentences. The principle does not describe a state of correspondence between meaning and form. It describes the process in which form (a linear arrangement of signifiers) is produced *on the basis of* meaning (a hierarchical arrangement of signifieds).

Let us look again at the grammatical and ungrammatical versions of (1) – (3). Instead of devising a grammar that would distinguish between them, let us try to *reconstruct* the processes of their production.

The *message-kernels* for the sentences involving the verb *break* should be:

(1a') *ASSERT(john)(break)*

(2a') *ASSERT(stick)(break)*

(3a') *ASSERT(I)(break)*

The set of extensions presented to the speaker, in all these cases, on the way towards the *basic message*, is the set of extensions associated with *break*. We have already looked at it before:

Extensions of event-type (<i>break</i>):	* <i>patient</i>	_____
	<i>agent</i>	_____
	<i>instrument</i>	_____
	<i>goal</i>	_____
	<i>reason</i>	_____
	<i>manner</i>	_____
	* <i>tense</i>	_____
	<i>time</i>	_____
	<i>place</i>	_____

In (1a'), then, the speaker assigns the signified **stick** as the obligatory *patient*, the signified **John** (which is already the topic-entity) as the *agent*, and the signified **past** as the *tense*. In (2a'), the speaker assigns the signified **stick** (which is already the topic-entity) as the obligatory patient, and **past** as *tense*. The *agent* is not obligatory here, and for a very good reason: The event signified by *break* is of the type *change-of-state*, and the conventions obligate the speaker to specify the identity of the entity which changed its state, the *patient*, and nothing more. In (3a'), the speaker cannot specify the complex signified **possessed body-part** (John on John's leg) as *patient*, because the complex signified is a variation on the semantic notion of *theme* – it is not semantically connected to change-of-state signifieds on the symbolic landscape in the first place. (3a') is thus an ungrammatical sentence because the message it signifies is not made possible by the conventions of message construction. The basic messages produced from (1a') and (2a') are then developed into their complete messages (we have already seen how that works). The signifieds of the two complete messages are then copied onto their respective signifiers, which are then linearly ordered, connected and phonologized. The end results are the full sentences in (1a) and (2a).

The *message-kernels* for the sentences involving the verb *hit* should be:

(1b') *ASSERT(john)(hit)*

(2b') *ASSERT(stick)(hit)*

(3b') *ASSERT(I)(hit)*

The set of extensions presented to the speaker, in these cases, would be the one associated with *hit*:

	* <i>theme</i>	_____
{ <i>one of these</i> } <i>obligatory</i>	* <i>agent</i>	_____
	* <i>instrument</i>	_____
	<i>goal</i>	_____
Extensions of event-type (<i>hit</i>):	<i>reason</i>	_____
	<i>manner</i>	_____
	* <i>tense</i>	_____
	<i>time</i>	_____
	<i>place</i>	_____

Two elements are different here: First, the set of extensions specifies an obligatory *theme*, rather than a *patient*. Second, the set of extensions obligates the speaker to specify either an *agent* or an *instrument* – and for a very good reason: *Hit* is a *surface-contact* verb, which means that it necessarily involves the contact between two entities, one of which is a *theme*. (Again, this as a prescriptive demand: 'Whenever you tell us about an event of hitting, make sure that you tell us *what hit what*'.) In (1b'), then, the speaker assigns the signified **stick** as the obligatory *patient*, the signified **John** (which is already the topic-entity) as the obligatory *agent*, and the signified **past** as the *tense*. In (2b'), the speaker assigns the signified **stick** (which is already the topic-entity) as the obligatory patient, **past** as *tense*, and nothing more. The speaker thus fails to specify an obligatory extension, and the sentence cannot be produced. In (3b'), the speaker specifies the *theme* with a

complex signifier whose semantics is thematic. The process launched with (1b') and (3b') is then continued, to produce the sentence in (1b). and (3b).

Each grammatical sentence, then, is the product of a specific *convergence* of prescriptive conventions. The sentence in (1a), 'John broke the stick', is what it is because (i) the signified *break* is semantically-attached, by convention, to *agents* and *patients*; (ii) John is a suitable *agent*; (iii) the stick is a suitable *patient* for *break* (very specifically so, unbreakable entities would not do); (iv) the conventions of language obligate the speaker to specify a *topic-entity*; (v) John is a suitable topic-entity; (vi) past is a suitable *tense*; (vii) each of these signifieds is by convention marked by a specific signifier; and (viii) the linearization conventions of the language specify how these signifiers should be arranged in the linear order.

Complex constructions, then, of which sentences are a specific case, are the products of different convergences of conventions. Frequently-used constructions are remembered by speakers as such. Because of that, as Michael Tomasello and others have shown, constructions play a crucial role in the process of acquisition. Language, however, is not the set of its constructions (or, in Croft's terms, the population of its utterances). It is the social technology that produces them.

4. Island Constraints and the Convergence of Conventions

The problem of long-distance dependencies, also known as *island-constraints*, has taken center-stage in Generative Grammar (and as a result also elsewhere), virtually since its inception.³⁹ It is indeed based on some of the most intriguing observations about the structures of sentences. As the contrasts in (4) and (5) demonstrate, there seem to be serious constraints on the formulation of questions and other complex sentences:

- (4) a. Who did the girl hit? (Answer: The girl hit *the boy who delivered the pizza*)
b. *What did the girl hit the boy who delivered? (Answer: The girl hit the boy who delivered *the pizza*)

³⁹ See Ross (1967), Chomsky (1973), Dean (1991), Erteschik-Shir and Lappin (1979), Szabolcsi (2006) and Hofmeister and Sag (2010).

(5) a. Whom did your obsession with Madonna annoy? (Answer: Your obsession with Madonna annoyed *your father*)

b. *Whom did your obsession with annoy your father? (Answer: your obsession with *Madonna* annoyed your father)

Generative grammarians have always taken observations of this type to be the best demonstrations of the autonomy of syntax. The constraints involved in these observations were thus described and explained as purely formal limitations on the possible arrangements of syntactic elements with respect to each other.⁴⁰ Over the years, generative grammarians have developed increasingly complex formal theories designed to explain all this, which were unfortunately only capable of capturing (some of) the phenomena with formal stipulations which did not seem to be independently motivated in any real sense.

⁴⁰ Very informally, the issue was formulated this way: The question word at the beginning of a question sentence (*who*, *what*) is structurally related to a certain syntactic position in the parallel assertive sentence – in the above examples, the position of the *answer* in the sentences in parenthesis. In (4a), for example, the word *who* is structurally related to the position of its answer, the noun-phrase *the boy who delivered the pizza*, which is the direct object of the verb *hit* in the answer sentence. Originally, this structural relationship was thought of as the indication that a *transformation* had taken place: In early Generative Grammar, question sentences were produced, by transformation, from assertive sentences. In (4a), then, the word *who* was originally situated, in *deep structure*, in the object position of the assertive sentence, and it was then moved to its new position, in *surface structure*, by a question transformation. The relevant constraints were thus formulated as constraints on transformations: Certain structural positions in assertive sentences disallow the movement of words from within them to the outside: They are *islands*. The position of the answer in (4a), *the boy who delivered the pizza*, is not an island. The position of the answer in (4b), *the pizza*, is an island: The noun-phrase *the pizza* is structurally embedded within another, larger noun-phrase, *the boy who delivered the pizza*. Complex noun-phrases of this type are barriers to transformations, and because of that, (4b) is ungrammatical. The position of the answer in (5a) allows for the relevant transformation, and because of this, (5a) is grammatical. The answer in (5b), the noun-phrase *Madonna*, is structurally embedded within the subject of the answer sentence, *your obsession with Madonna*, and subjects, like complex noun-phrases, are barriers to transformations. This is why (5b) is ungrammatical.

Functionalists, cognitivists and pragmaticists have attempted to find more substantial ways to explain the relevant phenomena, proposing two types of theories. The first attempts to explain the phenomena in terms of the cognitive task involved in language *processing*. Paul Dean, for example, suggests that restrictions on transformations are correlated with the cognitive notion of *attention*: When listening to question sentences, and other such complex structures, listeners have to simultaneously attend to two parts of the syntactic structure, and figure out which part of the sentence the question word stands for, and this strains their limited working memory. The constraints on transformations thus ensure that sentences are only generated if they are not too complicated to process, if their structures do not overstrain the attention of the listener. The second type of theory, most clearly articulated by Nomi Erterschik-Shir and Shalom Lapin, argues that constraints on transformations are correlated with the pragmatic notion of *dominance*: A syntactic element in a sentence is dominant if, by uttering the sentence, the speaker intends to direct the attention of the listener to the meaning of the relevant syntactic element. Syntactic elements can be moved away from their original syntactic positions only if they are potentially dominant.

These two theories, then, attempt to explain the relevant syntactic facts on the basis of the communicative experiences of the speaker and the listener, and because of that, they suffer from the same problems encountered by the functionalist and cognitivist theories noted earlier: They attempt to explain very clear, fine-grained syntactic facts on the basis of analogue notions such as attention and intent, and because of that, they are only capable of providing first approximations towards the phenomena, not to capture them at the level of specificity that they require. Moreover, explaining the syntactic facts on the basis of the cognitive capacities and intentions of speakers and listeners carries a very clear implication: Speakers and listeners with different capacities and intentions should have different judgments with respect to the relevant facts. Listeners with a better capacity for attention, for example, should judge more complicated transformations as grammatical than those who find the relevant cognitive task more demanding. This is clearly not the case.⁴¹

⁴¹ A technical clarification is required: There are two general types of Island constraints – strong and weak (Szabolcsi 2006). Weak islands are highly variable and easily violable. As Hofmeister and Sag (2010) show, on the basis of experimental data, their patterns of acceptability are

Most importantly, the theories implicitly accept the generative formulation of the entire problem, as the grammatical problem of transformation constraints. They consequently attempt to construct an experientially-based grammar for the relevant syntactic phenomena, and this, as we have already seen, is exactly what cannot be done. Let us, then, think these two types of theories not as attempts at direct explanations of the phenomena, but as attempts to describe the experiential *difficulties* involved in communicative events in which the relevant sentences are produced. Dean's theory is already formulated in these terms: It is a description of a certain type of difficulty involved in the process of comprehension. The listener identifies the question-word, understands that the speaker is asking something, and he or she now has to figure out what the speaker's question is *about*. Erterschik-Shir's theory may be re-interpreted as a description of one of the challenges involved in the process of production - the challenge of selecting and marking the particular semantic element, that the listener should eventually direct his or her attention to as the thing asked-about. What we should be looking for, then, is a set of prescriptive conventions for the regulation of communicative events of this type – conventions that were socially developed and mutually-identified *in order to overcome these difficulties*.

As a first step, let us look at the essence of the speech act of asking. As formal semanticists and pragmaticists have demonstrated, the speaker who asks actually does two things at the same time: He or she performs the *assertive* speech act, telling the listener about a certain eventuality, and also the *interrogative* speech act, asking the listener to indicate the identity of a certain semantic element of the eventuality, which the speaker marks as unspecified. Thinking about all this as a prescriptive convention, the combined speech-act may be formalized as the following requirement, made by the community and addressed to the speaker: 'If you wish to ask us something, tell us who and what we are talking about - which type of eventuality took place and who was

correlated with different determinants of processing effort. But the idea that island constraints tell us something of significance about the autonomy of linguistic structure has always been inspired by *strong* islands – such as the one we are dealing with here. English speakers' judgments with respect to (4) and (5) – *and* their patterns of performance - are clear-cut. Because of this, processing effort as such cannot explain them.

involved - and make sure that you specify for us a single element of this eventuality, and mark it as the thing you ask about'.

There is a non-trivial restriction involved in this requirement, which lies at the heart of the matter. What the restriction says to the speaker, in effect, is this: 'Make sure that you ask your listener about the event you are telling about; don't tell your listener about one event and then ask about another'. This, to be sure, is a very reasonable demand. Listening to someone talking about one event and then asking about another is confusing enough when the talking and the asking take place *one after the other*. It is indeed a challenge (in terms of attention) to keep track of the speaker's intentions in this type of situation. It is even more difficult, however, to do this when the two speech-acts take place at the same time. The conventional restriction, then, is there to minimize this confusion. It is a restriction imposed on the speaker by the community, in its role as a community of listeners. As a convention, however, it is prescribed for all speakers, all the time, whether or not they (or their listeners) actually find themselves straining their capacities, or losing sight of their intentions. This is already familiar: The restriction bears a *striking resemblance* to a certain essence in the experience of communication, but it is nevertheless, by its very functional definition, *autonomous* from this experiential essence. It is not there as a reflection of this essence. It is there to regulate it.

As a first approximation, then, what all this translates to in terms of the process of production is the following characterization of the *message-kernel* associated with asking:

ASK(unspecified element of eventuality)(topic)(eventuality)

This message-kernel allows the speaker, for example, to formulate such kernels as:

ASK(reason)(john)(break)

ASK(patient)(john)(break)

ASK(agent)(unspecified)(break) (a question about the *agent*, which is the topic)

ASK(time)(john)(break)

And so on.

As strange as this may seem at first sight, this very simple characterization already explains the facts in (4) and (5), and a host of similar observations. Let us take another look at our examples. The message-kernels for the grammatical sentences in (4a) and (5a) are the following:⁴²

(4a') *ASK(theme)(girl)(hit)*

(5a') *ASK(experiencer)(obsession)(annoy)*

These are good kernels. They may be developed into their respective basic messages, and then into their complete messages. The signifieds of the complete messages may then be replaced by their respective signifiers (including the signifier **who** for the unspecified theme and the unspecified experiencer), and these signifiers may then be linearly ordered (with an additional linearization convention which specifies that the question-signifier should be positioned as the first word of the sentence). (4a) and (5a), then, are grammatical.

The message-kernels for the ungrammatical (4b) and (5b), however, directly violate the entire prescriptive principle of the speech-act. The speaker is supposed to tell the listeners about an eventuality and ask about one of the semantic elements involved in *that* same eventuality, but the semantic elements that are supposed to be the things asked-about in (4b) and (5b) are not involved in the relevant eventualities at all. In (4b), the speaker tells of an eventuality of *hitting*, but the *pizza* is not involved in it. The *pizza* is involved in another eventuality, that of *delivering*, which is part of the additional information that the speaker was allowed to specify about the *boy*, the *specified* theme of the hitting event - in the *answer sentence*. As we have seen, language allows the speaker, whenever he or she specifies an entity in the process of message construction, to further specify whether the entity also participated in another eventuality, and in the answer sentence to (4b), the speaker decided to specify *deliver*. None of this, however, can happen in the parallel question sentence, because the major prescription governing the production of questions

⁴² The *experiencer* in (5a') is the thematic role of the cognitive entity which goes through an internal change-of-state as a result of the eventuality.

is the specification that the thing asked-about should play a thematic role in the asserted eventuality. The pizza does not get there in the first place.

The same is true of (5b). The speaker tells about the eventuality of *annoying*, but *Madonna* is not involved in this eventuality at all. What is involved in the eventuality of annoying is the *obsession*. Madonna is the object of the obsession, not of the annoyance, and she gets into the process of production only in the answer sentence, where she is specified as additional information about the obsession. In the question sentence, she does not get into the process at all. Madonna, and the pizza, in other words, are just like all the other things in the world which are not involved in the eventuality chosen by the speaker in the message-kernel: The conventions of language prescribe that they cannot be asked about. Madonna and the pizza can only be asked about if they themselves participate in the eventuality asserted by the speaker. The questions in (6) and (7), for example, are producible and thus grammatical, because the asserted eventualities are those of the *delivering* and the *obsessing*:

(6) What did the boy whom the girl hit deliver?

(7) Who did you obsess with in a way that annoyed your father?

Here, however, questions about the *girl* and the *father* are unproducible, because they are involved in the eventualities of *hitting* and *annoying*:

(8) *Who did the boy who hit deliver the pizza?

(9) *Who did you obsess with Madonna in a way that annoyed?

All of which means that no transformations are involved here. Each of the question-sentences we have looked at is indeed related to a corresponding answer-sentence, but the relationship is not structural – it is *semantic*. The pairs of sentences are, quite simply, *questions and answers about the same eventualities*. They are related to each other because they are founded on related semantics. They are produced independently, however, and their production processes are governed by different sets of conventions. The set of conventions for the production of questions is more restrictive than the parallel set for the production of assertive sentences, because it includes the additional restriction on the semantics of the unspecified element, the thing asked-about. The conventions of

the language allow the speaker to tell about one event and then tell about another in the same message ('the girl hit the boy who delivered the pizza'), but they do not allow the speaker to tell about one event and then *ask about another* (*'what did the girl hit the boy who delivered'). Because of that, the set of questions producible by language about a certain eventuality is always smaller than the parallel set of assertions *about the same eventuality*. This is a semantic fact: It is a fact about the linguistic conventionalization of communication. If we mistake it for a structural fact, and think about the question-sentences as generated, by transformation, from their assertive counterparts, we find ourselves having to explain why some semantic elements which may appear, in specified form, in the answer-sentences, cannot re-appear, transformed into question-words, in the question sentences. This is the wrong way to think of the entire problem.

Let us take this one step further. The observations at hand are known as *long-term dependencies*, because it seems that some question-sentences (and other types of sentences) which are more complicated than those in the ungrammatical examples above, are nevertheless perfectly grammatical. Look at (10), for example:

(10) What did Eve say that Mary thought that Bill told her that Dean was going to buy in the supermarket?

At first sight, this sentence seems to go against everything we have said. The asserted event in the sentence is that of *saying*, but the thing asked-about, the thing bought in the supermarket, is involved in the event of *buying*. To understand what is happening here, think about the line of events in the sentence – *saying, thinking, telling* and *buying*. All but the last denote what we shall call *epistemic events*: Events in which *a cognitive entity makes a judgment about events*. Other epistemic events are *remembering, informing, knowing, doubting, explaining, assuming, believing*, and so on and so forth. What all of these eventualities have in common is a certain stereotyped semantic relationship between a cognitive entity and an event, the factuality of which the cognitive entity asserts, or doubts, or affirms, or declares, and so on. Sentences like (10) may thus be lengthened at will, but only as long as all the eventualities involved in them, apart from the very last, are epistemic events. A sentence like (10) with other types of eventualities would have to look like (11):

(11)* What did Eve say that Mary broke the table that Bill told her that Dean was going to buy in the supermarket?

Or like (12):

(12) *What did Eve say that Mary met Bill who walked along with Dean who was going to buy in the supermarket?

These just don't work. All the events apart from the last one have to be epistemic – and for a very good reason: Consider a possible answer to (10), the assertive sentence 'Dean was buying a loaf of bread in the supermarket'. The sentence is asserted by the speaker: Its message-kernel is *ASSERT(Dean)(buy)*. What this means is that the speaker himself or herself establishes a certain epistemic relationship with the event of buying. The sentence says: '*I, the speaker, say: Dean was buying a loaf of bread in the supermarket*'. What the speaker *does* in asserting belongs in the class of epistemic events. Now, consider a longer version of this: 'Bill told Mary that Dean was buying a loaf of bread in the supermarket'. What the sentence says is: '*I, the speaker, say that Bill told Mary: Dean was buying a loaf of bread in the supermarket*'. The *telling* event, involving Bill and Mary, is of the same type as the event of *asserting*. It is an extension of the speaker's assertion. The speaker says what Bill told Mary *about the buying event* – which is still the asserted event. The same is true of the full-length answer to (10). It says: '*I, the speaker, say that Eve said that Mary thought that Bill told her: Dean was buying a loaf of bread in the supermarket*'. This means that the prescriptive principle involved in the speech-act of asking, presented earlier as the message-kernel *ASK(unspecified element of eventuality)(topic)(eventuality)*, should be re-formulated as *ASK(unspecified element of asserted-eventuality)(topic)(eventuality)*. It allows the speaker to choose an unspecified element of *any* eventuality – as long as it is the *asserted eventuality*. The revised message-kernel allows for the production of sentences like (4a) and (5a) and (10), and disallows the production of all the ungrammatical sentences.

The fact that the symbolic landscape of a language includes a conventionalized set of epistemic verbs allows speakers to produce utterances not just about what they think or believe, doubt or conjecture about the world – but also about what (they think) other people do. A language whose symbolic landscape includes such a set thus provides its

speakers with much more expressive power as far as *asserting* is concerned: Speakers may assert more things, tell their listeners more about the epistemic sources of what they say, and, probably most importantly, say things about what other people think without committing themselves to the truth of their thoughts. All this is positively functional, but how does it apply to *asking*? Well, a language that prescribes the first version of the message-kernel disallows speakers, when asking, to take advantage of the full expressive power of epistemic verbs. A language that relaxes the restriction, and allows speakers to ask about events they were told about by others, thus gives them more expressive power.

The revised convention, then, is a complex compromise between speakers' and listeners' interests. It explains all the observations we have looked at, and a host of other observations that the generative tradition has always thought of as the very best demonstrations of the autonomy of syntax. The restriction does not explain everything, of course: There are other types of observations (the restrictions associated with *wh-islands*, for example) which require additional explanation. The general structure of these observations, however, and the type of grammaticality judgments associated with them, strongly indicate that they can be explained on a similar basis, as reflections of additional prescriptions, socially-constructed and semantically-based, for the formulation of questions, and other complex sentences.⁴³

⁴³ Note that the clusters of conventions involved here maintain *dependency* relations between them which indicate that some of them could have only been added to the language *after* some of the others. Simple assertions and questions were probably there before complex ones, and as long as speakers only asked simple questions, the issue of island constraints did not arise. We may assume that at some point, speakers began to experiment with epistemic state descriptions. In their simple form, however, the norms of asking did not allow for the construction of questions about indirectly-reported events. The interrogative kernel had to be revised. Then, when speakers began to develop the system of embedding, the same problem appeared again. This time, however, nobody seems to have taken the trouble to fix the problem, for three interrelated reasons: Questions of this type are extremely cumbersome and confusing; allowing them would produce a entire set of new problems; and most importantly, the change would not alter the expressive power of the language. The questions can already be formulated in the more traditional way, simply by treating the qualifying event as the major reported event of the message. There is no reason to ask 'what did the girl hit the boy who delivered?' when you can already ask 'what did

The generative argumentation about the autonomy of syntax has always been accompanied by a sort of meta-observation, intended to highlight the importance of the relevant empirical facts for the overall theory of language. It has been claimed that native speakers are genuinely surprised when presented with ungrammatical sentences such as (4b) and (5b). They testify that they never came across anything similar, but they nevertheless immediately know that the sentences are ungrammatical. What this was supposed to imply was that the speakers could not have learned about the ungrammaticality of these sentences from their past experiences, which means that they must have grasped the ungrammaticality on the basis of their innate knowledge of language. The argument does not hold, but the observation is actually correct: We really do not come across ungrammatical sentences like (4b) and (5b). This is significant not because it indicates that language is innate, but because we do come across so many other types of ungrammatical sentences. Many types of ungrammaticality are familiar to us, but (4b) and (5b) are not. They really are surprising.

What this testifies to is this: In our regular lives as speakers, not as linguists, all the sentences we ever encounter as candidates for grammaticality judgment are produced by speakers who *intended to say something*. We are thus familiar with two types of ungrammaticality. Speakers may either intend to follow the conventions to the letter, but *fail* to do so, or they may deliberately break the conventions in order to communicate something (like in Grice's maxim violations). Characteristically, when speakers fail to follow the conventions (because they only have partial command of the language, or because they are exhausted, and so on), they manifest a certain reduced capacity with respect to the entire cluster of conventions they have to follow. In situations of this type, we find instances of ungrammaticality spread all over the sentence. We rarely, if ever, come across situations in which speakers manage to produce complex sentences, following all the conventions of the language to the letter – *except one*. A speaker who produces a sentence like (4b) or (5b) by mistake must have perfect command of the entire cluster of conventions involved in the production of these sentences (the conventions of the symbolic landscape and message construction, of linearization and phonologization,

the boy who the girl hit deliver?' Island-questions, then, were left out of the arsenal of languages, a fact that nobody even had to have noticed – until linguists began to invent them *as examples*.

and so on) – and a total misunderstanding of the conventional nature of the speech act of asking. This is a highly unlikely state of affairs. The other type of ungrammaticality we are familiar with is the one produced by the speaker in a deliberate fashion. It is difficult to think of a communicative intent that would be served by sentences like (4b) and (5b) – unless, of course, we think of them as instances of *poetry*. Poetry is one of the strategies adopted by innovative speakers in their struggle to push the expressive envelope of language further: It is a deliberate attempt to violate the social conventions of the language, to go beyond them. It is actually quite easy to imagine (4b) or (5b) appearing as lines in an avant-garde poem, and it is also easy to see that, as part of a poem of this type, they would be *supposed* to surprise, to shock, to disturb the conventional order of language and thus reveal its constraining nature. This is exactly what happens to speakers who are asked by linguists to provide a judgment about a sentence of this type: The sense of shock is actually more acute, because the sentences are not presented as poetry – but as potential instances of regular speech. Innateness has nothing to do with this.

References

- Anderson, J. M. (2006). The non-autonomy of syntax. *Folia Linguistica*, 3-4, 223-250.
- Anderson, J. M. (1997). *A Notional Theory of Syntactic Categories*. Cambridge University Press.
- Anderson, M. L. (2003). Embodied Cognition: A Field Guide. *Artificial Intelligence* 149, 91–130.
- Arbib, M. (2003). The evolving mirror system: A neural basis for language readiness. In: Christiansen and, M. and S. Kirby (eds.), *Language Evolution*. Oxford University Press.
- Armstrong, S., L. Gleitman and H. Gleitman (1983). What some concepts might not be. *Cognition*, 13, 263-308.
- Aronoff, Mark, Meir, Irit, Padden, Carol A., & Sandler, Wendy (2008). The roots of linguistic organization in a new language. *Interaction Studies: A Special Issue on Holophrasis Vs. Compositionality in the Emergence of Protolanguage*, 9, 1, 131-150.
- Austin, J. L. (1962). *How to Do Things with Words*. Harvard University Press.
- Bach, K. (2006). Implicature vs. explicature. Paper given at the Granada workshop on *Explicit Communication*, in honor of Robyn Carston.
- Bates, E., Bretherton I. and L. Snyder (1988). *From First Words to Grammar: Individual Differences and Dissociable Mechanisms*. Cambridge University Press.
- Bates, E., Dale, P. S. and D. Thal (1995). Individual differences and their implications for theories of language development. In Fletcher, P. and B. MacWhinney (Eds.), *Handbook of Child Language*. Blackwell Publications.
- Bickerton, D. (1990). *Language and Species*. University of Chicago Press.
- Bickerton, D. (1981). *Roots of Language*. Karoma Publishers.
- Boroditsky, L. (2003). Linguistic Relativity. In Nadel, L. (Ed.) *Encyclopedia of Cognitive Science*. MacMillan Press.
- Botha, R. (2003). *Unraveling the evolution of language*. Elsevier, Oxford.
- Botha, R and C. Knight (2009). *The Cradle of Language*. Oxford University Press.
- Bucholtz, M. and K. Hall (2008). All of the above: New coalitions in sociocultural linguistics. *Journal of Sociolinguistics* 12, 4, 401-431.

- Bybee, J. L. (2006). *Frequency of use and the organization of language*. Oxford University Press.
- Carston, R. R. (2004). Truth-conditional content and conversational Implicature. In Bianchi, C. (ed.). *The Semantics/Pragmatics Distinction*. Stanford University: CSLI.
- Chomsky, N. (1973). Conditions on transformations. In: S. Anderson and P. Kiparsky (eds.), *A Festschrift for Morris Halle*. Holt, Reinhart & Winston.
- Chomsky, N. (1956). *Syntactic Structures*. Mouton.
- Choi, S. & Bowerman, M. (1991). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition*, 41, 83-121.
- Clark, H. (1996). *Using Language*. Cambridge University Press.
- Comrie, B. (1981). *Language Universals and Linguistic Typology*. Chicago University Press.
- Christiansen, M. H., C. Collins and S. Edelman (2009). *Language Universals*. Oxford University Press.
- Christiansen M. H., and S. Kirby (eds.), (2003a). *Language Evolution*. Oxford University Press.
- Christiansen, M. and S. Kirby (2003b). Language evolution: The hardest problem in science? In Christiansen and Kirby (eds.), *Language Evolution*. Oxford University Press.
- Corballis, Michael (2003). From hand to mouth: The gestural origins of language. In Christiansen and Kirby (eds.), *Language Evolution*, Oxford University Press.
- Croft, W. (2002). *Typology and Universals*. Cambridge University Press.
- Croft, W. (2001). *Radical Construction Grammar*. Oxford University Press.
- Croft, W. (2000). *Explaining Language Change: An Evolutionary Approach*. Longman: Linguistic Library.
- Croft, W. (1995). Autonomy and functionalist linguistics. *Language*, 71, 3, 490-532
- Croft, W. (1991). *Syntactic Categories and Grammatical Relations: The Cognitive Organization of Information*. University of Chicago Press.
- Deane, P. (1991). Limits to attention: a cognitive theory of island phenomena. *Cognitive Linguistics* 2.1–63.
- Deacon, T. W. (1997). *The Symbolic Species: The Co-Evolution of Language and the Brain*. Norton Publications.

- Derrida, J. (1977). Signature Event Context, *Glyph 1*.
- Deutscher, G. (2005). *The Unfolding of Language: An Evolutionary Tour of Mankind's Greatest Invention*. Metropolitan Books.
- Donald, M. (1991). *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition*. Harvard University Press.
- Dor, D. (2011). The evolution of language as a communication technology. Paper to be given at the *Social Origins of Language* workshop, February 17th, UCL.
- Dor, D. (2005). Towards a semantic account of that-deletion in English. *Linguistics* 42-3, 345-382.
- Dor, D. (2000). From the autonomy of syntax to the autonomy of linguistic semantics: Notes on the correspondence between the transparency problem and the relationship problem. *Pragmatics & Cognition*, 8 (2), 325–356
- Dor, D. and E. Jablonka (2010). Plasticity and canalization in the evolution of linguistic communication: an evolutionary developmental approach. In: Larson, R. K., V. Deprez and H. Yamakido (eds.) *The Evolution of Human Language: Biolinguistic Perspectives*. Cambridge University Press.
- Dor, D. and E. Jablonka (2001). How language changed the genes. In Tabant J. Ward. S. (editors). *New Essays on the Origin of Language*. Mouton de Gruyer.
- Dor, D. and E. Jablonka(2000). From Cultural Selection to Genetic Selection: A Framework for the Evolution of Language. *Selection*, 1, 33-56.
- Dunbar, R. (1998). *Grooming, Gossip, and the Evolution of Language*. Harvard University Press.
- Eckert, P. (2005). Variation, convention, and social meaning. Paper presented at the Annual Meeting of the Linguistic Society of America.
- Eckert, P. (2000). Three waves of variation study: The emergence of meaning in the study of variation. www.stanford.edu/~eckert/PDF/ThreeWavesofVariation.pdf
- Elman, J., A. Karmiloff-Smith, E. Bates, M. Johnson, D. Parisi, and K. Plunkett (1996). *Rethinking Innateness: A Connectionist Perspective on Development*. The MIT Press.
- Ehrman, M. E., Leaver B. L., and R. L. Oxford (2003). A brief overview of individual differences in second language learning. *System* 31, 313–330
- Erteschik-Shir, N. and S. Lappin (1979). Dominance and the functional explanation of island constraints. *Theoretical Linguistics* 6.43–84.

- Evans, N. and S. C. Levinson (2008). The myth of language universals: Language diversity and its importance for cognitive science. *Behavioral and Brain Sciences*, 32, 429-492.
- Fillmore, C. (1967). The grammar of hitting and breaking. In Jacobs, R. and Rosenbaum, P. *Readings in English Transformational Grammar*. pp. 120-133.
- Fodor, J.A., M. F. Garrett, E.C.T. Walker and C.H. Parkes (1980). Against Definitions. *Cognition* 8, 63-367.
- Fodor, J. A. and E. Lepore (1996). The red herring and the pet fish: why concepts still can't be prototypes. *Cognition*, 58, 253 – 270
- Frawley, W. (1992). *Linguistic Semantics*. Lawrence Erlbaum Associates.
- Fitch, W. T. (2010). *The evolution of language*. Oxford University Press.
- Gardenfors, P. (2004) Cooperation and the evolution of symbolic communication. In D. Kimbrough Oller and U. Griebel (eds.), *Evolution of Communication Systems: A Comparative Approach*. The MIT Press.
- Gardiner, A. (1932). *The Theory of Speech and Language*. Oxford : The Clarendon Press.
- Gathercole, S. E., and A. D. Baddeley (1993). *Working Memory and Language*. Lawrence Erlbaum Associates.
- Ginsburg, S. and E. Jablonka (2007a), The Transition to Experiencing: I. The Evolution of Limited Learning. *Biological Theory*, 2(3), 218–230.
- Ginsburg, S. and E. Jablonka (2007b), The transition to experiencing: II. The Evolution of Associative Learning Based on Feelings. *Biological Theory*, 2(3), 231–243.
- Givón, T. (2002). *Bio-linguistics: the Santa Barbara lectures*. John Benjamins.
- Goldberg, A. (2006). *Constructions at Work: the nature of generalization in Language*. Oxford University Press.
- Goldberg, A. (1995). *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago University Press.
- Grice, H. P. (1967), *Logic and conversation*. Harvard University Press.
- Hauser, M. D., Chomsky, N. and Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how does it evolve? *Science*, 298, 1569–1579.
- Hofmeister, P. and I. Sag (2010). Cognitive constraints and island effects. *Language*, 86, 2,366-415.
- Hopper, P. and E. Traugott (2003). *Grammaticalization*. Cambridge University Press.

- Hurford, J. (2007). *The Origins of Meaning: Language in the Light of Evolution*. Oxford University Press.
- Itkonen, E. (2008). The central role of normativity for language and linguistics. In J. Zlatev, T. Racine, E. Sinha, & E. Itkonen, *The Shared Mind: Perspectives on Intersubjectivity*. John Benjamins.
- Itkonen, E. (2003). *What is Language? A Study in the Philosophy of Linguistics*. University of Turku Press: Publications in General Linguistics 6.
- Jablonka, E. and M. Lamb (2005). *Evolution in Four Dimensions*. The MIT Press.
- Jackendoff, R. (1999). Possible stages in the evolution of the language capacity. *Trends in Cognitive Science* 3, 272-279.
- Jackendoff, R. (1990). *Semantic Structures*. MIT Press.
- Jackendoff, R. (1983). *Semantics and Cognition*. MIT Press.
- Johnstone, B. (2000). The Individual Voice in Language. *Annual Review of Anthropology*, 29, 405-24.
- Kenaan, H. (2002). Language, Philosophy and the Risk of Failure: Rereading the Debate Between Searle and Derrida. *Continental Philosophical Review* 35, 117-133.
- Kisch, S. (2008). 'Deaf Discourse': The Social Construction of Deafness in a Bedouin Community. *Medical Anthropology*, 27, 3, 283–313.
- Knight, C. (2007). Language co-evolved with the rule of law. *Mind & Society* DOI 10.1007/s11299-007-0039-1
- Knight, C. (2008). 'Honest Fakes' and Language Origins. *Journal of Consciousness Studies*, 15, 236-248.
- Knight, C. (1998). Ritual/Speech coevolution: A solution to the problem of deception. In Hurford, J. R., S. Kennedy and C. Knight (eds.), *Approaches to the Evolution of Language*. Cambridge University Press.
- Knight, C., M. Studdert-Kennedy and J. Hurford (eds.) (2000). *The Evolutionary Emergence of Language: Social function and the origins of linguistic form*. Cambridge University Press.
- Kovas Y., Hayiou-Thomas M. E., Oliver B., Dale P.S., Bishop D. V. M. and R. Plomin (2005). Genetic influences in different aspects of language development: The etiology of language skills in 4.5 year-old twins. *Child Development*, 76, 3, 632-651.
- Lakoff, G. and M. Johnson (1980). *Metaphors We Live By*. Chicago University Press.

- Langacker, R. (2005). Construction Grammars: Cognitive, Radical and Less So. In: Francisco J. Ruiz de Mendoza Ibanez and M. Sandra Pena Cervel (eds.) *Cognitive Linguistics: Internal Dynamics and Interdisciplinary Interaction*. Mouton de Gruyter.
- Larson, R. K., V. Deprez and H. Yamakido (eds.), (2010). *The Evolution of Human Language: Biolinguistic Perspectives*. Cambridge University Press.
- Levelt, W. (1993). *Lexical access in speech production*. Blackwell Publications.
- Levelt, W. (1989). *Speaking: From Intention to Articulation*. MIT Press.
- Levin, B. (1993). *English Verb Classes and Alternations. A preliminary investigation*. University of Chicago Press.
- Levin, B., and M. Rappaport-Hovav. (1995). *Unaccusativity*. MIT Press.
- Levinson, S. C. (2000). *Presumptive Meanings: The Theory of Generalized conversational Implicature*. Cambridge, MA: MIT Press.
- Levinson, S. C. (1997). From outer to inner space: Linguistic categories and nonlinguistic thinking. In Nuyts J. and E. Pederson (eds.), *Language and Conceptualization*. Cambridge University Press.
- Levinson, S. C. (1996). Relativity in spatial conception and description. In J. J. Gumperz, & S. C. Levinson (Eds.), *Rethinking linguistic relativity*. Cambridge University Press.
- Lewis, D. K. (1969). *Convention: A Philosophical Study*. Harvard University Press.
- Li, P., & Gleitman, L. (2002). Turning the tables: Spatial language and spatial cognition. *Cognition*, 83, 265-294.
- Lieberman, P. (2007). The Evolution of Human Speech: Its Anatomical and Neural Bases. *Current Anthropology* 48, 1, 39-66.
- Lieberman, P. (1991). *Uniquely Human: The Evolution of Speech, Thought and Selfless Behaviour*. Harvard University Press.
- Lucy, J. (1992a), *Grammatical Categories and Cognition: A Case Study of the Linguistic Relativity Hypothesis*. Cambridge University Press
- Lucy, J. (1992b), *Language Diversity and Thought: A Reformulation of the Linguistic Relativity Hypothesis*. Cambridge University Press.
- Lucy, J. and S. Gaskins (2001), Grammatical categories and the development of classification preferences: A comparative approach. In S. Levinson and M.

- Bowerman (eds.), *Language Acquisition and Conceptual Development*. Cambridge University Press.
- Meillet, A. (1921). *Linguistique Historique et Linguistique Generale*. Paris: H. Champion.
- Meir, I., M. Aronoff, W. Sandler and C. Padden (2010). Sign Languages and Compounding. In: Scalise, S. and I. Vogel (Eds.), *Compounding*. John Benjamins.
- Newmeyer, F. (2007). In Defense of the Autonomy of Syntax. Paper given at the *Scandinavian Dialect Syntax Meeting*.
- Newmeyer, F. (1998). *Language Form and Language Function*. MIT Press.
- Osherson, D. and E. Smith (1981). On the Adequacy of Prototype Theory as a Theory of Concepts. *Cognition* 9, 35-58.
- Piattelli-Palmarini, M. (1989). Evolution, selection and cognition: from 'learning' to parameter setting in biology and in the study of Language. *Cognition*, 31, 1-44.
- Pinker, S. (1994). *The Language Instinct: How the Mind Creates Language*. HarperPerennial.
- Pinker, S., and P. Bloom (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, 13, 707-726.
- Poulisse, N. (1999). *Slips of the tongue: speech errors in first and second language production*. John Benjamins.
- Prat, C. S., Keller, T. A. and M. A. Just (2007). Individual differences in sentence comprehension: A functional magnetic resonance imaging investigation of syntactic and lexical processing demands. *Journal of Cognitive Neuroscience* 19, 1950-1963.
- Prinz, J. and L. Barsalau (2000). Steering a Course for Embodied Representation. In: E. Dietrich and A. Markman (eds.) *Cognitive Dynamics: Conceptual Change in Humans and Machines*. MIT Press.
- Rich Harris, J. (1998). *The nurture assumption: Why children turn out the way they do*. The Free Press.
- Rosch, E. (1975). Cognitive Representations of Semantic Categories. *Journal of Experimental Psychology* 104, 3, 192-233.
- Ross, J. (1967). *Constraints on Variables in Syntax*. MIT dissertation.
- Sacks, H. (1995). *Lectures on Conversation*. Blackwell Publishing.
- Sandler, W. and D. Lillo-Martin (2006). *Sign Language and Linguistic Universals*. Cambridge University Press.

- Sapir, E. (1921). *Language: An Introduction to the Study of Speech*. Harcourt, Brace.
- de Saussure, F. (1916, 1966). *Course in General Linguistics*. McGraw-Hill Books.
- Savage-Rumbough, Sue, and R. Lewin (1994). *Kanzi: The Ape at the Brink of the Human Mind*. John Wiley.
- Schegloff, E.A. (1992). Repair after next turn: The last structurally provided defense of intersubjectivity in conversation. *American Journal of Sociology*, 98:1295-1345
- Schegloff, E.A., G. Jefferson, H. Sacks (1977). The preference for self correction in the organization of repair in conversation. *Language* 53: 361-82.
- Searle, J. (1977). Reiterating the Differences: A Reply to Derrida. *Glyph* 1
- Senghas, A. and M. Coppola (2001). Children creating language: How Nicaraguan Sign Language Acquired a Spatial Grammar. *Psychological Science*, 12, 4, 323-328
- Senghas, A., S. Kita and A. Ozyurek (2004). Children Creating Core Properties of Language: Evidence from an Emerging Sign Language in Nicaragua. *Science*, 305, 17.
- Senghas, R. J., A. Senghas and J. E. Pyers (2005). The Emergence of Nicaraguan Sign Language: Questions of Development, Acquisition and Evolution. In: Langer, J., S. T. Parker and C. Milbrath (eds.). *Biology and Knowledge Revisited: From Neurogenesis to Psychogenesis*. Erlbaum.
- Shapiro, L. (2010). *Embodied Cognition*. Routledge.
- Shen, Y. and D. Gil (to appear), Language, thought and the animacy hierarchy: Experimental studies of hybrids in Hebrew, Indonesian and Minangkabau. *Pacific Linguistics*.
- Slobin, D. (1996). From 'thought and language' to 'thinking for speaking'. In J. J. Gumperz, & S. C. Levinson (Eds.), *Rethinking linguistic relativity*. Cambridge University Press.
- Sperber, D. and D. Wilson (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Steels, L. (2001). Grounding symbols through evolutionary language games. In: Cangelosi A. and Parisi D. (Eds.), *Simulating the Evolution of Language*. Springer Publications.
- Szabolcsi, A. (2006). Strong vs. weak islands. In: M. Everaert and H. van Riemsdijk (eds.), *The Blackwell Companion to Syntax*, 4, 479–531. Blackwell Publishing.

- Traugott, Elizabeth Closs and B. Heine (eds), (1991). *Approaches to grammaticalization*. John Benjamins.
- Tomasello, M. (2009). *Why We Cooperate*. The MIT Press
- Tomasello, M. (2008). *Origins of Human Communication*. The MIT Press.
- Tomasello, M. (2003). *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Harvard University Press.
- Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., and Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5):675--691.
- West-Eberhad, M. J. (2003). *Developmental plasticity and evolution*. Oxford University Press
- Whorf, B (1956), *Language, Thought, and Reality: Selected Writings of Benjamin Lee Whorf*. The MIT Press.
- Wilson, D. and R. Carston (2006). Metaphor, relevance and the 'emergent property' issue. *Mind & Language* 21: 404-433
- Wittgenstein, L. (1953). *Philosophical Investigations*. The Macmillan Company.
- Zlatev, J. (2009). The Semiotic Hierarchy: Life, Consciousness, Signs and Language. In: *Cognitive Semiotics, Special issue on Anthroposemiotics vs. Biosemiotics*, 4.
- Zlatev, J. (2008). The Dependence of Language on Consciousness. *Journal of Consciousness Studies*, 15, 6, 34-62.
- Zlatev, J. (2007). Intersubjectivity, mimetic schemas and the emergence of language, *Intellectica* 2-3, 46-47, 123-152.
- Zlatev, J., T. Racine, C. Sinha, and E. Itkonen (2008), *The Shared Mind: Perspectives on Intersubjectivity*. John Benjamins.